

Characterising the effect of environment on galaxy evolution

by

Egidijus Kukstas

A thesis submitted in partial fulfillment
of the requirements of Liverpool John Moores University
for the degree of Doctor of Philosophy

April 2020

Declaration

The work presented in this thesis was carried out at the Astrophysics Research Institute, Liverpool John Moores University. Unless otherwise stated, it is the original work of the author.

While registered as a candidate for the degree of Doctor of Philosophy, for which submission is now made, the author has not been registered as a candidate for any other award. This thesis has not been submitted in whole, or in part, for any other degree.

Egidijus Kukstas
Astrophysics Research Institute
Liverpool John Moores University
IC2, Liverpool Science Park
146 Brownlow Hill
Liverpool L3 5RF
UK

APRIL 2020

Abstract

Galaxies are not the ‘island Universes’ they were once thought to be. Instead, they form a part of a larger structure called the ‘cosmic web’ which consists of dark matter, gas, and stars in order of decreasing fraction of the total mass budget. As part of it, galaxies are both influencing the cosmic web and are influenced by it. There is strong evidence demonstrating that galaxies located in dense environments (such as group or clusters) exhibit suppressed star-formation rates, red colours, early-type morphologies, and older stellar populations than their counterparts in the general ‘field’. This feature is often referred to as environmental galaxy quenching and, while there are many possible processes proposed as being responsible for this transformation, the detailed understanding of how it takes place is still lacking. In this thesis I propose a new method of characterising the link between galaxies and the environment responsible for quenching them, as well as improving on existing methods of studying galaxy transformation in dense environments. Cosmological, hydrodynamical simulations are used extensively to understand and interpret the results. Further work is necessary to make them capture the complex physics in groups and clusters - I highlight some of the limitations.

I begin by proposing a map-based method involving spatial cross-correlations between gas measures and a low-redshift galaxy survey ($z \leq 0.15$). This approach avoids the issues associated with membership assignment, and also directly links the underlying measure of environment and galaxy properties. I demonstrate that it can be applied to current observations from SDSS, Planck, and ROSAT surveys, yielding strong cross-correlation signals between gas pressure/density and galaxy density/quenched fraction. Hydrodynamical simulations, EAGLE and BAHAMAS, both reproduce the observed signal with some variation due to feedback implementations. The simulations can also be used to understand the measurements: I use BAHAMAS to demonstrate that most of the signal in Sunyaev-Zel’dovich effect – quenched fraction cross-correlation originates from

quenched satellites in groups and clusters. The same exercise shows that BAHAMAS over-quenches satellite galaxies.

Next, I investigate the performance of three hydrodynamical simulation codes (BAHAMAS, EAGLE AGNdT9, TNG300) at $z \sim 1$. There is evidence to suggest that quenching mechanisms may be different at this regime relative to the nearby Universe. Simulations do not currently capture all the necessary processes of quenching at $z \sim 0$. If the processes change between the two epochs, there is a possibility that simulations perform better at higher redshifts. I make several predictions of stellar content in haloes as well as quenched fraction from all three simulations in preparation for observational counterparts from the GOGREEN survey. There is great variation in the predicted relations, demonstrating that models used in simulations are relatively unconstrained in their current form. Comparing to data that is currently available indicates that none of the three simulations fully capture quenching of galaxies in dense environments at $z \sim 1$. Further observations will be able to inform future implementations of feedback for better agreement.

Galaxy membership assignment is, potentially, a big source of uncertainty in observations, especially at higher redshifts. Simple aperture combined with velocity cut methods are commonly employed in studies of group and cluster galaxies. I investigate the potential biases introduced from one such method, demonstrating that a relatively large number of contaminants is introduced. This severely affects the dominant type of galaxies and, subsequently, measured quenched fractions. I identify the main sources of contaminants and make suggestions on how to minimise them.

Publications

In the course of completing the work presented in this thesis, the following papers have been submitted for publication in a refereed journal:

‘Environment from cross-correlations: connecting hot gas and the quenching of galaxies’
Kukstas, Egidijus; McCarthy, Ian G.; Baldry, Ivan K.; Font, Andreea S., 2020, MNRAS, Volume 496, Issue 2, pp.2241–2261.

This paper forms the basis of Chapter 2.

‘The GOGREEN Survey: A deep stellar mass function of cluster galaxies at $1.0 < z < 1.4$ and the complex nature of satellite quenching’
van der Burg, Remco F. J.; Rudnick, Gregory; Balogh, Michael L.; Muzzin, Adam; Lidman, Chris; Old, Lyndsay J.; Shipley, Heath; Gilbank, David; McGee, Sean; Biviano, Andrea; Cerulo, Pierluigi; Chan, Jeffrey C. C.; Cooper, Michael; De Lucia, Gabriella; Demarco, Ricardo; Forrest, Ben; Gwyn, Stephen; Jablonka, Pascale; **Kukstas, Egidijus**..., 2020, *Astronomy and Astrophysics*, Volume 638, Article A112.

‘The GOGREEN survey: the environmental dependence of the star-forming galaxy main sequence at $1.0 < z < 1.5$ ’

Old, Lyndsay J.; Balogh, Michael L.; van der Burg, Remco F. J.; Biviano, Andrea; Yee, Howard K. C.; Pintos-Castro, Irene; Webb, Kristi; Muzzin, Adam; Rudnick, Gregory; Vulcani, Benedetta; Poggianti, Bianca; Cooper, Michael; Zaritsky, Dennis; Cerulo, Pierluigi; Wilson, Gillian; Chan, Jeffrey C. C.; Lidman, Chris; McGee, Sean; Demarco, Ricardo; Forrest, Ben De Lucia, Gabriella; Gilbank, David; **Kukstas, Egidijus**..., 2020, MNRAS, Volume 493, Issue 4, pp.5987-6000

Acknowledgements

I would like to begin by thanking my supervisor, Ian McCarthy. First, for proposing such a unique and engaging project, then providing the support and guidance to follow it through. I appreciate the freedom you gave me to do my own thing, then providing a helping hand when I, eventually, found myself lost.

I must also thank: my second supervisor, Andreea Font, all of ex-gal group, and the GOGREEN collaboration for useful discussions, criticisms, and advice.

I am also thankful to everyone at the ARI for making it such a welcoming and inclusive environment. I came here on a recommendation that it's "a good place to be" and was not disappointed one bit. Everyone, from students to post-docs and staff members has contributed to what it is.

I would like to thank my parents for teaching me the value of education and enabling me to follow the career path I enjoy. Thank you to my aunt, Aurelija, for always telling me that I have what it takes to achieve my goals.

Lastly, I would like to thank my lovely Sophie for putting up with me during my grumpy moments. I hope some of your positivity has rubbed off on me.

Contents

Declaration	ii
Abstract	iii
Publications	v
Acknowledgements	vi
List of Figures	x
List of Tables	xii
1 Introduction	1
1.1 Galaxy distribution in the Universe	2
1.2 Contents of overdensities	4
1.3 Classifying the cosmic web environment	5
1.3.1 Characterising the full cosmic web	5
1.3.2 Quantifying galaxy environment	6
1.4 Galaxy transformation	9
1.4.1 Distribution functions	10
1.4.2 Physical processes affecting galaxy evolution	12
1.5 Environmental influence on galaxy evolution	13
1.5.1 Observational evidence	14
1.5.2 Environmental physics and processes	16
1.6 Numerical simulations	18
1.6.1 N-body and semi-analytic models	18
1.6.2 Hydrodynamical simulations	21
1.7 This thesis	26
2 Environment from cross-correlations: connecting hot gas and the quenching of galaxies	28
2.1 Introduction	28
2.2 Data and Map-making	33
2.2.1 Galaxy catalogues	33
2.2.1.1 SDSS DR7 spectroscopic sample	33
2.2.1.2 SDSS DR12 photometric sample	35
2.2.2 HEALPix map-making using galaxy catalogues	42
2.2.2.1 Galaxy density	43

2.2.2.2	Quenched fraction	44
2.2.3	Masking	44
2.2.4	Thermal Sunyaev-Zel'dovich effect map	46
2.2.5	X-ray map	47
2.2.6	Simulations	48
2.2.6.1	EAGLE	48
2.2.6.2	BAHAMAS	49
2.2.6.3	Light cones and simulated map making	50
2.3	Auto- and Cross-power spectra estimation	52
2.3.1	Formalism	52
2.3.2	Null-tests and error estimation	55
2.4	Results	56
2.4.1	Galaxy overdensity power spectrum	57
2.4.2	Quenched fraction power spectrum	60
2.4.3	Quenched fraction-galaxy overdensity cross-spectrum	61
2.4.4	tSZ effect-galaxy overdensity cross-spectrum	62
2.4.5	X-ray-galaxy overdensity cross-spectrum	63
2.4.6	tSZ effect-quenched fraction cross-spectrum	65
2.4.7	X-ray-quenched fraction cross-spectrum	66
2.5	Discussion: isolating external from internal quenching	66
2.6	Summary and Conclusions	69
3	Comparison of simulated cluster galaxies at $z \sim 1$	72
3.1	Introduction	72
3.2	Simulations	74
3.2.1	BAHAMAS / MACSIS	75
3.2.2	EAGLE / Hydrangea	76
3.2.3	Illustris TNG300	77
3.3	OBSERVATIONS	78
3.4	RESULTS	79
3.4.1	Stellar mass content of haloes	80
3.4.2	Galaxy distribution as a function of radius	83
3.4.3	Galaxy stellar mass function: field vs. clusters	86
3.4.4	Halo quenched fraction	92
3.4.5	Quenched fraction profile	94
3.5	Summary and Conclusions	94
4	Testing the accuracy of observational galaxy selection with simulations	97
4.1	Introduction	97
4.2	Simulation description	99
4.3	Computing observable quantities	100
4.3.1	Observed redshift	100
4.3.2	Cluster velocity dispersion	101
4.4	Results	102
4.4.1	Phase-space diagrams	102
4.4.2	Effects on GSMF and quenched fraction	105
4.4.3	Radial profiles	108

4.5	Summary and conclusions	108
5	Summary, conclusions, and future work	112
A		115
A.1	SPH smoothing and shot noise	115
A.2	Contamination from AGN	118
B		122
B.1	Choice of star-forming–quiescent division	122
	Bibliography	125

List of Figures

1.1	A slice of the cosmic web	3
1.2	Examples of galaxy bimodality observed in colour-magnitude space and star formation-stellar mass.	11
1.3	Galaxy colour correlation with environment and mass, as well as satellite quenched fraction correlation with host halo mass.	15
2.1	Mass-to-light ratio – stellar mass relation used to find volume completion limits.	34
2.2	A comparison of parametrically-derived stellar masses and their catalogue values.	37
2.3	Galaxy stellar mass functions for spectroscopic and photometric samples.	38
2.4	An example of the colour–stellar mass relations used to calibrate the quenched–star-forming division for the SDSS photometric sample.	40
2.5	Quenched fraction - stellar mass relations for observational samples and simulations.	41
2.6	Maps used for cross-correlations.	45
2.7	Auto- and cross-power spectra for spectroscopic and photometric galaxy samples.	58
2.8	Cross-power spectra between galaxy properties and hot gas measures.	64
2.9	A test of the sensitivity of the method using simulated galaxies.	68
3.1	Halo mass distributions for GOGREEN and all the simulations used here.	79
3.2	Stellar mass content of haloes for all central galaxies.	81
3.3	Stellar mass content of simulated haloes in the cluster sample.	83
3.4	Galaxy number and stellar mass density radial profiles split by type.	86
3.5	Galaxy stellar mass functions split by type: comparison to GOGREEN.	91
3.6	Total halo quenched fraction and quenched fraction excess as a function of stellar mass.	93
3.7	Quenched fraction radial profiles.	95
4.1	Phase-space diagrams for three different galaxy membership selections in bins of halo mass.	103
4.2	Cluster GSMF measurements for observational and FoF selections, split by galaxy type.	106
4.3	Galaxy number density and quenched fraction radial distribution in clusters.	109
A.1	Distributions of SPH smoothing parameter, h , for spectroscopic and photometric samples; observations and simulations.	117

A.2	A demonstration of SPH smoothing and shot noise effects on the power spectra for the photometric galaxy sample.	119
A.3	$X - \check{N}_{\text{tot}}$ cross-power spectra computed using different masks for AGN sources in the RASS x-ray map.	121
B.1	Field galaxy distributions demonstrating the choice of division between star-forming and quenched galaxies.	123
B.2	Cluster galaxy distributions demonstrating the choice of division between star-forming and quenched galaxies.	124

List of Tables

2.1	A summary of galaxy sample properties for the spectroscopic and photometric SDSS volume-limited samples.	42
3.1	A summary of simulation properties used.	78

Chapter 1

Introduction

Our home planet, the Earth, is the third closest planet orbiting around a fairly average star we call the Sun. The Sun lies about two-thirds of the way out in the Orion arm of the Milky Way, a collection of billions of stars, gas, and dust arranged in a disc-like structure with spiral patterns of density. Our home galaxy, the Milky Way, is one of many galaxies in close proximity of one another that make up a structure called the Local Group. The Local Group is connected to other such structures with bridges of galaxies, and those other structures are connected to neighbouring structures of their own, forming a three-dimensional web of voids, sheets, filaments, groups, and clusters - the cosmic web. It has taken many observations and a lot of theoretical work to arrive at this description of the Universe we use today. Ongoing work is constantly refining this description and adding new knowledge.

Galaxies, rather than being randomly distributed in space, occupy different environments in the cosmic web which, in addition to the normal galaxy evolutionary path, influences the way a galaxy evolves over time. A galaxy which happens to occupy a dense environment is, more often than not, observed to be ‘early type’; that is, it is more likely to be elliptical in shape with primarily old, red stars in it. A galaxy like our Milky Way, occupying a less dense environment, is more likely to be ‘late type’, i.e. blue disc or spiral, with primarily young populations of stars. These relations have been known for the best part of a century but there is still no complete description of how they come about. Which aspect(s) of environment is responsible for this transformation? How long does it take? Does the dominant process change over time? These are all

questions still yet to be answered. This thesis aims to contribute to the understanding of environmental processes by proposing a new way of measuring its influence as well as using state-of-the-art numerical simulations to interpret the observed trends.

In the following sections I attempt to summarise what we know about the galaxy distribution in the cosmic web and the influence it has on galaxy evolution, as well as ways this influence is studied. In this introductory chapter, I begin by introducing the concept of cosmic web and its composition in Sections 1.1 and 1.2. I then give an overview of how it is often classified using current algorithms in Sec. 1.3. Having placed galaxies in the context of a cosmic web, I then focus on general galaxy evolution in Sec. 1.4, describing the observed distributions and scaling relations. In Sec. 1.5, I describe galaxy transformation in the context of environment, presenting observational evidence for it and processes which may be able to explain it. Finally, in Sec. 1.6 I introduce and describe numerical simulations as a tool for studying processes responsible for galaxy evolution, especially in dense environments.

1.1 Galaxy distribution in the Universe

Currently, the most accepted description of the Universe is the Λ CDM model which describes the Universe dominated by two components: a cosmological constant, ‘ Λ ’, driving the accelerating expansion (Peebles & Ratra, 2003) and ‘Cold Dark Matter’ (White & Rees, 1978) which dominates the mass budget. Within this description, small perturbations in density at the early stages of the Universe grow and collapse to form ever larger structures, a process referred to as hierarchical growth. This hierarchical growth leads to the development of the cosmic web. This is best visualised by taking a slice from an N-body simulation box and introducing a colour gradient to particle density. One such image is shown in Figure 1.1, which shows the dark-matter density field; black colours represent low-density regions called ‘voids’ and yellow shows the densest parts - ‘cluster’ centres. Baryonic matter traces this density field; gas flows towards towards the minima in the gravitational field, which compresses the gas and enhances its cooling rate. If sufficiently compressed, the gas can collapse and form stars within galaxies.

Support for this picture of structure assembly has been around for many decades now. Edwin Hubble successfully measured distances to ‘nebulae’ to establish that they were




FIGURE 1.1: A slice of the cosmic web, centred on a massive cluster, as seen in dark matter. Colour gradient indicates density with purple showing the less dense regions and yellow showing the overdensities. Image credit: Virgo collaboration, the Millenium project.

not of local origin (Hubble, 1925). Not long after, Shapley & Ames (1926) found that ‘clouds’ of galaxies in the constellations of *Virgo* and *Coma Berenices* had similar apparent magnitudes and angular diameters. This led to the conclusion that they are physically close to one another in what are today known to be nearest clusters of galaxies. The list of known clusters has since been greatly expanded, most notably with the Abell catalogue (Abell, 1958; Abell et al., 1989). The field of structure mapping has been revolutionised with large-area, reasonably deep galaxy surveys such as the 2 degree Field Galaxy Redshift Survey (Colless et al., 2001), Sloan Digital Sky Survey (SDSS, York et al. 2000), and the Dark Energy Survey (DES, Abbott et al. 2018) which have mapped the full structure and not just the most dense parts. Our desire to classify things has led to the definition of five major groups in the order of increasing density and, therefore, galaxy abundance: voids, sheets, filaments, groups, and clusters.

1.2 Contents of overdensities

The extent of our knowledge about each component of the cosmic web is heavily biased by the ease or difficulty of observing it. Early observations started with ‘nebulae’, confirming that they were in fact other galaxies outside the Milky Way because they emitted a substantial amount of light at optical wavelengths and it so happened to be what instruments at the time were sensitive to. The same very much applies to today’s observations: galaxies are very much the dominant tracer of total matter despite contributing relatively very little to the total mass of it.

It was first shown by Zwicky (1933) and Smith (1936) for Coma and Virgo clusters, respectively, that galaxies only make up a small fraction of the total mass and that there is far more matter that we cannot see. More than 80% of the mass budget is taken up by ‘dark matter’, which does not interact with either light nor other forms of matter except through gravity. Contributions from numerical simulations have shown that dark matter is relatively smooth and space filling, taking up a smoothly varying profile at peaks in the density field. Galaxies, therefore, occupy haloes of dark matter and a group or cluster is a large halo with several galaxies embedded in it rather than a collection of galaxies. Haloes were then shown to be self-similar and follow a common two-parameter density profile (the ‘NFW’ profile, Navarro et al. 1995).

Observationally, it is not possible to make direct astronomical detections of dark matter because it has no observable emission of its own. The best indirect observation is the gravitational effect it has on surrounding matter - it is possible to infer the distribution of dark matter by the way it shapes spacetime gravitationally either through strong lensing (Soucail et al., 1987) or weak lensing and cosmic shear (Tyson et al., 1990).

The bulk of baryonic matter is in the form of ionised plasma that traces the cosmic web. Roughly 40 to 50% (Bykov et al., 2008) of this gas is thought to be in the form of warm-hot ($10^5 - 10^6$ K) intergalactic medium (WHIM) contained within filaments based on numerical simulations (Cen & Ostriker, 1999; Davé et al., 2001). This diffuse material held in thin filaments is extremely difficult to observe directly, but recent attempts suggest that it does indeed exist (Nicastro et al., 2008; Tanimura et al., 2019; Umehata et al., 2019). The other part of ionised gas is held within haloes and, as a result of the increased gravitational forces, has temperatures of $10^7 - 10^8$ K (Voit et al., 2005).

This intra group-/cluster-medium (IGM/ICM¹) is sufficiently hot that it has strong X-ray emission through the bremsstrahlung process and X-ray emission lines (Sparke & Gallagher, 2007). It can also be detected through its effect on Cosmic Microwave Background (CMB) photons as they stream towards us. Low energy photons can be inverse-Compton scattered by relativistic electrons in the ICM plasma, which result in temperature distortions in the measured angular power spectrum of the CMB. This feature is called the thermal Sunyaev-Zel'dovich (tSZ) effect and carries information about the line-of-sight pressure of the ICM gas (Staniszewski et al., 2009). This type of observation will be used in Chapter 2. The last baryonic component is the still cooler gas which results from ICM and WHIM cooling due to their emission. This leads to gas condensing into clumps (Simionescu et al., 2011), eventually either forming stars or being deposited onto galaxies and cluster centres (Fabian, 1994). The conversion of relatively cool gas into stars is a vast field in itself and not a subject of this thesis, but it is fascinating to remark that less than 15% of all baryonic matter is cool, and baryonic matter makes up less than 5% of the total contents of the Universe, yet stars and cool gas have dominated the focus of astronomy for centuries. Dark matter and hot gas may not make up galaxies but their influence on galaxy formation and evolution is profound.

1.3 Classifying the cosmic web environment

Identifying and appreciating the complex nature of the cosmic web is important but, if one is interested in its effect on galaxies, then the need to simplify and encapsulate the properties in simple parameters becomes somewhat of a necessity. If one was to choose a galaxy and ask what the major sources of influence on its evolution are, determining whether it lives in a crowded cluster or sparse sheet may be sufficient to describe the environmental influence.

1.3.1 Characterising the full cosmic web

While the bulk of cosmic web is made up of dark matter that we cannot see, galaxies trace out this structure and reflect the underlying density field. Studying the best methods of estimating cosmic field structure from galaxies is a field in itself, with an

¹for clarity, 'ICM' will now refer to hot gas in both groups and clusters, despite their differences in halo mass.

industry of algorithms striving to provide the most accurate results (see Libeskind et al. 2018 for a review). Some examples of such algorithms include: DIScrete PERsistent Structures Extractor (DisPerSe, Sousbie 2011; Sousbie et al. 2011), a formalism designed to take advantage of the topological structure of the cosmic mass distribution and is particularly good for finding filaments in simulations, and the Multiscale Morphology Filter/Nexus (Aragón-Calvo et al., 2007, 2010) framework, which works by classifying matter distribution on the basis of local variations in the density, velocity or gravity field encoded in the Hessian matrix. It has been successfully applied to SDSS data to find filaments (Jones et al., 2010). NEXUS+ (Cautun et al., 2013) extended this method to include finding of knots (i.e. clusters).

1.3.2 Quantifying galaxy environment

While exploring the influence of filamentary environment on galaxies is becoming possible with latest spectroscopic data, and will most certainly become more viable in the future, the most influential environments are, without doubt, groups and clusters. Knots in the cosmic web are hosts to the highest environmental densities of galaxies, gas, and dark matter which makes them (i) easier to locate from galaxy data, (ii) their effect is subjected to the highest number of galaxies, and (iii) environmental processes are strongest. Methods used to find the densest environments are only concerned about finding the dense dark matter haloes and so the picture of the Universe they present is substantially simpler. This Universe is composed of haloes that can host a single galaxy, anywhere from several to tens in the case of groups, and potentially thousands in the case of clusters and superclusters. The exact lines of separation are somewhat arbitrary, there is no clear separation between them, it simply helps to categorise structures according to their size and mass when studying their effect on galaxies. The following few methods are concerned with finding such haloes.

The ‘Friends-of-Friends’ (FoF) percolation algorithm (Huchra & Geller, 1982; Davis et al., 1985; Eke et al., 2004) was specifically designed to find overdensities of arbitrary shape based on the separation distance of constituent members. These members can be galaxies or, in the case of numerical simulations, dark matter particles. The process involves connecting constituents with a ‘linking length’ of chosen size in either two or three dimensions. Constituents which lie within the linking length of the chosen starting

particle/galaxy are said to be members and belong to the same halo. This method is relatively simple and works rather well provided that an appropriate linking length is chosen. A linking length that is too long will group distinct objects into one and a linking length that is too short will lead to fragmentation of haloes. FoF algorithms are very commonly used in numerical simulations where there is a need to find haloes which will host galaxies from a volume of loose dark matter particles.

Another concept commonly employed in simulations is the ‘spherical overdensity’ that involves defining a sphere centred on the region of interest (such as a central galaxy). All particles/galaxies within the sphere are then said to be members of the group. The radius, r , is commonly chosen so that the sphere encloses some multiple, Δ , of the mean or critical density of the Universe. To make the choice physically motivated, Δ is often chosen to be such that the radius of this sphere is equal to the virial radius of the halo (Peebles, 1980), i.e. $\Delta_c = 200$ (Gunn & Gott, 1972; Lahav et al., 1991; Lacey & Cole, 1993) and $r_{200c} = r_{\text{vir}}$. r_{200c} and M_{200c} (mass enclosed within a sphere whose density is 200 times the critical density of the Universe) are commonly used in observational studies as well as theoretical.

FoF and spherical overdensity algorithms work very well on numerical simulation data when the number of particles is very large and their full three-dimensional positions are available. This is not the case for magnitude-limited galaxy redshift surveys. The transverse and line-of-sight linking lengths are no longer equal, complicating the process. The resulting membership often suffers from lack of completeness of group members, interlopers, group fragmentation, and group blending. Mock catalogues from simulations are often used to inform the choice of linking lengths for observations (Robotham et al., 2011), but this only alleviates one of several issues. An improved method which builds on the standard FoF approach is often referred to as the ‘halo finding algorithm’. The end result is a ‘halo catalogue’: a classification of galaxies into either centrals (of a group or their own halo) or satellites. The method starts with (i) the traditional FoF algorithm in order to find the centres of potential groups and to assign initial membership; (ii) the total luminosity of any given group is estimated including the magnitude limit of the survey being used; (iii) by assuming a model mass-to-light ratio for the group, the total mass, radius, and virial radius can be estimated; (iv) having defined the extent of any given group, a probability can be assigned to catalogue galaxies of being members of their nearest group; (v) new group centre is recomputed using these members and the

process is repeated until stable centres are established. The most well-known example of this method is the Yang et al. (2005, 2007) SDSS DR7 (Abazajian et al., 2009) halo catalogue. Other catalogues exist for different surveys with variations of the method outlined here.

The choice of an environmental measure is often dictated by the quality and size of data as well as the goals of the study. In many cases it is not necessary to find which galaxy is a central, a satellite, or what halo mass they occupy. When all one wants to know is “how big of an overdensity does a galaxy occupy?”, there are other estimators which may be more applicable. These work on the assumption that galaxy density reflects the underlying dark matter and gas density. The task then becomes to estimate how close galaxies are to one another (see Muldrew et al. 2012 for a review). There are two basic classes of such environment estimators: those computing the distance to N nearest neighbours and those computing the number of galaxies within a given volume. The former relies on a suitable choice of neighbour galaxies, N , and an estimate of distance to N^{th} nearest neighbour. The shorter the distance is the more densely galaxies are located to one another. This can be done either in projected space, with an introduction of a line-of-sight velocity limit to minimise projection effects, or three-dimensional space if accurate redshifts are known. The choice of N typically varies between 1 and 10 (Baldry et al., 2006; Brough et al., 2011), depending on the type of structures of interest. Variations of this basic method exist which attempt to reduce the number of interlopers (Marinoni et al., 2002; Cooper et al., 2005; Li et al., 2011). The second type and perhaps the simplest of all types of density estimates is the ‘fixed aperture’ method. This involves a choice of aperture in the transverse direction with a combination of velocity limits in the line-of-sight direction resulting in a cylindrical volume. The number of galaxies within this volume can then be counted and density estimated. Such measurements are often expressed as a density contrast with respect to the mean density that would be expected if galaxies were distributed randomly, rather than just raw density counts. This takes the form of:

$$\delta \equiv \frac{\delta\rho}{\rho} = \frac{N_q - \bar{N}_g}{\bar{N}_g}, \quad (1.1)$$

where N_g is the number of galaxies measured within the aperture, and \bar{N}_g is the mean number expected if they were distributed randomly. Size of aperture and velocity limits

vary heavily depending on the objects studied and quality of data (spectroscopic vs. photometric redshift estimates). This technique has been successfully used in numerous studies including Kauffmann et al. (2004), Croton et al. (2005), and Peng et al. (2010).

Another type of nearest-neighbour estimator (and quite relevant for this thesis) is the smoothing kernel (Monaghan & Lattanzio, 1985) approach commonly used in Smoothed Particle Hydrodynamic (SPH) simulations, but has been successfully applied to observational data (Park et al., 2007). In this approach the underlying galaxy field is smoothed with a function whose size varies depending on the distance h , distance to the N^{th} nearest neighbour. The smoothing function is then defined as follows:

$$W(r, h) = \frac{8}{\pi h^3} \begin{cases} 1 - 6 \left(\frac{r}{h}\right)^2 + 6 \left(\frac{r}{h}\right)^3, & 0 \leq \frac{r}{h} \leq \frac{1}{2}, \\ 2 \left(1 - \frac{r}{h}\right)^3, & \frac{1}{2} < \frac{r}{h} \leq 1, \\ 0, & \frac{r}{h} > 1, \end{cases} \quad (1.2)$$

where r is the separation between galaxies. The output of this process is either a smoothed density field of galaxies or a weighting to be applied to each galaxy. The main advantages of using SPH smoothing are that it implicitly conserves the quantity being smoothed (total mass or number of galaxies) and that it does not ‘over-smooth’ dense regions of interest when compared to other methods.

1.4 Galaxy transformation

Having discussed the idea of a cosmic web as a framework in which galaxies are embedded, let us now look at it from the perspective of galaxies. A galaxy is strongly affected by its immediate environment; it relies on fresh gas flows and its gas reservoir staying intact in order to form stars. In the following sections I will describe how galaxy properties change as a result of varying star formation.

Galaxies do not all look the same. This fact emerged very early on in galactic astronomy. With time and increasing number of observations, a visual classification based on optical morphology emerged: the Hubble tuning fork (Hubble, 1926). This scheme classified galaxies into two main categories: spheroidal ‘early type’ and spiral disc, ‘late type’. Late type galaxies can further be split up into those which contain a bar and ones

which do not, the exact reason for this is still an active field of research. It must be noted that Hubble's terminology does not imply temporal evolution (as explained by Baldry 2008), but simply classifies galaxies based on complexity in their appearance on photographic plates. This classification mirrors that used for stellar spectra, which also classifies objects based on visual complexity (due to absorption lines in stellar spectra, in this case). Since then, many other patterns in galaxy populations have been observed supporting the idea that galaxies evolve from late type to early type but follow many possible pathways. Multiwavelength imaging surveys together with either photometric or spectroscopic redshift estimates make rest-frame luminosity and colour measurements from ultraviolet to far-infrared possible. Such spectral energy distribution estimates allow for the possibility of fitting models of galaxy star-formation histories in combination with stellar-population models (Walcher et al., 2011; Conroy, 2013). This procedure can yield accurate stellar masses, star formation rates, and stellar population ages which allows for galaxy classification in other quantities (see Somerville & Davé 2015 for a more in-depth review).

1.4.1 Distribution functions

The co-moving number density of galaxies as a function of either stellar mass or luminosity have long been used to characterise galaxy evolution, especially over cosmic time. They often take the characteristic shape of a Schechter function (Schechter, 1976) which can be described by a normalisation, turnover, and an asymptotic slope towards lower stellar masses/luminosities.

In galaxy formation models, it is accepted that the number of galaxies on the low-mass end is controlled by star-formation feedback and the high mass end is determined by feedback from supermassive black holes (see Blanton & Moustakas 2009 for a review). Galaxy luminosity and stellar mass functions have been measured out to redshifts of $z \sim 4$. These measurements paint the following picture: (i) galaxies appear to be building up their mass over cosmic time, in line with predictions from hierarchical structure formation (Madau & Dickinson, 2014); (ii) the number density of massive galaxies (those above the turnover point in the Schechter function) increases rapidly from $z \sim 4 - 2$ but is then constant or decreases slightly from $z \sim 2 - 0$, suggesting that massive galaxies assembled in the early stages of the Universe and are no longer doing so (Marchesini

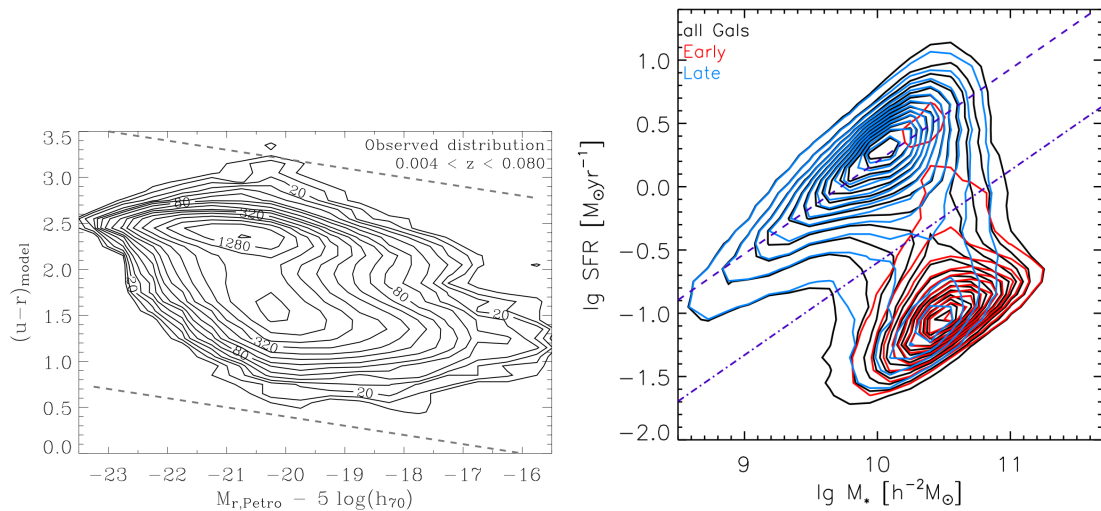


FIGURE 1.2: **Left:** galaxy colour bimodality using SDSS galaxies; contours indicate number density. *Reproduced from Baldry et al. (2004).*

Right: bimodality in the $\log_{10}(\text{SFR})$ - $\log_{10}(M_{*})$ space; contours indicate number density: black for all galaxies, blue for late-type, red for early-type. Dashed line is a fit to the star-forming sequence and dash-dotted line demonstrates that the same line scaled down divides the two populations relatively cleanly. *Reproduced from Liu et al. (2019).*

et al., 2009; Muzzin et al., 2013); (iii) number density of low-mass galaxies, on the other hand, increases more rapidly at $z \sim 2 - 1$ showing that they assembled later and more slowly (Cimatti et al., 2006).

Another important feature is the ‘galaxy bimodality’ which shows up in the distributions of several different quantities. It was first observed in the colour-luminosity relation (Baldry et al., 2004), where galaxies are observed to occupy either the ‘red sequence’ or ‘blue cloud’ (in optical colours). A similar distribution is observed in the specific star formation rate (sSFR) - stellar mass ($\log_{10}(M_{*})$) distribution (Kauffmann et al., 2003; Brinchmann et al., 2004) as well as just SFR- $\log_{10}(M_{*})$ (Liu et al., 2019). These two distributions are shown in Figure 1.2. It is possible to divide the two populations on any plane with a straight line into ‘star-forming’ and ‘quiescent’² populations. In particular, the right hand side figure shows that star-forming galaxies are more likely to be late-type, while the quiescent galaxies are nearly always early type. The same patterns apply to colour.

²throughout this thesis I use ‘quiescent’ and ‘quenched’ interchangeably to mean galaxies which are no longer actively forming stars, irrespective of the causes.

Deep surveys have confirmed that galaxy bimodality persists out to at least $z \sim 2$ (Brammer et al., 2011; Muzzin et al., 2013). Co-moving number and mass densities of individual populations have also been measured. Interestingly, the number/mass density of quiescent galaxies is observed to be *increasing*, whereas that of star-forming galaxies is *decreasing* (Bell et al., 2004, 2007; Faber et al., 2007; Brammer et al., 2011; Muzzin et al., 2013). This is contrary to what one might naively assume, that the number of star-forming galaxies is rising because new galaxies are forming from intergalactic gas and they are constantly forming new stars. The contrary is true: increasing number of star-forming galaxies are having their star formation ‘quenched’ as cosmic time progresses.

1.4.2 Physical processes affecting galaxy evolution

Galaxies grow by converting hydrogen gas from the cosmic web into stars. When the intergalactic medium collapses onto an overdensity, it is heated to temperatures of $T > 10^7 \text{K}$. It is fully ionised at such temperature and only cools through two-body interactions, such as bremsstrahlung. At temperatures in the range of $10^4 < T < 10^7 \text{K}$, electrons can recombine with atoms. Below 10^4K , cooling can occur through collisional excitation/de-excitation of atoms via metal line emission. Following this collapse, the gas settles into a pressure-supported gaseous halo. As this gas cools further, it supplies fresh fuel for star formation onto the disc of a galaxy (Kereš et al., 2005). Any process which can inject additional heat or disrupt this process of continuous gas flows will limit future star formation. There are a number of such processes:

(i) Star-formation feedback. Early galaxy formation models already acknowledged that uncontrolled gas collapse would lead to most of it already being concentrated at the centres of overdensities, contrary to observations (White & Rees, 1978; Dekel & Silk, 1986; White & Frenk, 1991). Energy generation from dense regions is necessary to reduce the efficiency of gas cooling and uncontrolled collapse. Feedback from stars is the ideal solution to this problem; their energy is injected into the surrounding gas through a number of mechanisms (e.g., photoheating, photoionisation, winds) and even thought to eject gas out of the galaxy (Hopkins et al., 2012).

(ii) Black hole-driven feedback. There is strong evidence to suggest that most (certainly those above the mass of the Milky Way) galaxies contain a supermassive black hole

(SMBH) at their centre (see a review by Kormendy & Ho 2013). The amount of potential energy that had to be released when growing these black holes exceeds the binding energy of galaxies they are hosted by. This energy had to be released in some form. There are many proposed physical mechanisms whereby the large amounts of energy and momentum originating from the SMBH can couple to the surrounding gas and galaxies, possibly regulating the growth of the black hole itself and surrounding galaxies. Active Galactic Nuclei (AGN) feedback processes are often divided into two classes: radiative ‘quasar’ mode and jet-driven ‘radio’ mode. The former is associated with strong heating of gas, winds capable of ejecting it, and radiation able to ionise and photodissociate the gas. The latter is characterised by highly-collimated radio jets, X-ray emitting bubbles, weak shocks, and sound waves. Indeed, AGN feedback has been the solution to reproducing the observed GSMF in hydrodynamic simulations. The amount of energy released by AGNs is certainly enough to eject/heat the surrounding gas but questions remain regarding how well winds/radiation can couple to it (see reviews by Fabian 2012; Heckman & Best 2014; Somerville & Davé 2015).

1.5 Environmental influence on galaxy evolution

The above processes can be classed as resulting from the ‘internal’ mechanisms of galaxy evolution, i.e. those acting from within the galaxy. Quite often, they are observed to scale with the galaxy’s stellar mass. For this reason, they are collectively referred to as ‘mass quenching’ processes (Peng et al., 2010). This trend can be seen quite clearly in the upper panel of Figure 1.3; at mean cosmic density (zero on the y-axis, which is overdensity relative to the mean), galaxy colours change from being completely blue dominated at low stellar masses to being fully red at masses above $\log(M_*) = 11.0$. Mass quenching is often suggested as an explanation for the cessation of massive galaxy growth at redshifts above $z > 2$ described in Section 1.4.1. In fact, there is growing evidence that mass and environmental quenching are dominant at different stages of universal galaxy evolution and stellar masses (Peng et al., 2012; Muzzin et al., 2012; Darvish et al., 2016; Balogh et al., 2016; Papovich et al., 2018; Pintos-Castro et al., 2019). Measures are often taken to attempt to separate one effect when studying the other. In the following, I discuss processes and mechanisms most commonly invoked in order to explain galaxy transformation in dense environments.

1.5.1 Observational evidence

The first observations of an environmental effect on galaxy colour and morphology was based on small photometric-only samples (Oemler, 1974; Davis & Geller, 1976; Dressler, 1980; Postman & Geller, 1984). As a result, only qualitative conclusions could be made. With the advent of large, spectroscopic galaxy surveys such as SDSS, numerous works came out quantifying these effects in substantial detail (Balogh et al., 2004; Hogg et al., 2004; Kauffmann et al., 2004; Blanton et al., 2005; Peng et al., 2010). Since then, some theoretical works attempting to explain the observations have also been published (Dekel & Birnboim, 2006; McGee et al., 2009; Wetzel et al., 2012, 2013; Bahé et al., 2013; McGee et al., 2014; Bahé & McCarthy, 2015). This combined effort has arrived at several conclusions: (i) Environment plays a major role in galaxy (especially low-mass) transformation in the local Universe (Peng et al., 2012). (ii) Environmental impact on SFR (and, by extension, colour) is stronger than on morphological transformation (Kauffmann et al., 2004; Blanton et al., 2005; Bamford et al., 2009). At fixed morphology, there is still substantial variation in SFR with environment, but the inverse is not true. (iii) SFR, colour, and morphology depend on *small* scale environment ($< 1\text{Mpc}$), with little influence beyond these scales (Hogg et al., 2004; Kauffmann et al., 2004; Blanton et al., 2005). (iv) In the halo formalism, it has been found that colour and star formation history depend heavily on the properties of the host halo (Bamford et al., 2009; Blanton & Berlind, 2007; Wilman et al., 2010). This same fact is illustrated in the bottom panel of Figure 1.3, where there is a clear scaling of satellite quenched fraction with host halo mass. (v) Environmental transformation (in colour) happens on timescales of $< 1\text{Gyr}$. This is supported by the observations of relatively few galaxies occupying the parameter space in between blue/red, star-forming/quiescent populations. In addition, variations in properties between samples of the same environment are very small (Balogh et al., 2004; Wetzel et al., 2012). If timescales were longer, we would see galaxies at different stages of the transformation. Environmental quenching is not instantaneous, however, because there *are* some galaxies in the green valley (Wetzel et al., 2013). The field remains an active area of research, however.




FIGURE 1.3: **Top:** galaxy colour distribution in the overdensity-stellar mass plane. Blue galaxies are predominantly low-mass and are found in under-dense environments. Red galaxies are either high-mass or occupy dense environments. *Reproduced from Peng et al. (2010).* **Bottom:** Quenched fraction as a function of stellar mass. Centrals show a monotonic increase with stellar mass, whereas satellites show an additional quenching feature which scales with host halo mass. *Reproduced from Wetzel et al. (2012).*

1.5.2 Environmental physics and processes

Environment clearly plays a major role in galaxy evolution in the local Universe. While the response of galaxies has been well measured, the details of *how* dense environment acts to suppress star formation are still to be fully understood. There are a number of proposed mechanisms; some act quickly and others are more gradual, some are driven by gravity and others by hydrodynamic interactions. A full review of such mechanisms can be found in Boselli & Gavazzi (2006) and a summary is presented below.

Tidal interactions. The increased density of galaxies and, hence, decreased separation distance is bound to lead to tidal interactions between galaxy pairs. If the typical galaxy radii are not too small relative to inter-galaxy separation distance, then tidal forces may overcome the gravitational bounding of gas, dust, stars, and dark matter (Spitzer & Baade, 1951). Such encounters are often referred to as ‘*galaxy harassment*’ and can progressively turn disc-like galaxies into spheroidals with most of their gas concentrated in the centre (Moore et al., 1996, 1998, 1999). In galaxy clusters, however, the relative velocities of galaxy pairs are significantly higher than those in the field. Therefore, galaxy encounters are more frequent but they act over much shorter intervals of time. Thus, the effect of the perturbation are much less severe (Fujita, 1998). Galaxy-galaxy tidal interactions will be more effective in low-mass groups where velocity dispersion is significantly lower. Therefore, substantial quenching is possible before galaxies even reach the virial radius of massive clusters: ‘*preprocessing*’. (Mihos, 2004; Fujita, 2004; Dressler et al., 2004). This is very much in line with Λ CDM cosmology, which predicts large numbers of smaller dark matter haloes relative to the rare but massive clusters (White & Rees, 1978; White & Frenk, 1991). Preprocessing means that a substantial number of quenched satellites in clusters we see at present may have accreted as part of a group, where they were quenched. It is one of the main mechanisms invoked to explain quenching at large cluster-centric radii (Kodama et al., 2001; Fujita, 2004; Wetzel et al., 2012; Bahé et al., 2013, 2019).

Tidal interactions between a galaxy and the overall overdensity, on the other hand, have a much higher probability of perturbing the galaxy, inducing gas flows which lead to bar formation, nuclear and disc star formation (Merritt, 1984; Miller, 1986). Tidally perturbed disc galaxies can experience declining rotational curves at large radii due to the increased non-circular velocities of stars (Valluri, 1994). The resulting rise of kinetic

pressure in the interstellar medium can induce star formation (Elmegreen & Efremov, 1997). Tidal interactions have also been used to explain environmental transformation at distances larger than the virial radius of a cluster (Balogh et al., 2000; Treu et al., 2003). A process whereby gas supply from the extended halo to the inner disc is removed due to the presence of a more massive neighbour, preventing further star formation is called ‘*Starvation*’ or ‘*Strangulation*’ (Larson et al., 1980; Bekki et al., 2002). It is capable of suppressing star formation over longer timescales (several Gyr) due to the innermost reservoir of gas being used to sustain star-formation after infall. A signature of starvation/strangulation is delayed quenching which is often reflected in stellar population ages (McCarthy et al., 2008).

Ram pressure. A galaxy moving at $\sim 1000\text{km/s}$ through the intra-cluster medium will experience a force opposing its motion like a wind: ram pressure (Gunn & Gott, 1972). If this force is sufficiently strong to overcome the gravitational pull of the galaxy then it will be capable of removing the gas from it, a process referred to as ‘*ram-pressure stripping*’ (RPS). This will lead to the formation of a tail extending away from the galaxy and opposite to the direction of its motion. This concentration of gas may lead to rapid collapse and temporary star-formation, quickly depleting the gas reservoir (Tonnesen & Bryan, 2012). Such gas tails are commonly observed in cluster galaxies through HI emission as ‘jellyfish galaxies’ (Poggianti et al., 2016, 2017). Evidence of star formation in the tails of stripped gas have also been observed (Ramatsoku et al., 2019; Tonnesen & Bryan, 2012). RPS can result in the removal of only the hot gas halo, in which case only the supply of gas is cut off and star-formation is still possible for some time. This is another case of *starvation/strangulation*, induced by ram-pressure rather than tidal forces. A slightly less extreme version of this is viscous stripping, as proposed by Nulsen (1982). There, stripping happens due to viscosity momentum transfer from the outer layers of the interstellar medium (ISM). This can lead to gas stripping on similarly short timescales and with a similar signature to ram pressure stripping.

Thermal conduction. The intergalactic medium (IGM) is significantly hotter than the gas surrounding a typical galaxy ($10^7 - 10^8\text{k}$ against $< 10^5\text{K}$) giving rise to a substantial temperature gradient across the boundary between the IGM and gas held by the infalling galaxy, ISM. The ISM will experience heating, as a result, causing the gas to evaporate and permanently escape the gravitational field: ‘*thermal evaporation*’ (Cowie & Songaila, 1977). Thermal evaporation is sensitive to the temperature of IGM, magnetic fields, and

to a lesser extent density. For these reasons, the effects of thermal evaporation on satellite galaxies are difficult to quantify.

1.6 Numerical simulations

Progressively better observations have advanced the field to where it is today. Better, deeper observations will progress it even further. There is a fundamental limitation to observing, however. There is no way of interacting with the system being observed, i.e. it is not an experiment. The solution to this problem would be to build one: have a system of dark matter, galaxies, gas, etc. and adjust the equations describing their interactions to make predictions which can then be compared to observations.

Without making some severe simplifying assumptions, the problem of galaxy assembly and evolution cannot be solved analytically. Just the motions of a system with 10^{11} stars is already too complex to model accurately, let alone the interaction of stars with dark matter, other galaxies, and the overall cosmic web. The problem becomes even more complicated when one wishes to capture the process of gravitational collapse of structures in the early Universe, gas flows in the cosmic web, the cooling and condensing of gas into galaxy discs, atomic and chemical evolution of baryons through star formation, as well as feedback from star formation and black holes. The use of computers to solve such non-linear equations has been ongoing for decades now.

1.6.1 N-body and semi-analytic models

Solving gravity

As simulations go, the simplest question that one may ask is whether it is possible to solve the equations of motion and reproduce the kind of structures that are observed. It helps that dark matter is thought to be collisionless and interacts only through gravity. Since it dominates the mass distribution, gravitational interactions are the only ones needed to be solved for. Furthermore, it is not necessary to model full general relativistic gravity, Newtonian gravity is a sufficiently accurate approximation. This comes about because linear structure growth is identical in the matter-dominated regime in both theories, and nonlinear large-scale structure induces velocities far below the speed of light (see

Vogelsberger et al. (2020) for a recent review of numerical simulations). Given the Newtonian approximation, the dark matter fluid can be approximated as finite ‘particles’ encompassing some amount of mass. Such models are commonly referred to as ‘N-body’ simulations. The first N-body simulation is probably the study of a tidal encounter between galaxies of Holmberg (1941). This was run on an analogue optical computer, each galaxy represented by 37 light bulbs, with a photocell and galvanometer measuring the inverse square law. As technology progressed, the number of particles increased so that larger structures could be modelled (Hockney & Eastwood, 1988; Bertschinger & Gelb, 1991). Modern N-body simulations use so many particles that even current machines would not be able to compute the direct summation of Newton’s law for upwards of 10^{11} particles (Angulo et al., 2012). For practical reasons, approximate methods such as the tree particle-mesh, where the direct summation of short-range interactions is approximated by a tree-like method, are employed to greatly increase the efficiency (Bode & Ostriker, 2003). On large scales, particles are mapped onto a three-dimensional grid (particle-mesh) which enables the solving of Poisson’s equation using Fast-Fourier-Transform techniques.

One great advantage of using particles (of fixed mass) to model a fluid is that the number of particles reflects the density of a region being simulated. There is no need to specify that a dense region be resolved with more particles, the scheme automatically does that. Massive structures, such as galaxy clusters, are automatically resolved with more particles as would be desired. However, since the total number of particles in a simulation is limited even in the most optimised schemes, it is inevitable that particles will represent an unphysical concentration of mass ($10^3 - 10^9 M_{\odot}$ for modern simulations). This then leads to unphysical effects at close encounters, such as particles being kicked out at high velocities (Baertschiger et al., 2002). The solution to this problem is to introduce ‘gravitational softening’, which involves replacing the r^2 term in the denominator of Newton’s law with a ‘softening kernel’ which tends smoothly to zero as $r \rightarrow 0$. The implication of such an approach is that structures smaller than the softening length cannot be resolved, and a spatial resolution limit is established. Since using less massive particles in a simulation would mean more particles available to resolve structures, mass of particles is directly related to the softening length. Mass resolution sets the spatial resolution. In modern simulations, mass resolution ranges from $10^3 M_{\odot}$ in single galaxy/cluster simulations to $10^9 M_{\odot}$ in cosmological volume simulations. The resulting

spatial resolution is $\sim 0.1 - 10.0$ kpc (see table 2 of Vogelsberger et al. (2020) for a full list of current simulations).

Adding galaxies through prescriptions

Solving for gravity and reproducing hierarchical structure formation was a major step forward for numerical methods. The next challenge is to reproduce observable objects, i.e. galaxies. The simplest and most common way of achieving this is to make the (reasonable) assumption that galaxies form in dark matter haloes. Therefore, the end result of running an N-body simulation in combination with a FoF algorithm contains information of where the galaxies should be located (also how massive they should be). The next step is to come up with a set of prescriptions of how haloes should be populated with galaxies. Such an approach is commonly referred to as semi-analytic modelling (SAM, Kauffmann et al. 1993; Baugh 2006) due to the use of numerics for the first part and analytic prescriptions for galaxy placement as the second part. The key advantage of SAMs is that they are relatively inexpensive to run; the bulk of the cost is taken up by running the dark matter simulation and the cost of solving the analytic equations is negligible. More importantly, the analytic aspect is added to the final output of a simulation, so one dark matter simulation can be turned into many different SAMs (De Lucia & Blaizot, 2007; Guo et al., 2011; Fabello et al., 2012; Taranu et al., 2013). They have been successfully used to explain aspects of galaxy formation, especially those dominated by accretion such as assembly of central galaxies (De Lucia et al., 2006; Croton et al., 2006; De Lucia & Blaizot, 2007).

The SAM approach is somewhat limited, however, especially when their predictive power is considered. The models used to describe star formation, gas flows, heating, etc. are simplified versions of what takes place in nature. The same way that analytic approaches were insufficient to reproduce structure formation, they are, at best, approximations to the non-equilibrium processes associated with galaxy transformation. While the number of tuning parameters is kept to a minimum, there is still a large number of them. Their values are often obtained by fitting some observed relations. This leads to the second problem of predictive power: if model parameters do not correspond to physical quantities and merely reproduce desired scaling relations, then there is no guarantee

that the same parameter choices will lead to other quantities beyond the scaling relation to be captured.

Lastly, while dark matter dominates the total mass content of the cosmic web, baryonic matter still influences the properties of dark matter: back-reaction. It has been well demonstrated that the processes associated with baryons change the detailed matter distribution, especially in dense regions (Blumenthal et al., 1986; Duffy et al., 2010; van Daalen et al., 2011, 2020). In SAMs, however, baryons do not have any influence on the dark matter halo they occupy.

Subhalo abundance matching

Another approach of populating dark matter haloes with realistic galaxies is the Subhalo Abundance Matching (SHAM, Conroy et al. (2006); Vale & Ostriker (2006)). This method assumes a one-to-one mapping between simulated dark matter haloes and observed galaxies, while preserving rank-order. The assumption is made that most massive galaxies reside in the most massive haloes, matching the corresponding part of a galaxy stellar mass function to the halo mass function. Such a statistical, non-parametric approach reproduces the observed stellar mass function by design, something simulations are often calibrated to reproduce. SHAM has been successful in reproducing many observed galaxy statistics, including spatial clustering, satellite fractions, cluster luminosity functions and luminosity–velocity relations (Berrier et al., 2006; Conroy et al., 2006; Vale & Ostriker, 2006; Yang et al., 2009; Wetzel & White, 2010).

Despite the successes, SHAM is limited due to the need for a well-measured GSMF at the chosen redshift (only really available at low redshift with SDSS) and the fact that it contains no time-varying physical model - the process must be repeated every time a different redshift is desired. Another caveat is that it neglects to consider the possibility of satellites following a different relation to centrals (Neistein et al., 2011; Yang et al., 2012).

1.6.2 Hydrodynamical simulations

Another way of adding baryons to N-body simulations is including them from the very beginning as another form of matter. This involves solving additional forms of interaction

(e.g. pressure) together with gravitational effects. Including baryons requires a lot more than the addition of a second type of particle; while most gas in the Universe is composed of hydrogen and helium, some of it cools and collapses to form stars. Inclusion of baryons is far more numerically demanding than dark matter due to the large dynamic range, highly supersonic flows and large Reynolds numbers (Vogelsberger et al., 2020).

Numerics

Gas in cosmological, galaxy formation simulations is typically described as inviscid ideal gas following the Euler equations. In order to be used on computational problems, these equations need to be discretised, and there are three main ways of doing so: Eulerian, Lagrangian, and arbitrary Lagrangian–Eulerian.

Eulerian methods. Eulerian approach focusses on specific volumes (a mesh) in space through which the fluid flows as time passes. The task is then to solve what happens as fluid flows from one cell to another. Euler equations are expressed in such a way that one needs to solve the hyperbolic partial differential equations used to discretise the problem. Modern cosmological simulations aim to simulate structures over a huge dynamic range, from massive voids to centres of galaxies. For this reason adaptive meshes are commonly used where the mesh size can be reduced based on some refinement criterion. Such approaches are collectively called adaptive mesh-refinement schemes (AMR, Colella & Woodward (1984); Berger & Oliger (1984)). They were first developed for solving general problems involving hyperbolic partial differential equations, and then later were also applied to cosmological simulations (Fryxell et al., 2000; Teyssier, 2002; O’Shea et al., 2004).

Lagrangian methods. The Lagrangian specification of the field assumes an observer that follows an individual fluid parcel (with its own properties such as density) as it moves through space and time. There is no mesh. This scheme is closely related to the particle approach used for dark matter where the parcel is a particle. To give these particles extent, a smoothing kernel is applied which allows for particles to interact with one another. This SPH scheme is the same as described in Section 1.3.2. The interpolation over many particles results in a density field over which Euler equations can be solved. Since particle density automatically adapts in regions of interest, SPH codes are popular in cosmological simulations. (Springel et al., 2001; Wadsley et al., 2004; Springel, 2005).

Arbitrary Lagrangian–Eulerian methods. The last approach is a combination of the two already outlined, aiming to take advantage of the strengths of each method. For cosmological applications, the simulation is constructed from a Voronoi tessellation of a set of discrete mesh-generating points, which are allowed to move freely. Most importantly, due to the mathematical properties of Voronoi tessellation, the mesh continuously deforms and changes its topology as a result of the point motion, without ever leading to problematic mesh-tangling. At the same time, it takes advantage of quantity-conserving properties that a mesh-based algorithm provides. There are a number of different applications of this idea to cosmological simulations (Springel, 2010; Duffell & MacFadyen, 2011; Hopkins, 2015; Vandenbroucke & De Rijcke, 2016).

Both Eulerian and Lagrangian methods are commonly employed in current cosmological hydrodynamic simulations, with arbitrary Lagrangian-Eulerian methods mainly aimed at future iterations. They both have their strengths and weaknesses, the choice ultimately comes down to specific use cases and goals of the simulation. Detailed comparisons have been made over the last few years by O’Shea et al. (2004), Agertz et al. (2007), and Scannapieco et al. (2012). To summarise, AMR-based codes are more successful in modelling fluid discontinuity effects such as Kelvin-Helmholtz or Rayleigh-Taylor instabilities which are important for fluid mixing (Zavala et al., 2012). SPH-based codes have to include a form of artificial viscosity in order to achieve shocks (Cullen & Dehnen, 2010) and conduction (Price, 2008), for example. SPH implementations, on the other hand, have the advantages of being Galilean-invariant, explicitly conserving mass, energy, and entropy. Their Lagrangian nature also makes it possible to trace individual structures even in very crowded environments. AMR-based codes need complex workarounds to achieve the same result.

Sub-resolution models

Just like with dark matter and gravity, there is a finite number of particles/cells available for resolving cosmic structures. There comes a point below which processes can no longer be resolved. The hydrodynamic equations have to be complemented by equations for astrophysical processes taking place on scales smaller than the resolution: sub-resolution (or ‘sub-grid’) models. These models are not vastly different from those used in SAMs in their concept, but they differ by the fact that they use hydrodynamic gas properties

as inputs and there are relatively few free parameters to calibrate. Hydrodynamic simulations have more predictive power as a result, particularly with regards to how the gas responds to cooling, feedback, etc. It must be noted that these models are independent of the underlying hydrodynamics scheme: SPH or AMR code is chosen before any sub-grid models are considered. The hydrodynamics scheme is not adjusted in response. The main sub-resolution models implemented in simulations include:

Gas cooling. Gas radiates and dissipates its internal energy through cooling processes, such as collisional excitation and ionisation, inverse Compton, recombination and free-free emission. Since such processes take place on the atomic level, they cannot be resolved. As a result, simulated cooling processes are coupled to the energy equation using cooling functions that are either tabulated from previous calculations or extracted from chemical reaction networks. Simulations often assume that the gas is optically thin and in ionisation equilibrium. Few simulations model dust on the fly, in most cases its effects are computed in post-processing.

Interstellar medium. Since star formation takes place in molecular clouds of cool gas, modelling this phase correctly is very important. However, doing so is rather difficult because of its multiphase structure spanning a wide dynamic range of scales. Processes in dense gas tend to take place on timescales significantly shorter than the gap between timesteps. For these reasons, dense gas is often described by an effective polytropic equation of state (Springel & Hernquist, 2003; Dalla Vecchia & Schaye, 2012) relating temperature of the gas particles to density. Star formation conditions can then be implemented based on these quantities.

Star formation. Observations support a nearly universal star formation efficiency in molecular gas, where about 1% of the gas is converted into stars per free-fall time (Bigiel et al., 2011; Krumholz et al., 2012). Simulations, therefore, need a model which transforms a portion of the gas particles into collisionless star particles with an underlying initial mass function. The exact particles are chosen probabilistically, reflecting the underlying star formation rate. Most schemes contain a relation which models the Kennicutt-Schmidt law (Schmidt, 1959; Kennicutt, 1998), depending primarily on the mass of gas. Criteria for density threshold, virial parameter, and Jeans length are also introduced to ensure that star formation takes place in dense regions, prone to gravitational collapse.

Stellar feedback. Stars interact with their surrounding gas through the injection of energy and momentum resulting in a feedback loop regulating star formation. To achieve the low conversion fraction observed, stellar feedback must be capable of launching galactic-scale outflows to eject gas from galaxies (Springel & Hernquist, 2003; Pillepich et al., 2018b; Vogelsberger et al., 2013). There is a variety of subresolution schemes aiming to achieve an efficient generation of galactic winds. The main area of difference is in how energy and momentum from supernova explosions couple to the surrounding gas. Some implementations feature a stage of decoupling between feedback and the surrounding matter so that momentum can be injected at a distance without affecting the immediate surroundings of the originating source. The wide variety of different implementations demonstrates that this area is not well understood yet.

Supermassive black holes. Every galaxy appears to host a black hole at its centre (Gehren et al., 1984; Kormendy & Richstone, 1995). However, the means by which black holes first form is not well understood, let alone modelled in simulations. As a result, they are artificially seeded once gas density reaches a specified threshold (Springel et al., 2005). They then accrete mass often based on an Eddington-rate-capped Bondi–Hoyle-like accretion rate. Due to limited spatial resolution, measures have to be taken to maintain realistic accretion rates in the centres of galaxies (Booth & Schaye, 2009; Hopkins et al., 2011).

Active galactic nuclei feedback. Accretion onto black holes leads to feedback in the form of electromagnetic radiation, relativistic jets, and non-relativistic outflows (Krolik, 1999). These outflows play a major role in regulating star formation in massive galaxies and, most likely, environmental quenching of satellites due to their role in heating the ICM. This feedback is commonly (but not always) divided into two modes: quasar and radio. Quasar mode feedback is associated with the radiatively efficient mode of black hole growth and is often implemented through energy or momentum injection. An assumption is often made that the bolometric luminosity is proportional to the accretion rate, and a fixed fraction of this luminosity is deposited into the neighbouring gas. Radio mode feedback is caused by highly collimated jets of relativistic particles, often assumed to be the cause of X-ray bubbles in the surrounding medium. Both modes are implemented separately and controlled by specific criteria, often calibrated to reproduce certain observations (Springel, 2005; Di Matteo et al., 2005). Some simulations do not differentiate between the two modes, arguing that resolution is not sufficient to resolve

the differences (Schaye et al., 2015). A substantial level of uncertainty about the exact implementation of AGN feedback remains.

1.7 This thesis

Having made the case for the existence and importance of environmental influence on galaxy evolution I will now give a brief outline of the work presented in this thesis. It can be summed up as presenting two main ideas: (i) improvements to the way the role of galaxy environment is studied at low redshift, where data is sufficiently abundant to allow for alternative approaches, and (ii) application of cosmological, hydrodynamical simulations to environment at redshifts $z \sim 1$. We have reached a stage in the field of galaxy environment where alternative methods and tools have become available that can help us unravel this question of great complexity. In what follows I present and advocate for the use of these tools.

Chapter 2 introduces the methods of spatial cross-correlation of two quantities and applies it to the problem of cause and effect in galaxy environment. It is widely accepted that either hot gas or gravity is behind galaxy quenching in dense environments, but clear connections have been difficult to establish. Such a method avoids the issues of galaxy membership assignment - a problem highlighted throughout this thesis. In this chapter I present statistically significant signals between gas density or pressure and galaxy density or quenched fraction. I also present the equivalent signals obtained using hydrodynamical simulations. This method promises great potential as this type of observation will become better and more common in the not too distant future.

Chapter 3 explores the possibility of using hydrodynamical simulations to understand environmental processes at $z \sim 1$. As this regime is not something they are specifically built to reproduce, most quantities are predictions and can be compared against observations. Not all quantities have been extracted from observations yet but, nevertheless, a comparison between different simulations can be done in preparation. I make this comparison, showing that three different codes make rather different predictions for halo content and star-formation activity at $z \sim 1$. We can learn quite a lot by comparing these to observations when they become available. Such exercises highlight areas in need of improvement for future generations of simulations.

Chapter 4 explores the possible bias being introduced by a typical observational galaxy cluster member selection. There are numerous ways to assign galaxies to clusters, all with their benefits and shortcomings. In this chapter, I compare the most common method: imposing cuts of an aperture and line-of-sight velocity on observed galaxy redshifts and compare it to the Friends-of-Friends selection of galaxies in a simulation. This analysis is linked to measuring the relative number of galaxies when split by type at $z \sim 1$, as was done in Chapter 3.

Finally, in Chapter 5 I summarise all of the preceding chapters, draw conclusions pertinent to galaxy environment, and propose possible extension to each chapter.

Chapter 2

Environment from cross-correlations: connecting hot gas and the quenching of galaxies

2.1 Introduction

Surveys such as the Sloan Digital Sky Survey (SDSS) and Galaxy and Mass Assembly (GAMA) (among others) have shown that in the local Universe, galaxies can be classified into two broad populations: star-forming and passive. Star-forming systems typically have blue colours (Strateva et al., 2001; Blanton et al., 2003; Baldry et al., 2004), late-type morphologies (Wuyts et al., 2011; van der Wel et al., 2014), young stellar populations (Kauffmann et al., 2003; Gallazzi et al., 2008), and high star formation rates (SFR) (Noeske et al., 2007; McGee et al., 2011; Wetzel et al., 2012). On the other hand, passive galaxies exhibit red colours, early-type morphologies, old stellar ages – all of which can be connected to their low star formation rates.

These effects manifest themselves in the form of a well-known galaxy bimodality (in colour, star-formation rate, or specific star-formation rate) when plotted as a function of stellar mass, persisting out to redshifts as high as $z \sim 4$ (Baldry et al., 2006; Brammer et al., 2009; Muzzin et al., 2013). There is a tight, positive correlation between stellar mass and star-formation rate for blue, star-forming galaxies known as the ‘main

sequence' (MS), which is in stark contrast to the passive population dominating the high-stellar mass end. The presence of a bimodality suggests that a relatively rapid transition occurs in the evolutionary sequence of a galaxy where it ceases its star formation, i.e. quenching.

It is yet undetermined which exact physical processes are responsible for this evolution. It is possible, however, to separate them out into two categories: internal and external processes (Peng et al., 2010). The former, secular mode, which can occur in all galaxies irrespective of external factors, is strongly correlated with stellar mass (Driver et al., 2006; Barro et al., 2017). Indeed, this mode is more pronounced at $\log_{10}(M_*/M_\odot) \gtrsim 10$, where Active Galactic Nucleus (AGN)-driven outflows (Nandra et al., 2007) and heating suppress star formation processes. At the low-mass end, stellar feedback is thought to regulate star formation activity. In galaxies with $\log_{10}(M_*/M_\odot) \lesssim 9$, which do not host a strong AGN, stellar-driven outflows eject the more loosely-bound gas (Dalla Vecchia & Schaye, 2008a) but may not remove it completely, allowing it to fall back onto the galaxy where it can become available for further star formation.

A galaxy's local environment can lead to additional quenching by inhibiting the supply or outright removing gas required to fuel star formation. It is now well established that over-dense environments host galaxies with suppressed star formation rates and the observational features outlined above (Oemler, 1974; Postman & Geller, 1984; Dressler et al., 1999; Blanton & Moustakas, 2009; Kimm et al., 2009; Peng et al., 2010; Wetzel et al., 2012). Even on an individual cluster scale, gradients in quenched fraction (f_q) have been observed to correlate with cluster-centric distance (Rasmussen et al., 2012; Wetzel et al., 2012; Haines et al., 2015; Barsanti et al., 2018) such that cluster centres are dominated by quenched galaxies whereas the outskirts host more star-forming galaxies. In addition, satellite galaxies are observed to be more metal-rich in their ionised gas and stars than those in the field (Pasquali et al., 2012; Bahé et al., 2017a; Maier et al., 2019). Relative to the field, enhanced quenched fractions extending beyond the virial radii of clusters have also been observed and attributed to 'pre-processing', whereby galaxies begin quenching as part of smaller groups prior to in-falling onto a cluster (Fujita, 2004; Lu et al., 2012; Wetzel et al., 2012; Bahé et al., 2013; Roberts & Parker, 2017).

There are many possible mechanisms for quenching satellite galaxies as they spiral in towards massive groups and clusters, either individually or as part of a smaller group.

For example, the cold, star-forming gas of infalling galaxies can be directly ram pressure stripped (RPS) due to the relative motion of the galaxy with respect to diffuse medium of the host system (Gunn & Gott, 1972; Abadi et al., 1999; Quilis et al., 2000; Poggianti et al., 2017; Brown et al., 2017; Barsanti et al., 2018). Material being circulated in feedback-driven galactic fountains may be even more susceptible to RPS (Bahé & McCarthy, 2015). Turbulent viscous stripping (Nulsen, 1982; Kraft et al., 2017) acts as an additional form of stripping alongside RPS. ‘Strangulation’ (or ‘starvation’) refers to the process through which the hot gas reservoir of an infalling galaxy, which supplies the fuel for ongoing star formation, is removed (Larson et al., 1980; Moore et al., 1999; Balogh et al., 2000; Kawata & Mulchaey, 2008; McCarthy et al., 2008; Peng et al., 2015). In this case, the infalling galaxy is quenched more slowly, on an ISM gas consumption time scale.

RPS, turbulent viscous stripping, and strangulation represent hydrodynamical processes. However, gravitational interactions can also result in the quenching of galaxies. For example, galaxy-galaxy mergers can cause an initial burst of star formation but leave the galaxy quenched in the end (Mihos & Hernquist, 1994a,b; Schawinski et al., 2014). ‘Harassment’ is a less dramatic process of inducing starbursts through repeated dynamical interactions (Farouki & Shapiro, 1981; Moore et al., 1996; Hirschmann et al., 2014), thus exhausting the gas reservoir. In addition, tidal interactions with the overall group/cluster potential well (Mayer et al., 2006; Chung et al., 2007) can disturb the cooling and accretion of gas making it less bound and, therefore, easier to be stripped by other quenching modes.

In addition to the above hydrodynamical and gravitational processes, it is possible that the local radiation (e.g., Kannan et al. 2016) and magnetic fields (e.g., Tonnesen & Stone 2014) may also play a role in the quenching of satellites. However, at present these possibilities are ill-constrained by observations.

The main limitation in our inability to robustly identify the main mechanism(s) behind environmental quenching (and its possible dependence on, e.g., time or galaxy mass) likely stem from our inability to observe each process independently. Instead, the focus has shifted to characterising when/where (rather than how/why) galaxies become quenched and then trying to use this information to test different physical models. For example, the currently favoured scenario is the ‘delayed-then-rapid’ scenario proposed

by Wetzel et al. (2013). This is a two-stage process, whereby a satellite galaxy experiences a slow form of quenching as it initially falls into a cluster (possibly through starvation). As it reaches the dense, inner part of the ICM, a short-timescale process (such as RPS) becomes much more efficient and begins to dominate - inducing a rapid form of quenching. Evidence for this two-stage scenario comes from studying quenching timescales (Muzzin et al., 2014; Peng et al., 2015; Fillingham et al., 2015; Foltz et al., 2018) or observing RPS in action through extended HI distributions around satellite galaxies in dense environments (Kenney et al., 2015; Poggianti et al., 2016; Jaffé et al., 2016).

The above-mentioned studies focussed on observing galaxy properties as a function of environment, i.e. an over-density probed via the same galaxies. There are a number of approaches used in estimating the local density; for example, group membership counts from a catalogue such as Yang et al. (2007), nearest neighbour distance (Park et al., 2007), number counts of neighbouring galaxies in a defined volume (usually cylindrical) (Kauffmann et al., 2004; Blanton & Moustakas, 2009), or taking their velocity into account as well through phase space diagrams (Pimblet et al., 2006; Oman & Hudson, 2016). Some studies have investigated environmental effects via angular cross-correlations, asking the question of whether galaxies are clustered differently when separated by their star formation rate (Hatfield & Jarvis, 2017) or morphology (Cervantes Sodi et al., 2016). However, while it is true that galaxy overdensity is a tracer of the underlying overdensity and environmental quenching correlates with it, probing the overdensity stands little chance of determining *which* of the processes are more dominant. This is because the majority of the different cluster mechanisms/components responsible for quenching correlate with the underlying matter density distribution. It is, therefore, necessary to observe multiple components at the same time and measure the correlation with galaxy properties in order to break the degeneracies present. This is easier said than done, however, as the diffuse ICM is extremely faint in X-ray emission and unbiased estimates of total overdensity from weak lensing are noisy as well as low in number, particularly on the scale of groups.

One recent attempt has been made using deep *Chandra* X-ray observations of low-redshift clusters in SDSS by Roberts et al. (2019). The authors find evidence for a threshold in ICM density which separates regions of gradually-increasing quenched fraction and sudden steepening of the trend closer to the cluster centre. This is interpreted

as the dominant quenching mechanism transitioning from a steady gas depletion (e.g. starvation) to a more abrupt gas removal process like direct ram pressure stripping of the ISM.

In this study, I introduce a new test of environment based around spatial cross-correlations between observables of the ICM/gravitational potential and large-area galaxy catalogues. This map-based approach offers the advantages of extracting a signal from otherwise low signal-to-noise observations by measuring over a large area of the sky, avoiding the complex process of finding groups and clusters, and potentially breaking the degeneracies between different cluster components. Such methods have already been employed in other areas of astrophysics, such as galaxy–CMB lensing to probe cosmology via the growth of structure (Giannantonio et al., 2016), thermal Sunyaev-Zel’dovich (tSZ) effect – X-ray emission – CMB weak lensing in order to measure halo bias in the clustering of dark matter (Hurier et al., 2019), and the unresolved γ -ray background – galaxy cluster cross-correlation to study the nature of this γ -ray emission (Hashimoto et al., 2019). As a first application of the method to environmental quenching, I focus on the hot gas–quenched fraction signal. I construct large-area maps of galaxy overdensity and quenched fraction from the SDSS spectroscopic and photometric surveys, compute auto- and cross-power spectra between galaxy survey quantities and two measures of ICM. To characterise the state of the hot gas, I make use of maps of the thermal Sunyaev-Zel’dovich (tSZ) effect as measured by the *Planck* satellite and X-ray emission from the ROSAT All-Sky Survey (RASS), respectively. I perform equivalent measurements on synthetic maps produced from state-of-the-art hydrodynamical simulations: EAGLE (Schaye et al., 2015; Crain et al., 2015; McAlpine et al., 2016) and BAHAMAS (McCarthy et al., 2017, 2018).

This chapter is organised in the following sections: I describe the galaxy catalogues, X-ray, and tSZ effect data and map-making procedures, as well as simulations in this study, in Section 2.2. The formalism of cross-correlating two discretised maps is outlined in Section 2.3. My main results are presented in Section 2.4 and discussed further in Section 2.5. Finally, I summarise my findings in Section 2.6.

Throughout, I adopt a flat Λ CDM concordance cosmology with $\Omega_m = 0.274$ and $H_0 = 70.5$ km/s/Mpc (Hinshaw et al., 2009).

2.2 Data and Map-making

2.2.1 Galaxy catalogues

For my galaxy samples I use the Sloan Digital Sky Survey (SDSS) DR7 (Abazajian et al., 2009) and DR12 (Alam et al., 2017) to construct two volume-limited samples. As described below, the DR7 is used to construct a shallower spectroscopic sample ($z < 0.06$), while the DR12 is used to construct a deeper ($z < 0.15$) photometric sample.

2.2.1.1 SDSS DR7 spectroscopic sample

The MPA-JHU¹ value-added galaxy catalogue (Brinchmann et al., 2004) provides derived galaxy properties from emission line analysis of SDSS DR7. I use their stellar masses, specific star formation rates ($\text{sSFR} = \text{SFR}/M_*$), observed `cModel` magnitudes, redshifts, and positions, and call this the ‘spectroscopic’ sample. To form my sample, I first select all objects identified as galaxies with the `TARGETTYPE = GALAXY` parameter in the `gal_info_dr7_v5_2.fits` file. Next, I select all galaxies with reliable (specific) star formation rates, i.e. `FLAG = 0` in the `gal_totsecsfr_dr7_v5_2.fits`. The catalogue is complete to Petrosian r-band magnitude $r \leq 17.77$ over most of the SDSS footprint, but to ensure a consistent sky coverage I use a more conservative limit of $r \leq 17.5$. Note that `cModel` and Petrosian magnitudes are sufficiently similar so that I can ignore the differences between them in terms of selection. Post selection I use the former as it provides a more reliable estimate of a galaxy’s total flux and has close to optimal noise properties (Stoughton et al., 2002).

In order to aid the interpretation of the observations and to make a straightforward and consistent comparison to simulations (described in Section 2.6), I opt for a volume- and stellar mass-complete sample. Ideally, I would like to probe the low-mass end of the galaxy population as these galaxies are more likely to be quenched due to environmental effects. With a flux-limited survey, however, a balance must be struck between the redshift limit and the lower mass limit if a volume-limited sample is desired. A lower stellar-mass limit is also introduced by simulations, as the resolution is finite and the low-mass galaxy properties become unreliable. BAHAMAS is the lower resolution simulation

¹<https://wwwmpa.mpa-garching.mpg.de/SDSS/DR7>

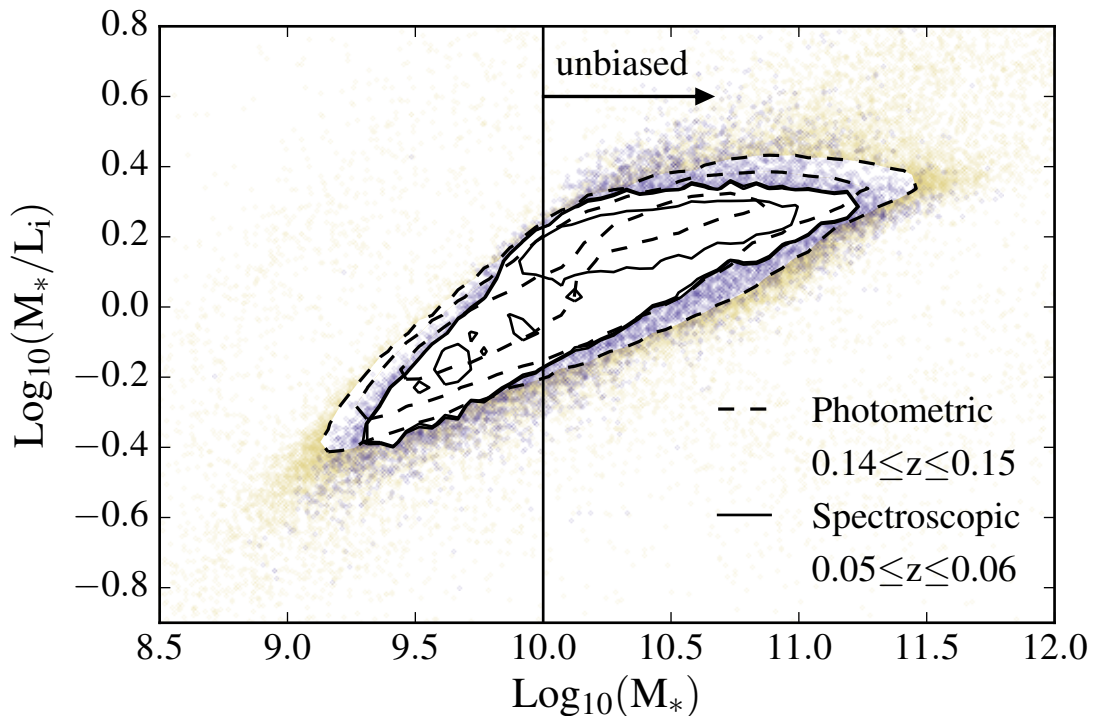


FIGURE 2.1: Mass-to-light ratio – stellar mass relation used to find volume completion limits. Navy points/solid contours and yellow points/dashed contours show thin slices in redshift of the spectroscopic and photometric samples, respectively. Redshift limits were fine-tuned such that no bias for galaxies above $\log_{10}[M_*/M_\odot] > 10$ is introduced. The same limits are determined by Baldry et al. (2018) for a fully spectroscopic sample of GAMA galaxies, where the drop-off is more pronounced.

of the two used in this study and has been shown to reproduce galaxy properties down to a stellar mass of $\log_{10}[M_*/M_\odot] = 10$. I therefore adopt this as my lower-limit in galaxy stellar mass.

To determine the limiting redshift for a given stellar mass cut in a volume-limited sample, I adopt the method described in Baldry et al. (2018). Specifically, I compute and examine the mass-to-light ratio in the i-band (M_*/L_i) against stellar mass (M_*) in a redshift slice (see Fig. 2.1); a clear drop-off in M_*/L_i can be seen for stellar masses which are no longer completely sampled given the r-band limit of the survey. I adjust my redshift upper limit such that the drop-off occurs at $\log_{10}[M_*/M_\odot]$ just below 10. For a volume-limited subsample (selected from the main SDSS sample) of galaxies with $\log_{10}[M_*/M_\odot] > 10$, the upper limit in redshift is $z = 0.06$. If a higher-redshift/larger volume is desired, the lower limit on stellar mass must be raised for the sample to remain volume-complete. But since the aim of this study is to characterise environment, it is necessary to cover

the range in M_* over which galaxies transition from blue/star-forming to red/quenched, so a lower stellar mass limit is preferred to one that probes larger volumes but with higher-mass systems.

I further split my sample into star-forming and quenched sub-samples by introducing a simple division in sSFR following Wetzell et al. (2012) at $\log_{10}[\text{sSFR}(\text{yr}^{-1})] = -11$. This division is then used to compute quenched fraction in Section 2.2.2.

2.2.1.2 SDSS DR12 photometric sample

In order to obtain a larger and deeper sample (in terms of limiting redshift) of galaxies, I also use a sample not restricted by the spectroscopic completeness of SDSS. For this purpose I use SDSS DR12 data, with photometric redshifts and parametrically-estimated stellar masses and refer to it as the ‘photometric’ sample.

I use photometric redshift estimates of Beck et al. (2016) obtained using a hybrid method of machine learning and template-fitting techniques. They achieve a normalised mean redshift estimation error of $\overline{\Delta z_{\text{norm}}} = 5.84 \times 10^{-5}$, where

$$\Delta z_{\text{norm}} = (z_{\text{phot}} - z_{\text{spec}})/(1 + z_{\text{spec}}), \quad (2.1)$$

a standard deviation $\sigma(\Delta z_{\text{norm}}) = 0.0205$, and an outlier rate of $P_o = 4.11\%^2$. Since spectroscopic redshifts are available for all galaxies with $r \lesssim 17.77$ in the main SDSS footprint (I use these in the spectroscopic sample), and photo z errors have the largest impact at low redshift, I use MPA-JHU spectroscopic redshift estimates where available. Therefore, although the sample is called ‘photometric’, the redshifts used are a combination of spec z and photo z . The stellar masses are estimated using the same method (described below) for all galaxies, independent of whether their redshifts are spectroscopic or photometric, in order to stay consistent throughout the sample.

Empirical stellar mass estimates are computed following Sedgwick et al. (2019), which itself is based on the method outlined in Taylor et al. (2011) and Bryant et al. (2015), and calibrated using SED-fitting data from the GAMA (Baldry et al., 2018) survey. The estimation relies on the correlation between mass-to-light ratio and colour. One

²Outliers are cases where the process has failed to return a redshift estimate within 3σ of the training validation value.

can write down an equation for stellar mass that depends only on distances, redshifts, and observed magnitudes and which folds in the k-correction:

$$\log_{10}(M_*/M_\odot) = -0.4i + 0.4D + f(z) + g(z)(g - i)_{obs}, \quad (2.2)$$

where i is the i-band observed `cModel` apparent magnitude, D is the distance modulus to the galaxy, z is the redshift estimate, $(g - i)_{obs}$ is the observed $g - i$ colour³, and $f(z)$ and $g(z)$ are fitted polynomial functions of redshift:

$$\begin{aligned} f(z) &= -15.15z^3 + 9.193z^2 - 1.687z + 1.104, \\ g(z) &= 26.40z^3 - 12.84z^2 + 0.5908z + 0.8237. \end{aligned} \quad (2.3)$$

To test the derived stellar masses, I match my photometric sample galaxies to MPA-JHU and GAMA⁴ catalogues by `ObjID` and compare catalogue (i.e., spectra-based) stellar masses to those resulting from eqn. 2.2. This can be seen in Figure 2.2. A comparison to MPA-JHU tests the parametric fit without photoz errors (spectroscopic redshifts were used for everything with $r \leq 17.5$) and matching to GAMA galaxies tests the validity of the estimate overall. From Fig. 2.2 it is apparent that while there is scatter of ~ 0.3 dex in the empirically-derived stellar masses, there does not appear to be any systematic bias. Furthermore, the visibly larger scatter at $\log_{10}[M_*/M_\odot] < 10$ has no consequence for my sample as I exclude these (grey shaded region) galaxies. Scatter across the boundary, i.e. contamination, is at the $\sim 5\%$ level and is approximately equal in both directions. Therefore, I conclude that photometric redshifts combined with empirical stellar mass estimates do not significantly bias my galaxy samples. When selecting my photometric sample, I adopt the r-band Petrosian magnitude limit of $r < 19.8$ from the GAMA DR3 survey, making the assumption that the rest of the SDSS field is complete at least down to this magnitude, and use the same redshift limit of $z \leq 0.15$ for a volume-complete sample of galaxies with $\log_{10}[M_*/M_\odot] \geq 10$ from Baldry et al. (2018).

³The colour is computed using `model` magnitudes in order to stay consistent between the two bands, as the aperture parameters are determined in the r-band and applied to all other bands.

⁴The GAMA DR3 footprint is, unfortunately, too small to achieve significant detections of the cross-correlations with the tSZ effect and X-ray data used in this study. Future high sensitivity X-ray (e.g., with eROSITA) and tSZ effect data (e.g., Simons Observatory, CMB-S4) will allow this issue to be overcome. Larger deep spectroscopic surveys (e.g., WAVES, DESI) are also expected to improve these detections dramatically. In the present study, I therefore use GAMA solely for the purpose of calibrating stellar mass estimates of the SDSS photometric sample.




FIGURE 2.2: A comparison of parametric stellar mass estimates using equation 2.2 to MPA-JHU estimates using line ratios, and GAMA SED-derived estimates. Blue with solid, black contours shows MPA-JHU-matched galaxies, whereas green with dashed contours corresponds to GAMA-matched sample; they represent 90th, 70th, and 50th percentiles. Galaxies in the photometric sample are matched to both catalogues by their SDSS object ID. Scatter in the stellar mass estimates is ~ 0.3 dex but there is no significant bias in the stellar mass estimates of the photometric sample. Grey shaded region indicates stellar masses which fall below the completeness limit and thus are not used in the sample. They are shown here for illustration purposes only.




FIGURE 2.3: Galaxy stellar mass functions for both samples, normalised by their respective footprint area. Solid and dashed curves correspond to the spectroscopic and photometric samples, respectively. The SDSS photometric sample slightly underestimates the GSMF at masses below $\sim 10^{10.7}M_{\odot}$ with respect to GAMA, achieving a good match otherwise. For completeness, a low-redshift version of photometric (dashed, blue curve) is shown, which is nearly identical to SDSS spectroscopic sample. Also shown are the GSMFs from EAGLE and BAHAMAS, which were calibrated on previous estimates of the GSMF. Shaded regions represent the 1σ confidence interval on the mean value of 10 simulated light cones. Overall the simulations yield similar GSMFs to those derived from my observational samples, although EAGLE tends to fall somewhat below the observations at the knee of the mass function (as also found previously).

As an additional test of my samples, I plot in Fig. 2.3 the galaxy stellar mass functions (GSMF) for the two SDSS samples and compare them to that derived from GAMA and from the simulations used in this study. Here I define the galaxy stellar mass function as the number of galaxies per decade in stellar mass per unit angular area on the sky, within the two redshift limits mentioned above. (The survey angular area is obtained from the HEALPix⁵ maps described below. For reference, the survey areas are: SDSS = 7849, GAMA = 153, EAGLE = 100, BAHAMAS = 625 square degrees.) Dashed curves are for the spectroscopic selection ($z < 0.06$) and solid curves for photometric selection ($z < 0.15$).

⁵<http://healpix.sourceforge.net>

Good consistency is obtained for both the SDSS photometric and spectroscopic GSMFs with those derived from GAMA. For completeness, although this sample is not used in any further analysis, I also show a low-redshift version of the photometric sample ($z < 0.06$), which yields a near identical GSMF to that derived from the spectroscopic sample. This suggests that scatter seen in Fig. 2.2 does not affect the stellar mass distribution statistics in a significant way. Close agreement is also achieved with the simulations. This is not particularly surprising, as the feedback prescriptions in both EAGLE and BAHAMAS were tuned to reproduce estimates of the local GSMF from previous studies (see Schaye et al. 2015; McCarthy et al. 2017). However, it is reassuring that neither my sample selection and stellar mass estimation on the observation side, nor my light cone making methods on the simulation side, introduce systematic biases. In the case of EAGLE, the simulations fall somewhat below the observations near the knee of the mass function for the deeper photometric selection. This is consistent with what was found previously in Schaye et al. (2015) (see their figure 4).

With no estimates of the SFR for galaxies without spectra, I cannot use sSFR as a means to separate star-forming and quenched galaxies for the photometric sample. Instead, I use the observed $(u - r)$ colour- M_* relation to divide the sample. In Fig. 2.4 I show this relation for GAMA-matched galaxies in the photometric sample, where galaxies are coloured according to quenched status (red=quenched, blue=star-forming) based on their sSFR. It is clear that most of the galaxies belonging to the ‘red sequence’ are quenched, as determined from their sSFRs. However, there is also a population of star-forming galaxies which inhabit the red sequence, presumably as a result of strong reddening by dust (Evans et al., 2018). I demonstrate later that this small contamination by star-forming galaxies is negligible for my cross-correlations.

Given that quenched galaxies lie almost exclusively on the red sequence, I can use a galaxy’s colour to assign a quenched flag. I do this by first sub-dividing the sample into three redshift bins with approximately equal number of galaxies in each. This is done in order to reduce the effect of k-corrections and account for any possible evolution of colour with redshift in the chosen range (although $0.008 \leq z \leq 0.15$ is small enough so that evolutionary effects are not significant). Next, a line is fitted to the red sequence in the GAMA-matched sub-sample to obtain a slope and intercept. I apply this relation to the full SDSS photometric sample, maintaining the slope of the relation but adjusting the intercept until the mean quenched fraction matches that of the GAMA-matched sample




FIGURE 2.4: An example of the colour–stellar mass relations used to calibrate the quenched–star-forming division for the SDSS photometric sample. SDSS galaxies are matched to GAMA and assigned a binary (red/blue) sSFR flag. Virtually all quenched galaxies (on the basis of their GAMA sSFR) lie on the ‘red sequence’. Some star-forming galaxies also lie on the red sequence (presumably as a result of strong dust reddening, see text). The black solid line shows the colour division which achieves the same mean quenched fraction and $f_q - M_*$ relation as for GAMA (see Fig. 2.5). The shaded region marks stellar masses which fall below my limit for volume-completeness.

(computed using sSFR) in each bin. The process is repeated for all three redshift bins, visually inspecting the $(u - r)$ colour– M_* relation, and assigning a quenched flag to each galaxy. The division is not perfectly clean, however. In particular, the number of galaxies which have colours consistent with the quenched population but are star-forming based on spectral information is $\sim 20\%$ over all three redshift bins.

The resulting distribution is verified in the left-hand panel of Figure 2.5, which shows the relation between quenched fraction and stellar mass. To within 50% the colour-determined $f_q - M_*$ relation derived from the SDSS photometric sample matches the sSFR-determined $f_q - M_*$ relation derived from the smaller GAMA spectroscopic calibration sample.




FIGURE 2.5: *Left:* Quenched fraction (f_q) as a function of stellar mass for the $0.008 \leq z \leq 0.15$ sample. f_q is computed in each $\log_{10}(M_*)$ bin using sSFR for GAMA, EAGLE, and BAHAMAS, and $(u-r)_{\text{obs}}$ colour for SDSS. The solid and dashed curves represent the observations and simulations, respectively, with shaded regions showing the 1σ scatter regions around the median of 10 light cones for the simulations. *Right:* Running quenched fraction, $f_q(< M_*)$, as a function of stellar mass. The SDSS photometrically-derived trend matches the spectroscopically-inferred relation from the smaller GAMA calibration set. The simulations predict quenched fractions that fall significantly below what is observed for stellar masses of $\log_{10}[M_*/M_\odot] > 10.5$. However, galaxies near the lower limit $\log_{10}[M_*/M_\odot] = 10.0$ dominate the sample (see text).

Computing f_q for the simulations (using SFR and M_* computed within a 30 kpc aperture) yields relations which fall significantly below what is observed for stellar masses of $\log_{10}[M_*/M_\odot] > 10.5$. (The result for EAGLE is consistent with that shown previously by Schaye et al. 2015 and Furlong et al. 2015.) Here I note that the different feedback schemes employed in EAGLE and BAHAMAS were not calibrated on this metric and were therefore not guaranteed to reproduce these observations. Nevertheless, this comparison illustrates that there are still some deficiencies in the feedback prescriptions of these simulations. Naïvely, one might expect these deficiencies to compromise comparisons of cross-correlations involving quenched fraction. However, it is important to note that for the selection employed in this study, the vast majority of the signal is dominated by galaxies near the lower stellar mass limit, where the simulations have

	Spectroscopic	Photometric
r-band limit	17.5	19.8
$M_{*,\min}$	$10^{10}M_{\odot}$	$10^{10}M_{\odot}$
z_{\min}	0.008	0.008
z_{\max}	0.06	0.15
z_{med}	0.047	0.118
N_{gal}	44799	953980
f_{q}^7	0.418	0.494

TABLE 2.1: A summary of galaxy sample properties for the spectroscopic and photometric SDSS volume-limited samples.

reasonable quenched fractions. This is just by virtue of the fact that the lower mass galaxies dominate the sample by abundance. To illustrate this, in the right hand panel of Figure 2.5 I show the cumulative quenched fraction as a function of stellar mass. This relation reveals that galaxies above $\log_{10}[M_*/M_{\odot}] \sim 10.7$ contribute very little to the total number of galaxies in the sample. I have also checked that the cross-correlations I present later (in Section 2.4) do not qualitatively change when I exclude high-mass galaxies with $\log_{10}[M_*/M_{\odot}] > 10.5$ from my samples⁶. I also point out that the auto- and cross-power spectra that I examine involve overdensities with respect to the mean quenched fraction, rather than the mean quenched fraction itself.

Several basic sample properties are summarised in Table 2.1.

2.2.2 HEALPix map-making using galaxy catalogues

In order to perform the cross-correlations between galaxy and hot gas properties I need to construct equivalent maps for the properties I am interested in. In terms of galaxy properties, in this chapter I focus on two quantities: total galaxy overdensity (\check{N}_{tot}), and quenched fraction (\check{f}_{q}). I adopt the ‘breve’ (‘ $\check{}$ ’) notation to denote excess-relative-to-the-mean quantities, such as overdensity: $\check{N}_{\text{tot}} = (N_{\text{tot}} - \bar{N}_{\text{tot}})/\bar{N}_{\text{tot}}$, where N_{tot} is the total projected galaxy surface density and \bar{N}_{tot} is its mean value. This ensures that maps for both galaxy measures are in the same $[-1, \infty)$ range, and have a mean value of zero.

⁶Given the differences at the high-mass end present in the left-hand panel of Figure 2.5, I could expect significant differences between the simulations and observations for cross-correlations involving exclusively high-mass systems. However, the relatively low abundance of high-mass systems results in noisy estimates of these cross-correlations at present. Deeper observations (e.g., with DES, LSST, Euclid) will resolve this issue in the near future.

⁷These values are final mean quenched fractions of the maps after all masks have been applied.

For the galaxy-based maps, I adopt the same `HEALPix` (Górski & Hivon, 2011) format and resolution (~ 1.716 arcmin, $N_{\text{SIDE}} = 2048$) as used for the tSZ effect and X-ray maps, which are described below. To compute the cross-correlations (described in Section 2.3), I use tools (e.g., `NaMASTER`) originally designed for analysis of contiguous fields in the `HEALPix` format, such as those regularly produced using cosmic microwave background data. The `NaMASTER` algorithm has the capability to deal with non-contiguous/incomplete fields to an extent, but I have found that even for my larger photometric sample, the galaxies are too sparsely distributed for the algorithm to give reliable results if I simply mask empty pixels. For galaxy density, this can be overcome by filling the empty pixels (within SDSS footprint) with a zero value, while masking everything outside the main footprint. However, this solution will not work for the quenched fraction, as zero-valued pixels would represent fully star-forming regions, and masking is not an option for already mentioned reasons. I therefore employ adaptive smoothing (see below) as a solution to this problem.

The downside of smoothing (adaptively or not) is that the power will be suppressed, or ‘smeared’ out, on scales smaller than the kernel size. This is illustrated in Figure A.2 where I compare an un-smoothed power spectrum to one which is derived from an adaptive SPH-smoothed map. To indicate the scale at which smoothing has a significant effect on the measurement (which I designate as a difference of 50%), I plot a black, solid line. However, I would like to stress that, although the power spectra are significantly affected on small scales, the comparison between the simulations and observations, which have both been smoothed in an equivalent way, is still valid even on small scales.

2.2.2.1 Galaxy density

Galaxies in a selected sample are smoothed in two-dimensional RA/DEC space with an adaptive smoothing kernel (SPH smoothing), the size of which is determined by the distance to $N_{\text{sph}}^{\text{th}}$ nearest neighbour. (The same scheme is used in numerical simulations to derive 3D density estimates of particles.) I choose $N_{\text{sph}} = 20$ and 10 for spectroscopic and photometric samples, respectively. (Through experimentation, I have found that these are approximately the minimum values that I can adopt for the two selections whilst retaining reliable estimates of the auto- and cross-spectra.) The same values are used for simulated analogues. This allows for sparsely populated regions to be filled in

with low density values, while dense regions are not over-smoothed so that the small-scale signal is preserved. SPH smoothing is described in more detail in Appendix A.1. The smoothed density field is then projected onto a flat grid spanning the full extent of SDSS DR7 main survey footprint, i.e. $RA = [100, 280]$, $DEC = [-20, 80]$ degrees. The resolution of this grid is twice that of the `HEALPix` pixels, i.e. 0.858 arcmin; this is done to ensure that all `HEALPix` pixels in the footprint are sampled and a contiguous footprint is obtained when I map the flat grid onto a `HEALPix` map.

The projection from the flat grid onto a `HEALPix` map is done by assigning pixel centre coordinates to the closest pixel centre in `HEALPix` via the inbuilt `ang2pix` function. Square pixel values which are assigned to the same `HEALPix` are summed together. On average, ~ 4.6 square pixels are assigned to one `HEALPix` pixel, however, 1.5% of `HEALPix` pixels are singly-occupied. So coverage is far from uniform due to geometry. The total number of galaxies is conserved at all stages of pixelisation⁸.

2.2.2.2 Quenched fraction

As galaxies in the total and quenched samples are expected to have different clustering properties and, therefore, different SPH smoothing kernels if smoothed separately, it is necessary to take measures to keep the kernel consistent between maps. For this purpose, I construct a map of quenched flags that is smoothed simultaneously with the map for the total galaxy sample. Each quenched galaxy is assigned a binary flag (1 = quenched or 0 = star – forming) which forms the basis of my quenched fraction maps. Regions with high density of galaxies will be smoothed with a relatively small kernel, averaging the binary flags to a number between 0 and 1. A small kernel preserves the compact regions of highly-clustered, quenched galaxies (in contrast to a fixed-size Gaussian kernel) resulting in a high quenched fraction value, which is where most of the signal is expected to originate.

2.2.3 Masking

In addition to the regions of the sky not observed as part of SDSS (which is the main limiting survey in terms of area coverage in this study), two other masks are applied to

⁸All maps and masks will be made available for download at: http://www.astro.ljmu.ac.uk/env_cor/

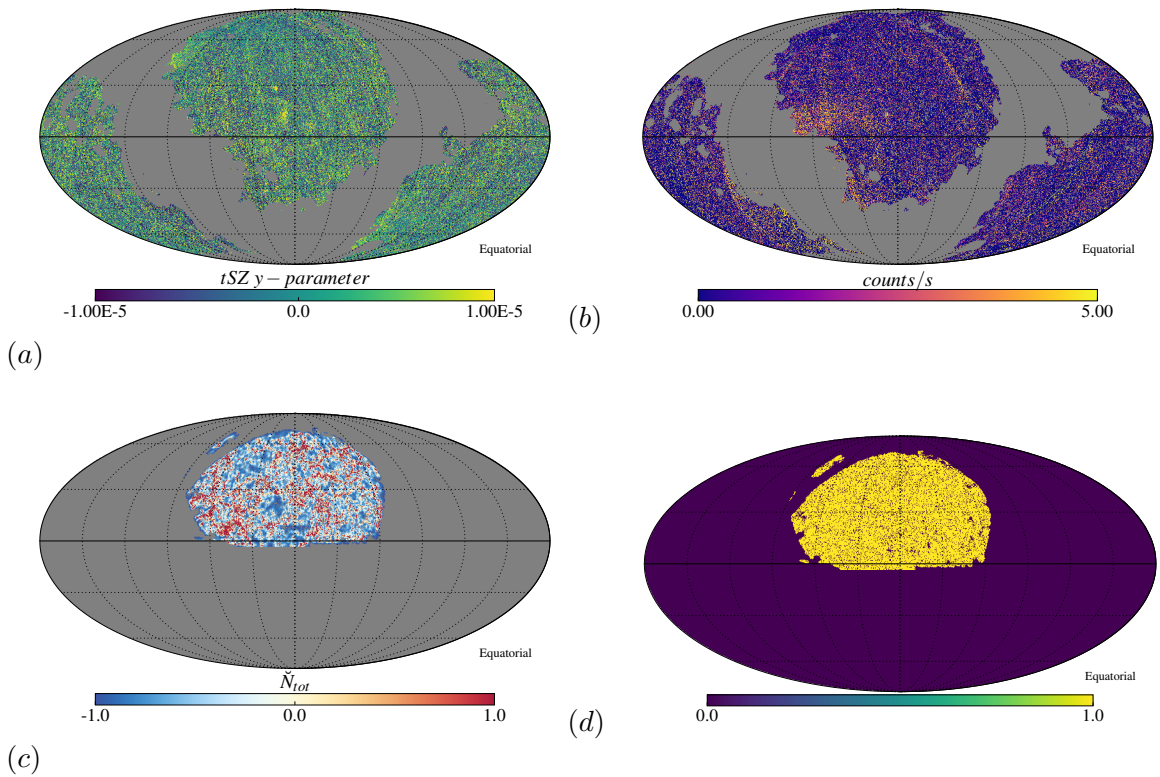


FIGURE 2.6: (a) A map of mean-subtracted thermal Sunyaev-Zel'dovich y parameter as produced by the MILCA component separation algorithm from *Planck* 2015 data. Milky Way and point source masks have been applied (grey regions). The map has also been transformed into equatorial coordinates and rotated so that the region overlapping with SDSS footprint is in the centre. (b) A map of mean-subtracted X-ray flux in counts/s from RASS. Milky Way and point source masks have been applied (grey regions). The map has also been transformed into equatorial coordinates and rotated so that the region overlapping with SDSS footprint is in the centre. (c) A map of SPH-smoothed SDSS galaxies in the spectroscopic sample (galaxy overdensity), projected onto the HEALPix sphere. The final 7849 square degree footprint is as a result of a combination of the Milky Way, point source, and main SDSS footprint coverage. (d) Final mask resulting from the combination of Milky Way, point source, main SDSS footprint, and zero exposure by ROSAT masks. This mask gives a sky fraction of 0.16 and is used in all correlations for respective galaxy samples in this paper. These and other maps are made publicly-available at http://www.astro.ljmu.ac.uk/env_cor/.

the maps prior to computing power spectra. First, I use the same Galactic mask⁹ as was used by the *Planck* collaboration (Planck Collaboration et al., 2014d), specifically, the $\sim 40\%$ sky Galaxy mask (M1), combined with the *Planck* point-source mask (M5), yielding a sky fraction of $\sim 58\%$. I then combine this mask with the main SDSS footprint, which covers $\sim 19\%$, giving a common area of $\sim 16\%$ of the sky. The *Planck* galactic+point source mask, therefore, does not have a significant effect on the total available sky fraction.

Finally, I also mask regions which were not observed by ROSAT as part of the RASS (i.e., those fields that have an exposure time value of zero). This masks a further 0.02% of the sky, which is negligible in terms of area but is necessary to avoid non-physical values. The final mask used in this study can be seen in Figure 2.6 (panel *d*). This mask is consistently applied to all maps followed by mean-subtraction - to ensure it remains zero. A map of galaxy overdensity for the SDSS spectroscopic sample is presented in panel (*c*) of the same figure.

2.2.4 Thermal Sunyaev-Zel'dovich effect map

The thermal Sunyaev-Zel'dovich (tSZ) effect (Sunyaev & Zeldovich, 1969) is a spectral distortion of the otherwise black-body CMB spectrum which is due to the inverse-Compton scattering of CMB photons by hot, free electrons (e.g., in the intracluster medium). The resulting change in intensity is directly proportional to the integrated line-of-sight electron gas pressure. The (frequency-independent) amplitude of the tSZ effect is characterised by the dimensionless Compton y parameter:

$$y = \frac{\sigma_{\text{T}}}{m_{\text{e}}c^2} \int P_{\text{e}} dl, \quad (2.4)$$

where $P_{\text{e}} \equiv n_{\text{e}} k_{\text{B}} T_{\text{e}}$ is the electron pressure (i.e., the product of the electron density and temperature); σ_{T} , m_{e} , and c are the conventional constants for Thomson scattering cross section, electron mass, and speed of light, respectively.

I use the publicly-available *Planck* 2015¹⁰ MILCA (Planck Collaboration et al., 2016b) tSZ effect map and masks. The data is stored in the HEALPix (Górski & Hivon, 2011) format of NSIDE = 2048, so there were minimal adjustments made before cross-correlating.

⁹COM_Mask.Compton-SZMap.2048.R2.00.fits

¹⁰<http://pla.esac.esa.int/pla/>

This map was shown to recover known clusters with a mean signal-to-noise ratio (S/N) of ~ 10 . In order to minimise radio continuum contamination from the Milky Way, a Galactic mask is used to mask 40% of the sky, with an additional $\sim 2\%$ covered by the radio point-source mask, as mentioned previously. A masked version of the tSZ effect map is presented in Figure 2.6 (panel *a*).

2.2.5 X-ray map

The hot gas in and around galaxies and groups and clusters of galaxies emits radiation at X-ray wavelengths via thermal bremsstrahlung and recombination lines. The observed X-ray surface brightness depends on the line-of-sight integrated electron density squared, as:

$$S_X = \frac{1}{4\pi(1+z)^4} \int n_e^2 \Lambda(T_e, Z) dl, \quad (2.5)$$

where z is redshift, n_e is the electron number density, and $\Lambda(T_e, Z)$ is the emissivity (or cooling function) in the relevant energy band, which only weakly depends on the temperature and metallicity (Z) of the gas for the energy range of 0.4 – 2.4 keV (Adam et al., 2017) sampled by ROSAT. Note how rapidly the surface brightness drops off with redshift ($S_X \propto (1+z)^{-4}$), making individual system analysis prohibitive for anything other than the most nearby objects and/or the most massive clusters.

I must point out the differing dependencies of the X-ray emission and tSZ effect on the properties of the hot gas (density and temperature). In principle, examination of cross-correlations between galaxy properties and these two observables should therefore help to break degeneracies in environmental processes (e.g., ram pressure stripping does not depend on the gas temperature).

Full-sky X-ray observations are derived from the ROSAT All-Sky Survey (RASS)¹¹ (see Voges 1993 and Voges et al. 1999 for survey description). The survey was completed using the Position Sensitive Proportional Counter (PSPC) instrument aboard the ROSAT satellite in ‘scanning’ mode and corrected for absorption by interstellar dust. The original RASS data is organised into 6.4×6.4 degree fields of the sky, which have been assembled into a full-sky map¹² in world coordinate system (WCS). This full-sky map has been conveniently projected into HEALPix format by the Centre d’Analyse de Données

¹¹<https://heasarc.gsfc.nasa.gov/docs/rosat/rass.html>

¹²<http://www.xray.mpe.mpg.de/cgi-bin/rosat/rosat-survey>

Etendues (CADE)¹³. The process is based on the drizzling library `Drizzlib`¹⁴ and the technical aspects of transforming from the WCS to `HEALPix` format are described in Appendix A of Paradis et al. (2012). This method guarantees photometric accuracy of the transformation with minimal data loss during the conversion from a local WCS FITS map to `HEALPix` format. I use the 0.4 – 2.4 keV count map (`RASS_IM2_1_2048.fits`) together with the exposure time map (`RASS_EXP_1_2048.fits`). To obtain a map of count rate, I simply divide counts by exposure and multiply by 5.24559 (see CADE website) to obtain the photon rate per `HEALPix` pixel, masking regions where exposure is equal to zero.

As already noted, in order to make consistent comparisons between the two environment maps (tSZ effect and X-ray), the zero-exposure X-ray mask is combined with the Galaxy and point-source masks from *Planck*. A map of RASS with the Galaxy, point source, and zero-exposure masks applied can be seen in Figure 2.6 (panel *b*).

2.2.6 Simulations

In order to decompose my detected signals and relate them to three-dimensional processes acting on single cluster scales, I utilise two different cosmological hydrodynamic simulations. Comparison to simulations is also beneficial in guiding my observational analysis, since the simulation-based maps do not have masks or incomplete coverage and allow us to experiment with the effects of smoothing and noise.

2.2.6.1 EAGLE

EAGLE (Evolution and Assembly of GaLaxies and their Environments, Schaye et al. 2015; Crain et al. 2015) is a suite of cosmological, hydrodynamic simulations designed to study the formation and evolution of galaxies at high resolution in a moderate-size box. Its standard (Ref-L100N1504) scientific run consists of a $L = 100$ comoving Mpc (cMpc) (on a side) box, with $N = 1504^3$ collisionless dark matter particles and an equal number of baryonic particles. The simulations were carried out with a modified version of the Lagrangian Tree-SPH code `GADGET3` (last described by Springel 2005). EAGLE adopts a *Planck* 2013 (Planck Collaboration et al., 2014a) Λ CDM cosmology.

¹³<http://cade.irap.omp.eu/dokuwiki/doku.php?id=rass>

¹⁴<http://cade.irap.omp.eu/dokuwiki/doku.php?id=software>

The EAGLE simulations include a number of subgrid treatments of processes that cannot be directly resolved in the simulations, including metal-dependent radiative cooling, star formation, stellar evolution and mass-loss, BH formation and growth, and stellar and AGN feedback. The efficiency of stellar feedback was calibrated to approximately reproduce the local galaxy stellar mass function and the sizes of galaxy discs, while the efficiency of AGN feedback was calibrated to reproduce the present-day scaling relation between the stellar mass of galaxies and that of their central BH (for further details see Crain et al. 2015). The feedback was not calibrated to reproduce the hot gas properties of galaxies or groups and clusters. It was recently shown in Davies et al. (2019) that simulations tend to predict X-ray luminosities that are somewhat too high compared to those observed for optically-selected samples of galaxies, whilst the simulations reproduce the large-scale tSZ effect flux well. I return to this point later.

2.2.6.2 BAHAMAS

The BAHAMAS (Baryons and Haloes of MAssive Systems, McCarthy et al. 2017, 2018) project is another set of cosmological, hydrodynamic simulations based on GADGET3, however its focus is on large-scale structure cosmology. As such, the simulations consist of much larger volumes but at significantly lower resolution than EAGLE. Specifically, the simulations primarily consist of 400 cMpc/h periodic boxes containing 2×1024^3 particles (dark matter and baryonic, in equal numbers). Here I use the run that adopts a WMAP9 (Hinshaw et al., 2009) cosmology with massless neutrinos.

BAHAMAS includes subgrid treatments for all of the same processes mentioned above for EAGLE, though with somewhat different parametrisations. The subgrid models were developed as part of the OWLS project (Schaye et al., 2010). The parameters governing the efficiencies of AGN and stellar feedback were adjusted so that the simulations reproduce the observed galaxy stellar mass function for $M_* \geq 10^{10} M_\odot$ and baryon content of groups and clusters, as dictated by the gas mass fraction–halo mass relation from high-resolution X-ray observations. As shown in McCarthy et al. (2017), the simulations match the local X-ray and tSZ effect scaling relations of galaxies and groups and clusters.

2.2.6.3 Light cones and simulated map making

To make like-with-like comparisons between the observations and simulations, equivalent maps of X-ray, tSZ effect, and galaxy properties from the simulation data are required. This is achieved by constructing light cones, extending from a point (the simulated observer) in a single direction of the simulation. A single simulation box does not have the required depth to accommodate a light cone out to $z = 0.15$ ¹⁵, so several simulation boxes have to be stacked in a line. Ideally, different sets of simulations would be used for each box constituting a light cone, however this would be too computationally expensive. As a compromise, randomly translated/rotated/reflected snapshots of the same simulation (but different redshift) are used. This minimises box-to-box correlations, so that they appear to be independent realisations of the simulated universe. Further detail on light cone construction can be found in McCarthy et al. (2018) (see also da Silva et al. 2000).

Ten different light cones are constructed for each simulation, representing 10 different lines of sight, allowing for an estimate of cosmic variance. Due to BAHAMAS having a simulation volume that is significantly larger than EAGLE, the field of view per light cone is 25 degrees on a side for the former, and 10 degrees for the latter. Cone-to-cone variance is, therefore, greater for EAGLE than BAHAMAS.

The desired quantities within a light cone now need to be projected and mapped onto 2D pixels. I follow McCarthy et al. (2018) when making the simulated tSZ Compton y maps. The parameter is computed directly from the properties of gas particles, first performing the integral in equation 2.4 and dividing that by the area of a pixel at the angular diameter distance of the particle (see McCarthy et al. 2018 for details).

To compute X-ray maps, I first compute high resolution ($dE = 2$ eV) synthetic X-ray spectra spanning the range 0.5-40.0 keV for each hot gas particle using the Astrophysical Plasma Emission Code (APEC; Smith et al., 2001) with updated atomic data and calculations from the AtomDB v2.0.2 (Foster et al., 2012). The spectrum of each gas particle is computed using the particle's density, temperature, and full abundance information.

¹⁵This is the maximum redshift of the galaxy catalogue, resulting in the maximum cross-correlation signal possible. Any additional depth in X-ray/tSZ maps (as is the case with observations) only results in additional noise on the cross-power spectra, this was tested by constructing light cones out to $z = 3$ and cross-correlating with a galaxy catalogue truncated to a lower z - noise increased with increasing discrepancy but the mean cross-power spectra did not change.

(Note that I exclude cold gas below 10^5 K which contributes negligibly to the total X-ray emission.) The spectra are appropriately redshifted using the redshift of the gas in the light cone and converted from intrinsic luminosity units into observed photon flux, taking account of cosmological dimming. As the RASS maps are provided in counts/s in the 0.4 – 2.4 keV, it is also necessary to fold in the instrumental response of the ROSAT PSPC instrument. To achieve this, I convolve the synthetic spectra with the effective area vs. energy function for the PSPC instrument, obtained from the WebPIMMS Count Rate Simulator¹⁶. This provides an estimate of the observed count rate in the 0.4-2.4 keV band as would be observed with ROSAT. I sum the contribution of each hot gas particle along the line of sight to the observed count rate, as was done for the tSZ effect.

When constructing the tSZ effect and X-ray maps, the initial native pixel size adopted was 10 arcseconds, which I rebinned to size of 1.7 arcmin in order to match the mean resolution of HEALPix pixels of tSZ effect and X-ray maps used in this study. The mapping of particles to a grid is done using a simple ‘nearest grid point’ interpolation method. Finally, the maps are smoothed with Gaussian beams of 10 arcmin for the tSZ effect maps for consistency with the *Planck* maps, and 1.8 arcmin for the X-ray maps for consistency with the PSF of the ROSAT PSPC instrument. In the case of *Planck* tSZ map, there is a dominant noise component contributing both positive and negative values in the map, whereas the physical signal only contributes positive values (due to photons being scattered only to higher energy levels). I, therefore, fit a Gaussian to the observed negative side of the pixel distribution, mirror it to the positive side, then sample this function to draw noise values for each pixel. This noise is then added to the simulated maps. Maps are mean-subtracted after noise addition to ensure a mean of zero, as for the *Planck* map.

In terms of galaxy catalogue-based maps, I select all galaxies in the light cones which have a mass exceeding $10^{10}M_{\odot}$ and lying within either $z < 0.06$ (spectroscopic) or $z < 0.15$ (photometric). Galaxies are defined to be either star-forming or quenched on the basis of their sSFR within a 30 kpc aperture, using the same threshold (10^{-11} yr^{-1}) as employed for the (spectroscopic) observations. Galaxies and their quenched flags are deposited into maps using the same SPH smoothing algorithm employed on the SDSS data. The value for N_{sph} (number of smoothing neighbours) was chosen such that a contiguous field is obtained with minimal smoothing. In this case, I use $N_{\text{sph}} = 20$

¹⁶<https://heasarc.gsfc.nasa.gov/cgi-bin/Tools/w3pimms/w3pimms.pl>

and 10 for spectroscopic and photometric samples, respectively, as was done for the observational data.

As I use overdensities/quenched fraction excess in the cross-correlations, it is necessary to compute the mean which I then use as a denominator in my calculations. This makes the normalisation of power spectra sensitive to the value of the mean. Given the degree of cosmic variance in the simulations, especially EAGLE, using the mean value of each light cone results in a substantially different normalisation of the power spectra. Since my goal is to use simulations in the interpretation of observed signals, it is desirable to have the power spectra as close as possible between simulations and observations. For this reason, I adopt the mean values of N_{tot} and f_q from SDSS maps (which are not limited by cosmic variance) and use them to compute excesses in the simulations.

2.3 Auto- and Cross-power spectra estimation

2.3.1 Formalism

I employ a two-point statistic to make a quantifiable measure of the correlation between two maps, each of which contains scalar quantities, using a quadratic estimator (Chiang & Chen, 2011), i.e.:

$$C_l^{jj'} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} j_{lm}^* j'_{lm} \equiv |j_{lm}| |j'_{lm}| \cos \Delta\phi_{lm}, \quad (2.6)$$

where j and j' are the two maps being considered, $C_l^{jj'}$ are the cross-power spectrum coefficients in multipole, l , space. $\Delta\phi_{lm}$ is the phase between j and j' , and takes values in the interval $[0, 2\pi]$; in a case where j and j' are the same signal, $\cos \Delta\phi_{lm}$ returns 1, and 0 if signals are uncorrelated. The effect is such that $\langle |j_{lm}| |j'_{lm}| \cos \Delta\phi_{lm} \rangle = 0$ in the case of spatially uncorrelated maps (where angle brackets indicate ensemble averages), otherwise it becomes a quantifiable measure of correlation between them. Note that if the two maps, j and j' , are identical, I obtain an estimate of the auto-power spectrum. If the maps differ, I estimate the cross-spectrum.

Values for j and j' could be computed directly from the maps if they were available for the entire sky, however, only some surveys have observed the entire sky and even then

there are foreground objects/contaminants which need to be masked. Masking has the effect of mode mixing, whereby eqn. 2.6 becomes an estimate of the biased *pseudo*-power spectrum and needs to be corrected for. The incomplete sky window-function can be represented as a position-dependent weighting with its own power spectrum:

$$W_l = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |w_{lm}|^2, \quad (2.7)$$

where w_{lm} are the spherical harmonic coefficients of the window function that is convolved with the underlying map of interest (Hivon et al., 2002). The spherical harmonic coefficients, in this case, take the following form:

$$\tilde{j}_{lm} = \int \Delta J(\hat{\mathbf{n}}) W(\hat{\mathbf{n}}) Y_{lm}^*(\hat{\mathbf{n}}) d\hat{\mathbf{n}} \approx \Omega_p \sum_p J(p) W(p) Y_p^*, \quad (2.8)$$

where $J(\hat{\mathbf{n}})$ is the scalar quantity captured in the map (and $\Delta J(\hat{\mathbf{n}})$ is its fluctuation from the mean), $W(\hat{\mathbf{n}})$ is the window function, and $Y_{lm}^*(\hat{\mathbf{n}})$ represents the spherical harmonics. In this equation I also write the expression for quantised maps where p represents a pixel, Ω_p is the pixel area, and the sum is over all data pixels in the map. The challenge is then to correct for the effects of partial-sky observations. Fortunately, there are existing publicly-available algorithms to do so.

NaMASTER¹⁷ (Hivon et al., 2002; Alonso et al., 2018) is an algorithm based on the direct spherical harmonic transform of the input maps. It operates entirely in the spherical harmonic space and, among many other functions, performs mask-correction by inverting the mode-mixing matrix relating the pseudo-power spectrum with the full-sky power spectrum estimator:

$$\langle \tilde{C}_l \rangle = \sum_{l'} M_{ll'} \langle C_{l'} \rangle. \quad (2.9)$$

In order to reduce windowing effects when performing the inversion, it is necessary to perform mask apodization prior to the computation. **NaMASTER** has multiple modes of apodization built in. For my combined mask, it was found that the mode ‘C1’ with apodization parameter $\theta_* = 0.04$ is the most optimal in recovering the full-sky power

¹⁷<https://github.com/LSSTDESC/NaMaster>

spectrum estimate. This mode involves multiplying all pixels by a factor f given by:

$$f(x) = \begin{cases} x - \sin(2\pi x)/(2\pi), & \text{if } x < 1 \\ 1, & \text{otherwise} \end{cases} \quad (2.10)$$

where $x \propto \sqrt{(1 - \cos\theta)(1 - \cos\theta_*)}$, and θ is the angular separation between the pixel and its closest masked neighbour.

Apodization was calibrated by making a simulated map from a known power spectrum, applying my mask, and demanding that the power spectrum be recovered within 1% error. For this purpose I made use of the best-fit Λ CDM CMB TT power spectrum provided by the *Planck* collaboration¹⁸, as it has many of the same statistical properties as the tSZ effect map. A full-sky map of $N_{\text{SIDE}} = 2048$ was generated and fed into the power spectrum estimator prior to computing all of the power spectra presented in this paper.

In terms of cross-correlations between hot gas (X-ray, tSZ) and galaxies, the methodology I have adopted, which is projecting galaxy surveys onto a HEALPix map for estimates of the Fourier-based cross-spectra with tSZ and X-ray data already in the HEALPix format, is fast becoming the standard practice (e.g., Makiya et al. 2018; Koukoufilippas et al. 2020; Pandey et al. 2020). However, given that the galaxies are treated as discrete objects, in principle one does not need to project the galaxies onto a map to analyse their clustering signal or cross-correlations with other signals. I have opted for a map-based approach not because there are obvious flaws with an object-based analysis, but mainly for convenience. Specifically, my map-based approach was motivated by: i) the available hot gas data I use (*Planck* Compton y and RASS X-ray counts) are in map form and, in the case of the *Planck* y map it has been smoothed to a fixed resolution of 10 arcmin; ii) from the point of view of galaxies, I am mainly interested not in individual properties but in ensemble quantities such as quenched fraction, so some spatial averaging is required (which is easily achieved within a map framework); iii) there are a wide variety of existing well-tested, publicly-available tools for analysing HEALPix maps (such as NaMASTER and PolSpice (Szapudi et al., 2001)); and iv) I can take advantage of existing software for projecting large cosmological hydrodynamical simulations onto maps (e.g., McCarthy

¹⁸<http://pla.esac.esa.int/pla/#cosmology>

et al. (2014, 2018)), thus allowing for a relatively straightforward comparison with the available data.

I expect that, if handled correctly, object-based and map-based analyses should converge on spatial scales above the pixel scale, which is where I limit my analysis to in the present study. It would be interesting to directly compare object-based and map-based approaches, but I leave this for future work.

2.3.2 Null-tests and error estimation

Given the noisy nature of the observational maps used here, there is always a possibility of obtaining a non-zero cross-power spectrum when there is no physical correlation. In order to ensure that my signals are real and not just driven by noise, I perform random rotations of one map relative to the other and compute power spectra for these combinations. Since no physical correlation is expected when the maps are rotated with respect to each other, any correlation that does persist is a result of noise. When rotating the maps, I ensured that galaxy/point source masks were fixed where necessary and a new combined mask was made prior to each computation. Each null power spectrum was inspected to check that it is consistent with zero over the entire ℓ range when binned in the same way as the signal power spectra. I indeed observe that each null power spectrum oscillates around zero, rarely having the same sign in several consecutive ℓ bins. I estimate the 1σ regions from 100 rotations for every correlation I compute and only plot the upper part of the region as all of my auto-/cross-power spectra are positive. The distance away from the null-test 1σ region provides a visual demonstration of the significance of any detection. The turquoise regions in Figs. 2.7 and 2.8 (below) represent the 1σ uncertainties in the auto- and cross-spectra as derived from the null tests.

I analytically estimate the statistical error bars on the observed and auto- and cross-spectra following the formalism of Tristram et al. (2005) (see also Hill & Spergel 2014; Hurier et al. 2015). For auto-spectra (\check{N}_{tot} and \check{f}_q) this involves computing (see eqns. 29-32 in Tristram et al. 2005):

$$(\Delta C_\ell^{jj})^2 = \frac{1}{f_{\text{sky}}} \frac{2}{(2\ell + 1)\Delta\ell} (C_\ell^{jj})^2, \quad (2.11)$$

where f_{sky} is the unmasked fractional area of the sky, $\Delta\ell$ is the width of a multipole bin centred on ℓ , and C_{ℓ}^{jj} is the power spectrum estimate.

Similarly, for the cross-spectra, the statistical errors are estimated using:

$$(\Delta C_{\ell}^{jj'})^2 = \frac{1}{f_{\text{sky}}} \frac{2}{(2\ell + 1)\Delta\ell} \left(C_{\ell}^{jj} C_{\ell}^{j'j'} + (C_{\ell}^{jj'})^2 \right), \quad (2.12)$$

where C^{jj} is auto-spectrum of the first map, $C^{j'j'}$ auto-spectrum of the second, and $C^{jj'}$ is the cross-power spectrum.

I note that the formalism of Tristram et al. (2005) was originally designed with the analysis of CMB maps in mind (including tSZ maps), as opposed to galaxy surveys. Therefore, as a check, I have also performed a ‘brute force’ estimation of the uncertainties by performing 100 random realisations of the galaxy density field under the assumption of Poisson statistics. That is, for each pixel I draw from a Poisson distribution whose mean is equal to the number of galaxies in that pixel. I generate 100 randomised realizations of the original SDSS galaxy density field in this way. For the galaxy density (auto-)power spectrum in particular, I find that the Tristram et al. (2005) and Poisson resampling uncertainties agree to typically better than a factor of 2 over the full multipole range, with both being ≈ 0.01 of typical power spectrum measurement (i.e., the galaxy power spectrum is very strongly detected). While the two methods of calculating uncertainties do not yield identical results (for undetermined reasons), none of the qualitative results or conclusions in this study are affected by my choice of error estimation technique. For specificity, I show the uncertainties calculated using the widely-employed Tristram et al. formalism throughout.

2.4 Results

I present my results in the form of panel plots in Figures 2.7 and 2.8, comparing the spectroscopic and photometric estimates of various auto- and cross-correlations side-by-side. This highlights the similarities and differences between the two samples as well as making it easier to spot changes in physical scale for different quantities.

In all cases, I restrict the multipole range between $100 < \ell < 2500$, which approximately corresponds to angular scales of $4.32 < \theta < 108$ arcmin¹⁹. On scales below $\ell \sim 100$ the observations become sparsely sampled and noisy, whereas the simulations reach their field of view limit leading to an abrupt truncation of the signal. At high multipoles, the observed and simulated power spectra are affected by the SPH smoothing kernel applied to the galaxy distribution and a beam present in the tSZ effect or X-ray observations, when computing cross-spectra involving those quantities. By $\ell \sim 2500$ these limiting factors are fully in effect and all power spectra smoothly decline towards zero with increasing ℓ . There is little information to be gained from examining such small scales, hence my limit.

I begin by plotting galaxy and quenched fraction auto-power spectra, followed by their cross-spectrum. Next, I introduce measures of gas environment by computing the tSZ effect/X-ray cross-spectra with galaxy overdensity and quenched fraction. These measure the connection between galaxy overdensity and hot gas pressure and density. The scales over which these quantities correlate indicates the angular scales over which the interplay between them acts. To further aid in this interpretation, I plot an approximate physical scale for these angular scales at the median redshift of the galaxy sample.

The cross-correlation of quenched fraction (as opposed to galaxy overdensity) with hot gas properties takes this one step further, by examining the star-forming state of galaxies. A statistically significant signal here would be the first time that a connection between hot gas properties and the quenched state of galaxies is measured directly without first selecting regions of the sky known to contain groups and clusters. This is important, as it is the local physical conditions that characterise the environment and not whether one has labelled that there is a group/cluster present.

2.4.1 Galaxy overdensity power spectrum

Panel (a) in Figure 2.7 shows the galaxy overdensity (i.e., auto- \check{N}_{tot}) power spectrum for the SDSS spectroscopic sample (navy points with error bars). Panel (b) shows the same quantity for the larger photometric sample. Note that the galaxy overdensity power spectrum is just the Fourier transform of the perhaps more familiar configuration-space (projected) two-point correlation function (2PCF). All power spectra in panels (a) and

¹⁹For reference, my pixel size is 1.716 arcmin and SDSS fibre angular resolution is 3 arcsec.

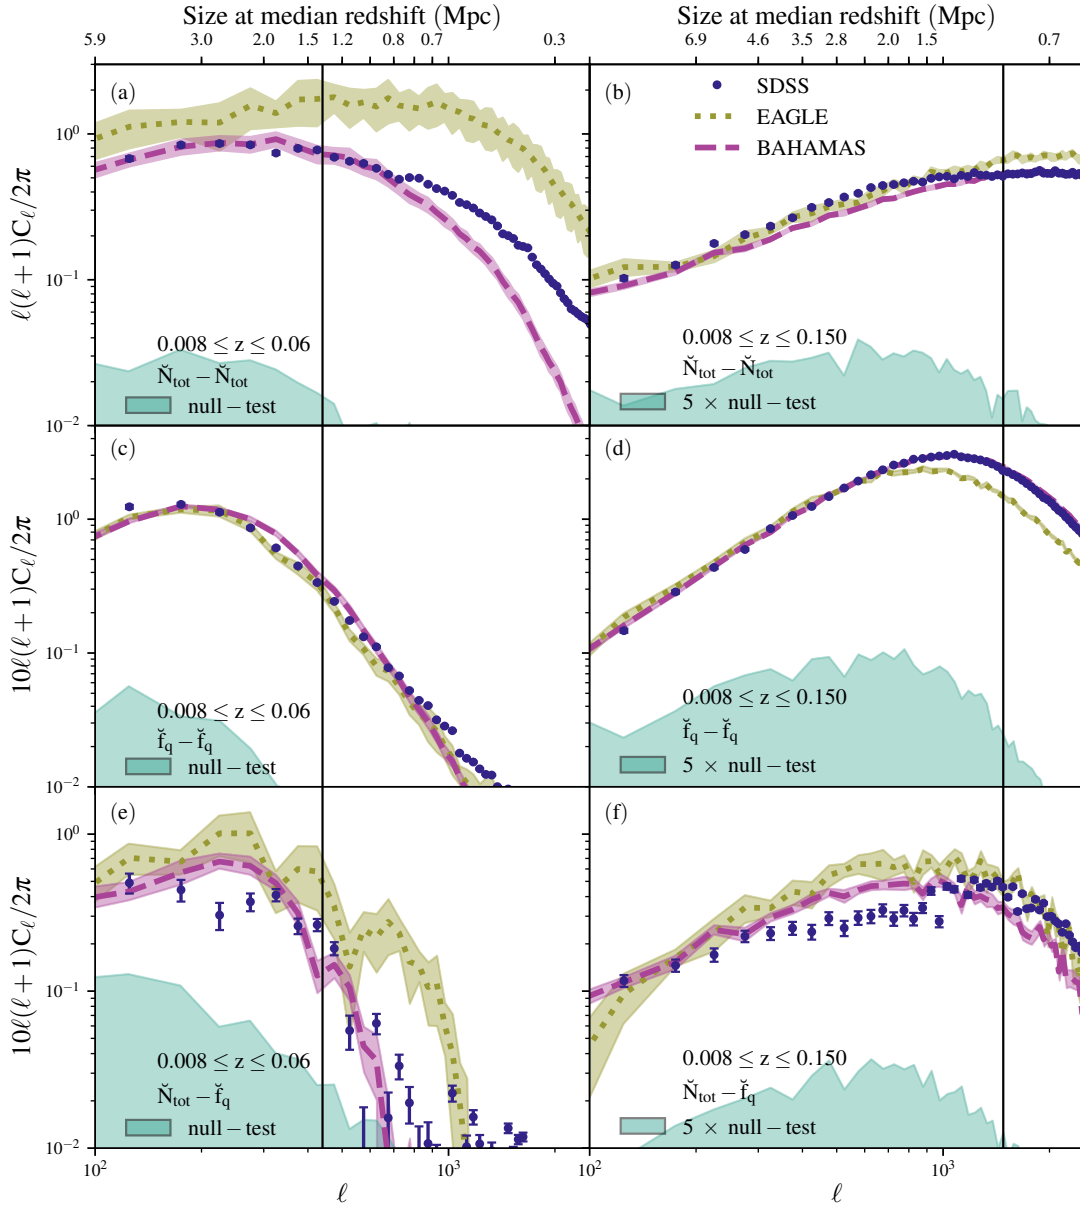


FIGURE 2.7: Auto and cross-power spectra involving the spectroscopic (*left*) and photometric (*right*) SDSS galaxy samples. Navy data points with error bars show the measured signal from the SDSS galaxy samples; pink, dashed curves and yellow, dotted curves show the predictions from BAHAMAS and EAGLE, respectively. Shaded regions indicate the 1σ confidence interval on the mean value of 10 light cones. Turquoise, shaded regions indicate the 1σ level of the null-signal (see text). Vertical, black lines indicate the scale at which SPH-smoothed power spectra deviate from their discrete counterparts by 50% (see Appendix A.1). The top x-axis indicates the approximate physical scale (transverse) assuming the median redshift of each sample. Top row (panels (a) and (b)) shows the total, shot noise-subtracted, galaxy overdensity power spectra; Middle row (panels (c) and (d)) shows the quenched fraction excess power spectrum; Bottom row (panels (e) and (f)) shows the quenched fraction-galaxy overdensity cross-spectra. Strong signals are detected in each case and are qualitatively consistent with the simulation predictions. In detail, the simulations reproduce the galaxy overdensity power spectrum and quenched fraction excess power spectrum remarkably well (particularly BAHAMAS), but both simulations predicted somewhat larger than observed amplitudes for the quenched fraction-overdensity cross-spectrum.

(b) have been shot noise-subtracted. (Please refer to Appendix A.1 for a discussion of shot noise estimation.) Also shown are the 1-sigma contours (turquoise shaded region) derived from the null tests (see Section 2.3.2) which give an additional estimate of the observational uncertainties due to noise, as are the predictions of the EAGLE and BAHAMAS simulations, for which I adopt the same selection criteria as in the observations. All power spectra (observed or simulated) are suppressed on small scales due to SPH smoothing. The scale at which this effect is 50% or more is indicated by a vertical solid, black line. Please refer to Appendix A.1 for more discussion of the effects of smoothing. Note that because smoothing is applied consistently to the observational and simulated maps, comparisons between the observations and simulations on scales smaller than the smoothing scale are still meaningful.

As expected, strong signals are detected on all scales that I sample, with the 1σ error bars being generally smaller than the data points and the null-test 1σ limit being at least an order of magnitude lower than the signal (and considerably larger than this for the photometric selection) on all scales. The observed power spectrum reaches a plateau at $\ell \sim 300$ (1000) for the spectroscopic (photometric) selections and declines thereafter due to smoothing effects. Note that differences are expected between the observed power spectra of the spectroscopic and photometric samples just on the basis that these are *angular* correlation functions and that the two samples have differing mean redshifts. A secondary effect is that the angular scale where the effects of smoothing become pronounced is reduced for the deeper photometric selection. This is just because the number of galaxies per pixel is increased and therefore the angular scale enclosing a fixed number of neighbours is decreased.

Both EAGLE and BAHAMAS reproduce the clustering of galaxies seen in the deeper photometric sample remarkably well. Qualitative agreement is also found for the spectroscopic sample comparison (i.e. similar shape and amplitude as the observed sample), although clear quantitative discrepancies can be seen between the observations and the two simulations. I attribute these differences to the larger degree of cosmic variance present in the simulations for the shallower (smaller volume) spectroscopic selection. (Note that, as the SDSS footprint is of *much* larger area than either the EAGLE or BAHAMAS light cones, the cosmic variance errors for the observed power spectra are negligibly small.)

The top panels of Figure 2.7 establish that galaxies in the simulations cluster in approximately the same way as in the real Universe. Note that this is only expected to be true if galaxies trace the correct haloes in the simulations (i.e., they have the correct stellar mass–halo mass relation, so that the selected galaxies have the correct bias with respect to clustering of matter) and the adopted cosmology is also broadly correct (so that the simulations have the correct matter clustering). The agreement of the predictions of the simulations with the observed galaxy overdensity power spectrum on small scales in particular may at first seem surprising, given the relatively large spread in predictions from hydrodynamical simulations for quantities like the total matter power spectrum at fixed cosmology (e.g., Chisari et al. 2019; van Daalen et al. 2020). However, it has been shown previously that many of the clustering statistics of galaxies can be reproduced relatively well by simple abundance matching techniques (see Conroy & Wechsler 2009 and references therein). As both EAGLE and BAHAMAS have been calibrated to reproduce the observed galaxy stellar mass function, the resulting stellar mass–halo mass relations from these simulations agree well with abundance matching methods (see Schaye et al. 2015; McCarthy et al. 2017). Consequently, the simulations should also reproduce the clustering statistics at least as well as abundance matching methods.

Since the spatial distribution of galaxies is correct, I can go further and ask whether the galaxies in the simulations respond to the environment in the same way as real galaxies.

2.4.2 Quenched fraction power spectrum

In panels (c) and (d) of Figure 2.7 I show the clustering of the quenched fraction (i.e., $\text{auto-}\check{f}_q$). Note that this is a measure of how quenched fraction itself clusters, independently of galaxy density. (Only the SPH smoothing kernel is common between maps of quenched fraction and overdensity.) Having said that, from previous studies the quenched fraction and galaxy density are known to be correlated quantities. It is therefore reasonable to expect that a similar correlation signal is observed here as in the top panels of Figure 2.7. Indeed, that is broadly the case; in both samples the power spectra exhibit a slow increase with decreasing angular scale until the power spectra turnover at small scales. Cosmic variance appears to be significantly reduced for \check{f}_q even in the spectroscopic case. Good agreement is obtained between the observed and simulated correlations, despite the fact that quenching is defined in terms of $(u - r)_{\text{obs}}$ colour

in the photometric sample and sSFR in the spectroscopic sample for the observations. (One might have worried that dusty red, but star-forming galaxies might have contaminated the colour-based quenched fraction at some level, but that does not appear to be the case.) BAHAMAS, which uses sSFR for both the shallow ($z < 0.06$) and deeper ($z < 0.15$) samples to determine quenched fraction, reproduces the observed clustering of quenched fraction remarkably well.

EAGLE exhibits an earlier turn off at $\ell \sim 600$ in panel (d) of Fig. 2.7, compared to $\ell \sim 1000$ for BAHAMAS and SDSS. Note that no such feature is visible in the EAGLE galaxy power spectrum at this scale (panel b), which rules out a difference in smoothing origin. We, therefore, conclude that this is a genuine issue. More generally, it is interesting that the clustering of quenched fraction (panels c and d) drops off at small scales faster than does the clustering of galaxies in general (panels a and b). As just mentioned, this cannot be a result of differences in smoothing, as the quenched fraction and galaxy overdensity have been smoothed in exactly the same way. Naïvely, one might have expected the opposite trend (i.e., that quenching becomes more prevalent on small scales). However, it should be kept in mind that the contribution of different types of systems can vary depending on the particular auto- and cross-spectrum being examined, as well as the scales under consideration. For example, the fact that the degree of cosmic variance in the simulations is relatively large for the galaxy overdensity power spectrum implies that it is dominated by relatively rare systems (e.g., clusters). The quenched fraction power spectrum, on the other hand, shows little variation from cone to cone (even for the spectroscopic selection using EAGLE), which strongly suggests that this correlation is dominated by relatively common objects (e.g., central galaxies near the lower mass limit of the sample). I discuss this further in Section 2.5.

2.4.3 Quenched fraction–galaxy overdensity cross-spectrum

As mentioned above, galaxy density is known to correlate with environmental quenching. Can I measure this correlation with my method? Panels (e) and (f) of Figure 2.7 show the quenched fraction–overdensity ($\check{f}_q - \check{N}_{\text{tot}}$) cross-spectrum. These strong detections confirm that, indeed, quenched fraction and galaxy overdensity are spatially correlated. The cross-spectrum has a very similar shape to the auto-spectra of its constituents: gradually rising at low multipoles, plateauing and then rapidly dropping to zero at small

scales. The scales at which the cross-power spectra turn over are intermediate to those seen the galaxy overdensity and quenched fraction power spectra. Interestingly, the degree of cosmic variance (cone-to-cone scatter) in the quenched fraction–overdensity cross-spectrum is significantly larger than for the quenched fraction power spectrum. This suggests that this correlation is picking out a population that is relatively rare (e.g., associated with massive systems). Indeed, I will show in Section 2.5 that this cross-spectrum is particularly sensitive to the quenching of satellite galaxies, whereas the quenched fraction power spectrum (auto-correlation) is significantly less so.

Relatively good agreement is obtained between observations and both of the simulations, although the overall shape and amplitude are by no means perfectly reproduced.

2.4.4 tSZ effect–galaxy overdensity cross-spectrum

Having established that a measurable signal can be obtained from correlations in galaxy properties alone, with overdensity being a commonly-used proxy for environment, I now turn my attention to direct environmental measures, specifically, the hot gas properties. Panels (a) and (b) of Figure 2.8 show the cross-spectrum between the tSZ effect y parameter and galaxy overdensity.

Since the *Planck* tSZ effect maps are convolved with a 10 arcmin beam, I do the same for the simulated tSZ effect maps. The scale of this beam, $\theta_{\text{FWHM}} = 10'$, corresponds to the angular scale of $\ell = 1080$, however, the effects become apparent on significantly larger scales. In the same way as for SPH smoothing effects, I indicate the scale at which beam-convolved power spectra deviate from beam-free by 50%. This occurs at $\ell = 670$. The small-scale decline in power is now dominated by *Planck* beam in the ‘photometric’ case.

An examination of the spectroscopic case in panel (a) reveals that even with a shallow galaxy sample a strong signal can be measured. Good agreement is achieved between the observations and BAHAMAS, whereas EAGLE predicts a slightly stronger correlation than is observed. Given the cosmic variance present in the spectroscopic case of $\check{N}_{\text{tot}} - \check{N}_{\text{tot}}$ in panel (a) of Fig. 2.7, it is reasonable to expect the same here. Indeed, the scatter between individual light cones is sufficiently large to account for the discrepancies between the observations and simulations. The power spectra are even biased in

the same way: EAGLE over-estimates the power on all scales, especially at high ℓ ; BAHAMAS agrees with observations for all but the smallest scales where it under-predicts the signal slightly.

Just as in the galaxy overdensity power spectrum ($\check{N}_{\text{tot}} - \check{N}_{\text{tot}}$) case, cosmic variance is greatly reduced in the photometric sample (panel *b*), where the measured cross-power spectra agree rather well between both simulations and observations.

2.4.5 X-ray–galaxy overdensity cross-spectrum

Panels (*c*) and (*d*) of Fig. 2.8 show the analogous cross-correlation for X-ray photon flux (as opposed to tSZ effect). While both the tSZ effect and X-ray flux are associated with the same hot gas, their connection to quenching could be quite different. However, there is no evidence for a qualitatively different correlation with galaxy overdensity. With the exception of normalisation, the cross-spectrum profiles look very similar for X-ray and tSZ effect.

EAGLE is once again significantly higher than the observations and BAHAMAS. EAGLE’s high amplitude, which is present in the photometric selection as well (and is therefore not due to cosmic variance), is expected in this case. It is already established that the AGN feedback present in EAGLE REF simulation is not sufficiently strong to remove an appropriate amount of gas from galaxy groups (see Schaye et al. 2015). Haloes, therefore, contain too much hot gas, which leads to excessive X-ray luminosities, as reported recently by Davies et al. (2019). While the ratio of EAGLE to observed X-ray luminosities (Anderson et al., 2016) at fixed halo mass is ~ 4 , it requires more complex modelling to introduce a correction factor into the cross-power spectrum with overdensity. As it is not the subject of this study, I simply report the measured signal as it is measured.

Despite the noisy nature of X-ray observations, strong detections are made for both samples. With the two simulations in hand, one of which agrees with observations while the other does not, it should be possible to decompose the signals and identify the physical cause (e.g., differences in feedback) that lead to these differences. This, in turn, should shed light on how exactly gas density/pressure are connected to the quenching of galaxies. This needs to be investigated in the future.

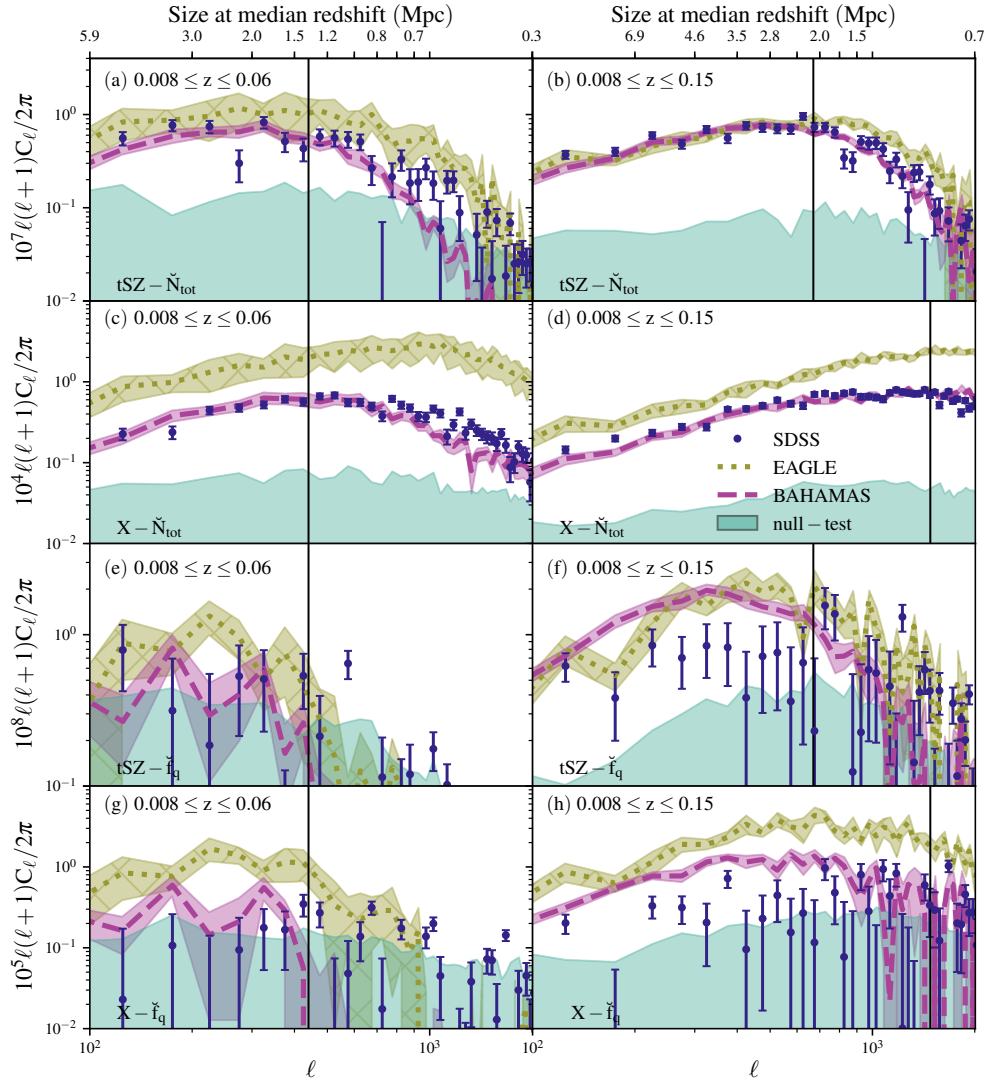


FIGURE 2.8: Cross-power spectra between galaxy properties and hot gas measures, for the spectroscopic sample (*left*) and photometric sample (*right*). Navy data points with error bars correspond to observed cross-spectra, dashed and dotted curves correspond to the predictions BAHAMAS and EAGLE, respectively. The shaded regions around the curves show 1σ confidence interval on the mean, while the turquoise regions at the bottom of each panel shows the 1σ level scatter around zero from 100 null-tests. Vertical, black lines indicate the scale at which SPH-smoothed power spectra deviate from their discrete counterparts by 50% due to either SPH smoothing or *Planck* beam effects (see Appendix A.1). The top x-axis shows the approximate physical scale assuming the median redshift of each sample. The top row (panels (a) and (b)) shows the tSZ effect–galaxy overdensity ($\text{tSZ} - \check{N}_{\text{tot}}$) cross-power spectra. The second row (panels (c) and (d)) shows the X-ray flux–galaxy overdensity ($X - \check{N}_{\text{tot}}$) cross-spectra. The third row (panels (e) and (f)) shows the tSZ effect–quenched fraction excess ($\text{tSZ} - \check{f}_{\text{q}}$) cross-spectra. The bottom row shows the X-ray flux–quenched fraction excess ($X - \check{f}_{\text{q}}$) cross-spectra. Strong detections are made for the hot gas–galaxy overdensity cross-spectra for both the spectroscopic and photometric selection. The hot gas–quenched fraction excess cross-spectra, on the other hand, are only well measured for the deeper photometric selection. Both EAGLE and BAHAMAS reproduce the tSZ–overdensity cross-spectra, but EAGLE predicts a larger than observed amplitude for the X-ray flux–overdensity cross-spectra (see text). Both simulations predict hot gas–quenched fraction cross-spectra that are higher in amplitude than observed, likely as a result of overly efficient quenching of satellite galaxies (see Section 2.5).

Note that the simulated X-ray maps contain emission from hot gas only, whereas the RASS X-ray map also contains point sources (e.g., AGN, X-ray binaries, stars, etc.) which I have not masked out. While X-ray AGN are more prevalent (by abundance) at high redshifts (e.g., Miyaji et al. 2001; Hasinger et al. 2005), they do exist locally as well and have been shown to spatially trace the normal galaxy population (e.g., Krumpel et al. 2012). I might therefore expect them to contribute to the observed X-ray cross-spectra and to potentially bias the comparison with the simulations, which do not model this effect. In Appendix A.2 I have explicitly checked the level of bias present in the recovered cross-spectra, by masking out AGN in the RASS point source catalogue and recomputing the observed cross-spectra. I find the level of bias present to be small (generally resulting in less than a 1-sigma change to the measurements at individual multipoles), such that none of the main conclusions of the present study are affected by neglecting their contribution.

2.4.6 tSZ effect–quenched fraction cross-spectrum

I present the cross-correlation power spectra between quenched fraction and tSZ effect y parameter in panels (e) and (f) of Fig. 2.8. Importantly, this correlation is independent of galaxy overdensity or whether a group/cluster is present. It is simply asking whether the quenching of galaxies knows about the local hot gas conditions. It is immediately evident that the signal-to-noise ratio of this cross-correlation is significantly lower than all the previous cases.

Visual inspection of panel (e) of Fig. 2.8 reveals that only the large scale contributions are (marginally) detected, with practically no signal present above $\ell \sim 500$. This is in line with all detected quenched fraction signals in this redshift bin; at that scale $\check{f}_q - \check{f}_q$ power spectrum in panel (c) of Fig. 2.7 is very much on the decline and $\check{f}_q - \check{N}_{\text{tot}}$ in panel (e) of Fig. 2.7 is similarly close to zero. The simulations more or less support this trend, although with a large degree of scatter.

Using the deeper photometric sample (panel (f) of Fig. 2.8), the tSZ effect–quenched fraction cross-spectrum is detected on most scales. The simulations produce similar correlations to each other over the entire range of scales, as in panel (b) of Fig. 2.8, however they both predict amplitudes that are too high relative to the observed tSZ effect–quenched fraction cross-spectrum. I discuss possible reasons for this difference

below, in Section 2.5. It is important to note, however, that this discrepancy is not due to the fact that observed maps contain clusters beyond the redshift limit of the galaxy catalogue and all simulated maps end at $z = 0.15$. Tests were carried out where simulated tSZ/X-ray maps extended much deeper than the simulated galaxy catalogue. The cross-correlation signal became more noisy (but still well below the cosmic variance level shown in the plots) but the overall amplitude did not change. Uncorrelated clusters beyond the redshift limit of the galaxy catalogue do not change the amplitude of the measured signal provided that the maps are properly mean-subtracted.

2.4.7 X-ray–quenched fraction cross-spectrum

Finally, I examine the X-ray–quenched fraction ($X\text{-ray}-\check{f}_q$) cross-power spectra in panels (g) and (h) of Fig. 2.8. The spectroscopic galaxy sample is not sufficiently deep to measure this signal for any analysis. While these detections are weak, a general behaviour of the correlation can still be seen, especially so for the deeper (photometric) of the two samples.

The trend of simulations overestimating the signal seen in panel (f) of Fig. 2.8 is also present here. EAGLE shows the same excess in signal relative to BAHAMAS as is seen in $X\text{-ray}-\check{N}_{\text{tot}}$ cross-correlation, which can be attributed to excessive X-ray luminosities in the former. Both simulations, however, predict cross-spectra that are in excess of what is observed.

2.5 Discussion: isolating external from internal quenching

As the simulations (particularly BAHAMAS) yield a reasonable match to the observed correlations, I can use them to gain some further insight into the physical drivers of the observed correlations presented above. I leave a detailed deconstruction of the auto- and cross-spectra for future work, commenting here only on the respective roles of internal and external quenching. In particular, thus far I have not made any distinction between central and satellite galaxies when dealing with sample selection, map making, or cross-correlation, in either the observations or simulations. This is partly due to the fact that this is a non-trivial task for observations, particularly those based on photometric redshifts. Here I note that upcoming large optical surveys (LSST, Euclid)

will be photometric only. However, I can easily separate simulated galaxies into centrals and satellites (as well as by a wealth of other available information) and see what this does to the predicted correlations.

In order to test the sensitivity of the measured signals to internal and external quenching, I explore the extremes of satellite quenched fraction. In particular, I artificially change the specific star formation rates of satellites in the BAHAMAS simulation such that they are either (1) all quenched; or (2) match the $f_q - \log_{10}(M_*)$ relation of centrals. (To achieve the latter, I randomly sample the sSFR distribution of central galaxies in a $\log_{10}(M_*)$ bin and assign sSFRs to a given satellite in the same bin.) This results in two samples where: (1) the environmental effects are maximally efficient (all satellites are quenched); or (2) they are non-existent (satellites are statistically the same as centrals). It is important to note that the population of centrals is unchanged in this process and the total galaxy density remains unaltered in the maps. To make the signals directly comparable to those previously measured, I have also used the same mean value f_q in computing the quenched fraction excesses. Thus, everything is measured relative to the default case presented in the previous plots. Fig. 2.9 shows a selection of cross-correlations taken from BAHAMAS where satellite quenched fraction has been modified as described above.

The top panel contains the quenched fraction power spectrum ($\check{f}_q - \check{f}_q$). The solid, green curve is the unmodified power spectrum from panel (b) of Fig. 2.7, the maroon, dashed curve is for central-matched- f_q (i.e., no external quenching), and coral, dotted curve represents the maximum environmental quenching case. While the quenched fraction power spectrum is sensitive in detail to external/environmental quenching (particularly on large scales), it is clearly mostly driven by internal quenching, as varying the external quenching in extreme ways only produces a relatively mild effect on the predicted power spectrum.

This behaviour is in strong contrast to the quenched fraction–galaxy overdensity cross-spectrum (middle panel), where the predicted signal is increased by a factor of ~ 2 (relative to the default case) when all satellites are quenched, and decreased by a factor of ~ 4 when environmental effects are completely absent.

Finally, in the bottom panel I show the predicted tSZ effect–quenched fraction cross-spectra for the different environmental quenching cases. This correlation is incredibly




FIGURE 2.9: The impact of changing the satellite quenched fraction to either fully quenched (‘max environment’) or so that it statistically matches that of centrals (‘no environment’) for BAHAMAS galaxies in the photometric sample. The solid curves represent the default (unmodified) power spectra shown in the previous plots. The dashed curves correspond to the case where all satellites are quenched and the dotted curves represent the case where satellite quenched fraction is adjusted to match centrals. *Top*: quenched fraction power spectrum. *Middle*: quenched fraction–overdensity cross-spectrum. *Bottom*: tSZ effect–quenched fraction cross-spectrum. The quenched fraction–overdensity and (particularly) the tSZ effect–quenched fraction cross-spectrum are particularly sensitive to the nature of satellite quenching, whereas the quenched fraction power spectrum (auto) is only mildly sensitive (being driven mainly by internal/mass quenching).

sensitive to the nature of external/satellite quenching: when satellite quenching is turned off the cross-spectrum is reduced by approximately an order of magnitude. The fact that the maximum external quenching is so similar (though not identical) to the default case, suggests that satellite quenching is particularly strong in the simulations. This is true for the EAGLE simulations as well (not shown). Since both BAHAMAS and EAGLE predict tSZ effect–quenched fraction cross-spectra that are in excess of the observed cross-spectrum from SDSS and *Planck*, this suggests that satellite quenching in the simulations, particularly in relation to the local hot gas properties, is too efficient. Interestingly, this conclusion seems consistent with the recent findings of Bahé et al. (2017b), who used the Hydrangea zoomed simulations of galaxy clusters (which were run with the EAGLE code) to examine the efficiency of satellite quenching with respect to the optical group catalogue-based findings of Wetzel et al. (2012) (see figure 6 of Bahé et al. 2017b). (See also Lotz et al. 2019 for similar conclusions based on the Magneticum Pathfinder simulations.) Ascertaining why the simulations are too efficient at quenching satellites should be a high priority.

2.6 Summary and Conclusions

The primary goal of this chapter is to establish a new methodology, based on spatial cross-correlations, for testing physical models of the environmental-dependence of galaxy evolution. In this regard, my study is very much a proof of concept. I have established that these auto- and cross-spectra can be detected and measured even in current data (and will therefore yield very strong signals in future surveys, e.g., LSST, Euclid, Simons Observatory, CMB-S4, eROSITA) and that state-of-the-art simulations yield a reasonable match to most, but not all, of the observed correlations. I also demonstrated that different power and cross-power spectra are sensitive to internal (e.g., AGN and stellar feedback) and external (e.g., ram pressure stripping, harassment, strangulation, etc.) quenching in different ways, allowing one in principle to constrain models for both simultaneously.

Below I summarise the main points:

- I constructed two volume-limited and stellar mass-complete samples from the SDSS. One is based on the DR7 spectroscopic sample ($z < 0.06$) and the other on

the DR12 photometric sample ($z < 0.15$) - see Fig. 2.3. Specific star formation rates (sSFRs) and colours were used to assign a star-forming/quenched status for the two samples, respectively (see Figs. 2.4 and 2.5).

- The SDSS samples were projected onto HEALPix images, to create maps of galaxy overdensity and quenched fraction overdensity. These are then used to compute angular power spectra and cross-spectra together with maps of the thermal Sunyaev-Zel'dovich effect from *Planck* and diffuse X-ray emission from the ROSAT All-Sky Survey. See maps in Fig. 2.6.
- I used the publicly-available software package **NaMaster** to compute the auto- and cross-power spectra from the HEALPix maps.
- Strong detections are reported for the auto- and cross-spectra involving galaxy properties only, including the galaxy overdensity power spectrum, the quenched fraction power spectrum, and the galaxy overdensity–quenched fraction cross-spectrum (see Fig. 2.7). Of these correlations, the galaxy overdensity–quenched fraction cross-spectrum is particularly sensitive to satellite quenching, whereas the quenched fraction power spectrum is considerably less so (see Fig. 2.9).
- Using synthetic observations of the EAGLE and BAHAMAS simulations, I show that, overall, the simulations reproduce the auto- and cross-spectra involving galaxy properties alone reasonably well, although they tend to predict a larger than observed amplitude for the overdensity–quenched fraction cross-spectrum. This suggests that satellite quenching may be too efficient in the simulations.
- Strong observational detections are also reported for the cross-spectra involving galaxy overdensity and either tSZ effect or X-ray surface brightness (see top panels of Fig. 2.8). The BAHAMAS simulations reproduce these cross-spectra well, whereas the EAGLE simulations predict a larger than observed amplitude for the X-ray–overdensity cross-spectrum, likely due to inefficient feedback on the scale of groups.
- I also report, for the first time, detections of the quenched fraction–tSZ effect and quenched fraction–X-ray flux cross-spectra (see bottom panels of Fig. 2.8). No information about galaxy overdensity or the presence of known galaxy groups/-clusters is used here.

- Both BAHAMAS and (particularly) EAGLE predict larger than observed amplitudes for the quenched fraction–tSZ/X-ray cross-spectra. As these cross-spectra are remarkably sensitive to the nature of satellite quenching (see Fig. 2.9), these results again suggest that the quenching of satellites in the simulations, particularly in relation to local hot gas properties, is too efficient in the simulations.

In a future study, I plan to examine the theoretical predictions in more detail, by deconstructing the power and cross-power spectra into contributions from, e.g., galaxy stellar mass, halo mass, redshift, central/satellite designation, host halo mass for satellites, and so on. This should yield further insight into the successes and failures of the simulations in reproducing the observed correlations reported here.

Finally, I point out that my methodology is not limited to linking galaxy quenched fractions to hot gas properties. One can easily substitute out quenched fraction for any galaxy-based property (e.g., a morphology-based quantity such as disk-to-total ratio, Sérsic index, or concentration) and/or substitute out hot gas properties for some other direct measure of environment, such as weak lensing potential. The advent of wide field galaxy and large-scale structure surveys now present us with a multitude of ways to directly link galaxies with their environments and this will only improve in the coming years with new surveys coming online. The increased statistics should also allow one to explore what I term ‘environmental tomography’ (in analogy to cosmic shear tomography), whereby the auto- and cross-spectra can be evaluated in redshift bins, to probe the redshift evolution of the correlations and also achieve a closer-to-3D view of the local physical environment. Such data, when compared carefully to simulations, holds the promise of developing a detailed physical picture for the environmental evolution of galaxies.

Chapter 3

Comparison of simulated cluster galaxies at $z \sim 1$

3.1 Introduction

Observational galaxy surveys of large, statistical samples have conclusively demonstrated that galaxies can be broadly classified into two populations (Strateva et al., 2001; Blanton et al., 2003; Baldry et al., 2004; Taylor et al., 2015). This bimodality was first observed at low redshift in the colour-magnitude diagram with one population having red colours, early-type morphologies (Wuyts et al., 2011; van der Wel et al., 2014), and old stellar populations (Kauffmann et al., 2004; Gallazzi et al., 2008) while the other showing blue colours, late-type morphologies, and young stellar populations. The same bimodality can also be observed in the specific star formation rate (sSFR)-stellar mass plane where the ‘blue cloud’ is strongly correlated with the ‘main sequence’ of galaxies with high star formation rates (SFR) and the red population has a low-SFR counterpart (Noeske et al., 2007; McGee et al., 2011; Wetzel et al., 2012). Since then, the same bimodality has been observed at higher redshifts in different colours (Baldry et al., 2006; Brammer et al., 2009; Whitaker et al., 2011; Muzzin et al., 2013). The presence of this bimodality suggests that galaxies do not spend very long transitioning between the two populations. The physical causes for the transition are likely to depend on the galaxy’s mass and location in the cosmic web and are the subject of much active research.

Peng et al. (2010) have demonstrated that the processes responsible for the blue-red transition can be separated into two categories: internal (caused by processes within the galaxy, such as stellar and AGN feedback) and external (induced by the environment a galaxy occupies). Even at $z = 0$, where observations of galaxies are the most complete, there is no consensus yet on exactly how the environment quenches satellites. There are many proposed processes involving hydrodynamic interactions between gas in the galaxy and its host, such as ram-pressure stripping (Gunn & Gott, 1972; Quilis et al., 2000; Barsanti et al., 2018) or strangulation/starvation (Larson et al., 1980; Balogh et al., 2000; McCarthy et al., 2008; Peng et al., 2015), or some relying on gravitational interactions alone, e.g. galaxy-galaxy mergers (Mihos & Hernquist, 1994a,b; Schawinski et al., 2014) and harassment (Farouki & Shapiro, 1981; Moore et al., 1999; Hirschmann et al., 2014).

Observations at higher redshift are significantly more difficult to obtain (particularly defining clean galaxy group/cluster populations) but data that is available at $z \sim 1$ suggests that quenching mechanisms in dense environments are different to those in the local Universe (Balogh et al., 2016; Kawinwanichakij et al., 2017; Papovich et al., 2018; van der Burg et al., 2018; Pintos-Castro et al., 2019). These mechanisms appear to be similar to those regulating the star formation rate of central galaxies in isolation, just acting faster.

Cosmological hydrodynamical simulations have made great strides in the last decade, they can now self-consistently solve the equations for cosmic evolution starting from Gaussian perturbations in density field and finishing with present-day, realistic-looking galaxies embedded in voids, sheets, filaments, and clusters (Schaye et al., 2015; McCarthy et al., 2017; Pillepich et al., 2018b). They allow the study of processes which would otherwise be impossible due to observational limitations, such as mergers, three-dimensional dynamics, gas interactions, etc. all as a function of time.

So far, simulations have been primarily concerned with ‘field’ galaxies, i.e. the ensemble average of galaxies in all types of environments. Relatively little attention has been paid to processes which take place exclusively in dense environments, specifically, the effect of feedback from star formation and Active Galactic Nuclei (AGN) on galaxies which are under heavy influence from their neighbours. It has already been demonstrated that calibrations which reproduce the observed field quantities do not simultaneously

reproduce the observed environmental influence on satellite galaxies, even at $z = 0$ (Bahé et al., 2017a). Given that environmental processes at $z \sim 1$ appear to act like those in the field but at a more rapid rate, it is possible that simulations which have been tuned to recover the field quantities may perform better at high redshift.

With better spectroscopic observations of dense environments at high redshift, such as the GOGREEN survey (Balogh et al., 2017), it will become possible to test the performance of hydrodynamic simulations at high redshift. If they do indeed simulate the processes correctly, then they will be a very useful tool in understanding the details of physical interactions taking place. Understanding the successes and limitations of hydrodynamic models at high redshift will lead to a better understanding of processes at present day since any deviation from reality is compounded as processes become more complex. At the same time, observational studies require more knowledge about the systems than can be obtained from data alone. Simulations can be used in, for example, testing cluster membership (see Ch. 4) and background subtraction schemes, calibrating property estimation techniques (e.g., velocity dispersion) and linking low and high redshift observations in a consistent way (e.g., via progenitor histories in Λ CDM).

Since stellar and AGN feedback are modelled quite differently in some simulations, we can hope to learn about the underlying physics by comparing a variety of different simulations. They all aim to reproduce the same observables at low redshift, but how do they compare where none of them have been tuned for that specific purpose? In this chapter I aim to perform this comparison and provide predictions that can be tested when observations become available. The structure of this chapter is as follows: In Section 3.2 I describe the three main models considered; in Sec. 3.3 I give a brief description of the GOGREEN survey used to perform some of the comparisons; in Sec. 3.4 I present the tests; and in Sec. 3.5 I summarise the findings and present the conclusions drawn.

3.2 Simulations

While the aim for them all is to obtain a close representation of the real Universe, different simulation groups model galaxy formation processes in a variety of ways and, consequently, the simulations differ in their detailed predictions and in how well they can reproduce current observations. As a result, noticeable differences are observable

in galaxy properties, especially in regimes which the simulations were not specifically calibrated to reproduce. In this section I briefly outline the features of each simulation code and describe their calibration strategies, as well as highlight any differences which would lead to discrepancies in observed quantities. Each simulation is accompanied by at least one release publication to which the interested reader is directed for more information.

All simulations used in this study share the way in which haloes and galaxies are identified from simulation particles. That is, objects such as haloes, subhaloes, and their properties are obtained by running friends-of-friends (FOF) and SUBFIND algorithms (Davis et al., 1985; Springel et al., 2001; Dolag et al., 2009) on snapshots at common timesteps in each simulation. All of the analysis presented in this chapter is done using catalogue-level data.

Although the cosmological parameters adopted by each simulation are different, this should not significantly affect the results of this study. I am only concerned about processes within massive haloes, where hydrodynamic and gravitational processes should dominate over any cosmological effects.

3.2.1 BAHAMAS / MACSIS

The BAHAMAS (BArYons and HaloEs of MAssive Systems, (McCarthy et al., 2017)) is a set of Smoothed Particle Hydrodynamic (SPH) simulations carried out using a significantly modified version of P-GADGET-3 (last described by Springel (2005)) and available in a variety of different cosmologies. In this study, I make use of the fiducial simulation which adopts the WMAP 9-year cosmology (Hinshaw et al., 2013). The simulations were performed in a periodic cube of length $L = 596 \text{ cMpc}$ and 2×1024^3 particles with masses of $\sim 4.45 \times 10^9 h^{-1} M_\odot$ and $\sim 8.12 \times 10^8 h^{-1} M_\odot$ for dark matter and baryons, respectively.

A number of sub-grid physics models are used for physics which cannot be resolved using simulation particles. Radiative cooling/heating rates are computed following (Wiersma et al., 2009a); star-formation follows the implementation of (Schaye & Dalla Vecchia, 2008); stellar evolution and chemical synthesis follows the model of Wiersma et al. (2009b); stellar feedback is implemented following the model of Dalla Vecchia &

Schaye (2008b); finally, AGN feedback follows the model of Booth & Schaye (2009). Parameters in these models were calibrated to reproduce the observed global galaxy stellar mass function (GSMF) and the amplitude of group/cluster gas mass fractions ($M_{\text{gas},500,x\text{-ray}}/M_{500,x\text{-ray}}$) at $z \sim 0$. The aim of the calibration was to ensure the simulations have the correct total baryon content in collapsed haloes, so that the simulations realistically capture the effects of baryons on the matter power spectrum $P(k)$, which is the basis of most large-scale structure tests of cosmology.

MACSIS (MASSive ClusterS and Intercluster Structures, (Barnes et al., 2017)) is an ensemble of 390 ‘zoom-in’ simulations centred on individual haloes drawn from a 3.2 Gpc N-body simulation. These re-simulations use the same hydrodynamic and sub-grid models, and were run at the same resolution as outlined above and the resulting haloes perfectly supplement the sample of haloes available in the main box, especially on the massive end.

The BAHAMAS model has been demonstrated to reproduce reasonably well the evolution of GSMF (McCarthy et al., 2017) and cluster scaling relations (Barnes et al., 2017). However, it appears to overquench low-mass galaxies (see Ch. 2 and Kukstas et al. 2019) in the local Universe.

3.2.2 EAGLE / Hydrangea

The EAGLE (Evolution and Assembly of GaLaxies and their environments) (Schaye et al., 2015; Crain et al., 2015; McAlpine et al., 2016) is another simulation built using an evolution of P-GADGET-3. The hydrodynamics solver differs from that used for BAHAMAS (which used the solver of Springel & Hernquist 2003), in that it uses the pressure-entropy SPH formalism of Hopkins (2013), artificial viscosity switch (Cullen & Dehnen, 2010), artificial conductivity switch (Price, 2008), and time-step limiter of Durier & Dalla Vecchia (2012).

EAGLE uses a similar set of sub-grid physics models as BAHAMAS for: cooling rates (Wiersma et al., 2009a), star-formation (Schaye & Dalla Vecchia, 2008), chemical synthesis (Wiersma et al., 2009b), stellar-feedback (Dalla Vecchia & Schaye, 2008b), and AGN feedback (Booth & Schaye, 2009). Where the two differ is in the calibration of model parameters: EAGLE is tuned to reproduce the global GSMF, galaxy-black hole

mass relation, and galaxy sizes. Several calibration instances are available, I choose ‘AGNdT9’ model because it provides a better match to observed gas fraction-total mass as well as X-ray luminosity-temperature relations. This suggests that the ICM is modelled more accurately, which is important for this study. EAGLE adopts cosmological parameters from Planck Collaboration et al. (2014b). I will refer to this model simply as ‘EAGLE’ throughout this chapter as it is the only variant used.

EAGLE AGNdT9 was calibrated on a 50 cMpc periodic box (with $N = 756^3$ and $m_{\text{baryon}} = 1.81 \times 10^6 M_{\odot}$ and $m_{\text{DM}} = 9.7 \times 10^6 M_{\odot}$), meaning that very few haloes above $\log_{10}(M_{200c}) \approx 14.0$ exist. For this reason, I supplement the sample with the Hydrangea (Bahé et al., 2017a) suite of zoom-in re-simulations using the AGNdT9 physics. It is part of the ‘Cluster-EAGLE’ project (Barnes et al., 2017). Hydrangea uses the same model and resolution as AGNdT9, allowing for the two to be combined seamlessly. These haloes are selected from the same 3.2 Gpc N-body simulation as MACSIS haloes.

3.2.3 Illustris TNG300

The TNG300 simulation, part of the IllustrisTNG project (Springel, 2005; Naiman et al., 2018; Nelson et al., 2018; Marinacci et al., 2018; Pillepich et al., 2018b), uses a radically different type of hydrodynamic treatment to those described above. It uses a deformable-mesh magneto-hydrodynamics (Pakmor et al., 2011; Pakmor & Springel, 2013) and gravity solver, AREPO (Springel, 2010). In addition to this, it implements sub-grid physical models for: gas radiative and metal-cooling, star-formation and dense interstellar medium, stellar population evolution and chemical synthesis, stellar feedback, as well as formation of black holes and the subsequent feedback mechanisms from them. A significantly more detailed description of the models can be found in the two methods papers: Pillepich et al. (2018a) and Weinberger et al. (2017). It adopts cosmological parameters consistent with Planck Collaboration et al. (2016a).

The TNG model was calibrated to reproduce several observed trends, including: the present-day GSMF, black hole mass–stellar mass relation, galaxy size–stellar mass relation, and gas fractions of galaxy groups. The simulation parameters were also adjusted to better reproduce the cosmic star formation rate density. While attempts were made to calibrate TNG on group gas-fractions, this was done on relatively small calibration volumes. As a result, the most massive clusters in TNG300 have somewhat too high

Name	L (cMpc)	N	$m_b (M_\odot)$	$m_{DM} (M_\odot)$
BAHAMAS	596	2×1024^3	1.16×10^9	6.36×10^9
EAGLE	50	2×752^3	1.81×10^6	9.7×10^6
TNG300-1	303	2×2500^3	1.1×10^7	5.9×10^7

TABLE 3.1: A comparison of the periodic-box simulations used in this study: length of the cubic simulation box in co-moving megaparsecs, number of simulation particles (baryonic and dark matter), and mass of each type of particle. Zoom-in simulations use the same resolution parameters.

gas fractions (Barnes et al., 2019). The standard TNG300(-1) simulation has 2×2500^3 particles, with $m_{\text{baryon}} = 1.1 \times 10^7 M_\odot$ and $m_{\text{DM}} = 5.9 \times 10^7 M_\odot$. The key aspects of each code are summarised in Table 3.1.

3.3 OBSERVATIONS

For the comparison in section 3.4.3, both cluster and field observational data used to compare to the simulations has been performed as part of the Gemini Observations of Galaxies in Rich Early Environments (GOGREEN, Balogh et al. (2017)). The survey targets 21 systems that cover the halo mass range: $5 \times 10^{13} \leq M_{200c} \leq 7.7 \times 10^{14}$ and span a range in redshift $1.0 < z < 1.5$. More specifically, I use the cluster and field samples presented in van der Burg et al. (2020). The cluster sample consists of 11 clusters with $M_{200c} > 1 \times 10^{14}$ (mean $\log_{10}(M_*) = 14.5$) with a mean redshift of $z = 1.23$. These clusters were observed with deep spectroscopy and multiband photometry providing a mixture of spectroscopic and photometric redshift estimates for galaxies down to stellar masses of $10^{9.5} \lesssim M_*/M_\odot$.

In order to isolate cluster-specific quenching effects on galaxies, a ‘field’ sample from the COSMOS/UltraVISTA survey is used as a control. The DR1 (Muzzin et al., 2013) catalogue consists of galaxies from a 1.69 deg^2 field, with photometric redshifts in the range of $1.0 < z < 1.4$, complete down to stellar mass of $10^{9.5} M_\odot$. Completeness corrections have been applied as per Sec. 4.2 of van der Burg et al. (2020). Both cluster and field galaxies were characterised into star-forming and quiescent based on their rest-frame U-V and V-J colours to take reddening by dust into account (see equation 6 of van der Burg et al. (2020)).




FIGURE 3.1: Halo mass distributions for all simulations from which clusters have been selected, and GOGREEN (green squares). Periodic box-based simulations are not sufficiently large enough to fully sample the GOGREEN cluster halo mass range. Supplementing BAHAMAS with MACSIS completely alleviates this problem.

3.4 RESULTS

In this section, I present a series of tests comparing stellar masses and star formation rates in clusters between all three simulations (and observations where available). All three implementations have been calibrated to reproduce the global ‘field’ galaxy stellar mass function at $z \sim 0$. No particular attention was paid to cluster-specific galaxy mass or SFR distributions. All three simulations have been shown to reasonably reproduce the observed field GSMF at higher redshifts. They may, however, differ in cluster environments.

As can be seen in Figure 3.1, the halo mass range probed by each simulation is very different at $z \sim 1$. BAHAMAS/MACSIS has an abundance of haloes at $\log_{10}(M_{200c}/M_{\odot}) > 14$, allowing for haloes to be binned and median value (as well as 1σ scatter region) to be plotted. TNG300 and Hydrangea, on the other hand, lack the most massive examples - these are always plotted as individual haloes. EAGLE does not have any haloes with

$\log_{10}(M_{200c}/M_{\odot}) > 14$ at all - it is used to compute field quantities only. The distribution functions are also very different. For these reasons, I perform different selections for each test in order to either minimise potential biases or match available observations. As haloes are not stacked, Sec. 3.4.1 uses all available haloes with $\log_{10}(M_{200c}/M_{\odot}) > 14$; Sec. 3.4.2 matches the mean halo mass between simulations for a consistent comparison; in Sec. 3.4.3, GOGREEN halo mass distribution is matched as closely as possible. However, in all instances, only galaxies with $\log_{10}(M_{*}/M_{\odot}) > 10$ are included in the analysis.

3.4.1 Stellar mass content of haloes

Before proceeding to select representative clusters, it is important to examine the level of agreement in the field galaxy properties between simulations. In this instance, any difference in stellar content of haloes in the field will influence the result found in clusters. Panel (a) of Figure 3.2 shows the stellar mass - halo mass relation for all three simulations and halo abundance matching results of Moster et al. (2018) and Behroozi et al. (2019) at $z \sim 1$. Stellar masses are measured within a 30 kpc aperture for all centrals with halo mass above $\log_{10}(M_{200c}/M_{\odot}) = 11.0$. All curves show the running median in bins of halo mass. At the low-mass end, TNG300 best matches the two literature results with the lowest stellar mass content. EAGLE and BAHAMAS predict higher stellar-to-halo mass ratios for relatively low-mass haloes, in that order. All curves peak at $\log_{10}(M_{200c}/M_{\odot}) \sim 12.1$ and decline with a constant logarithmic gradient. The 50 cMpc simulation box is known to be limited in the number of massive objects present, as a result, there are no haloes more massive than $\log_{10}(M_{200c}/M_{\odot}) = 13.0$ in the case of EAGLE. Between BAHAMAS and TNG300, there is a noticeable difference in the logarithmic gradient of each curve: BAHAMAS agrees well with the results of Behroozi et al. (2019) and has a very similar slope to Moster et al. (2018), whereas TNG300 shows a substantially less steep decline. High-mass, central galaxies in TNG300 are significantly more massive for the same halo mass.

Galaxies above $\log_{10}(M_{*}/M_{\odot}) = 11$ are known to extend beyond 30 kpc (Pillepich et al., 2018a), which says more about the applicability of the aperture than galaxy formation physics. However, stellar masses measured within an aperture of 30 kpc have been shown to match observations (Schaye et al., 2015). To demonstrate the effects of this, I plot

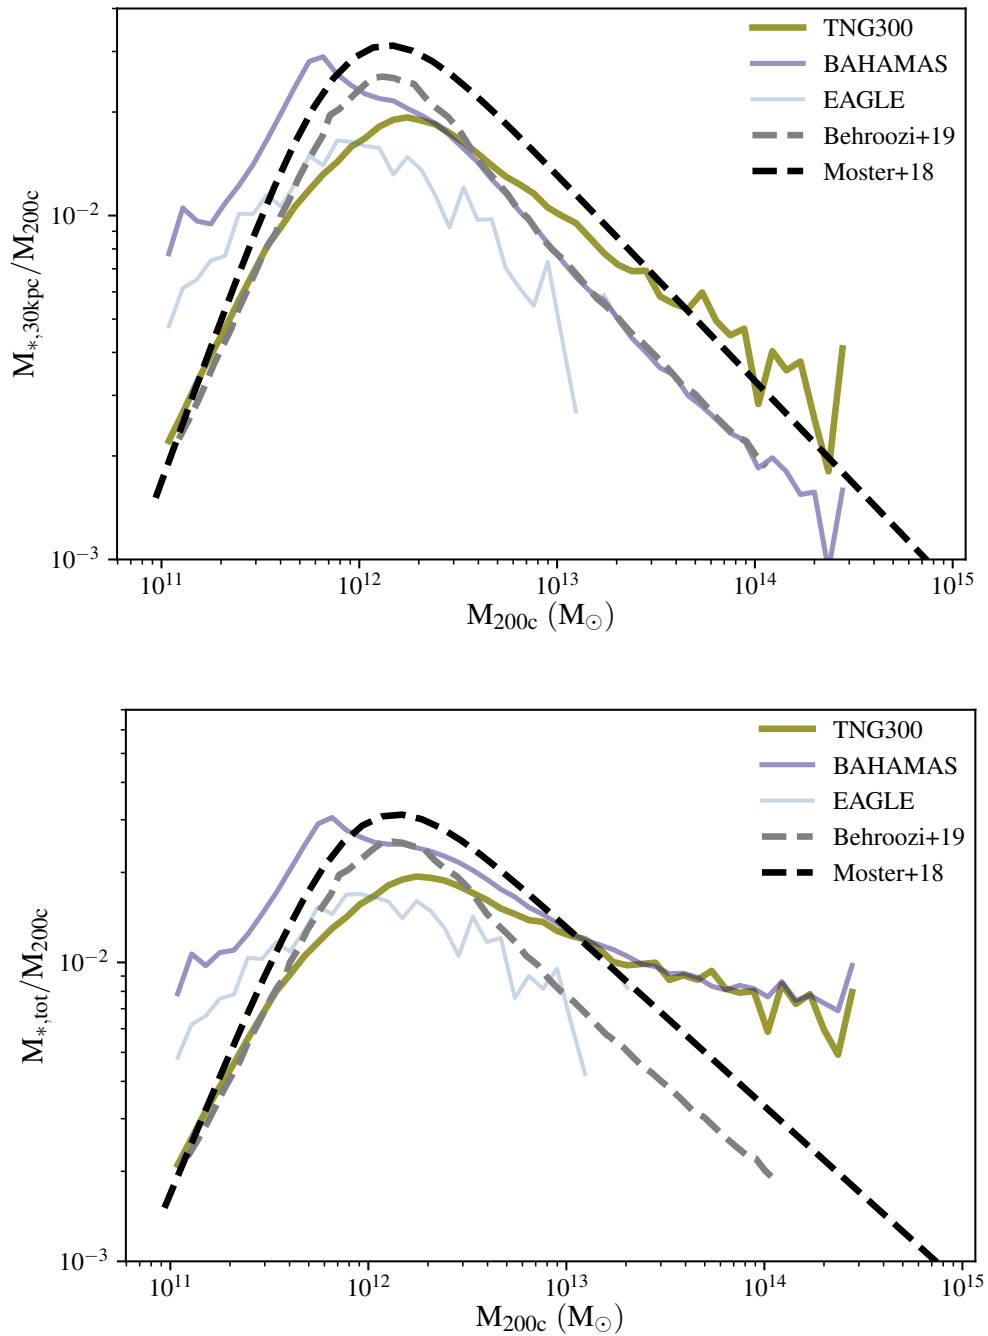


FIGURE 3.2: **Top:** Central stellar to halo mass relation, as a function of halo mass: all central galaxies in the simulation volume. Stellar mass is measured within a spherical aperture of $R = 30$ kpc. Solid lines show the three hydrodynamical simulations used here: yellow for TNG300, navy for BAHAMAS, and blue for EAGLE. Dashed lines represent abundance matching results from Moster et al. (2018) (black) and Behroozi et al. (2019) (grey). All three simulations differ quite substantially, but these differences are of the same order as literature results from abundance matching. High-mass centrals are the most massive in TNG300, while low-mass centrals are most massive in BAHAMAS. EAGLE’s low-mass centrals have intermediate masses, and the box size is too small to contain any of the most massive galaxies. **Bottom:** The same, but using the total subhalo stellar mass rather than aperture measurement. High-mass BAHAMAS galaxies now match TNG300, indicating that massive galaxies are less compact in BAHAMAS. Aperture measurements are more appropriate for comparisons to other methods but total measurements provide a more like-for-like comparison in this study. Abundance matching curves are the same as above and for reference only.

the same curves using the *total* stellar mass estimates in panel (b). Behroozi and Moster curves are carried over for reference. As expected, the low-mass end of this relation is unchanged by the aperture: all stellar particles are within 30 kpc of the centre. All three hydro simulation curves move to higher M_*/M_{200c} values at the high-mass end, but BAHAMAS is affected by it the most. This suggests that galaxies in BAHAMAS are less compact than TNG300, and vice versa. Stellar masses measured within a 30kpc aperture better match the observed stellar masses, as has been reported for $z \sim 0$.

I shall now examine cluster stellar mass content. For this section only, I select all haloes with $\log_{10}(M_{200c}/M_{\odot}) > 14$ and measure stellar mass content within a radius of R_{200c} . In panel (a) of Figure 3.3 I plot the total stellar mass against M_{200c} for all galaxies within R_{200c} of any given halo, when stellar mass is measured within a 30 kpc 3D aperture centred on the galaxy. TNG300 and Hydrangea more-or-less agree on the power-law relation between stellar content and halo mass with TNG300 having more haloes and, as a result, more scatter in this regime. BAHAMAS/MACSYS is very clearly offset by ~ 0.2 dex towards lower $M_{200c,*}$ at all halo masses. These haloes are hosts to galaxies of lower stellar masses or are less compact (i.e. 30 kpc aperture cuts out a significant part of stellar mass) relative to the other two simulations. I test this in panel (b) by using the total stellar mass estimate instead, which includes all star particles associated with the FoF group within R_{200c} . Here, the offset is much smaller (below 1σ scatter) which suggests that BAHAMAS/MACSYS galaxies are substantially larger for the same host halo mass and more of the mass resides in the intra-cluster medium (ICM), instead. This is consistent with our findings for all central galaxies in the simulation.

The compactness of galaxies relative to a fixed aperture should be most pronounced for the central ‘Brightest Cluster Galaxies’ (BCGs). In the lower row of Figure 3.3 I compare stellar mass estimates of centrals with a 30 kpc aperture (panel c) and without (panel d) as a function of halo mass. There is no noticeable difference between TNG300 and Hydrangea BCG stellar masses in either case, however BAHAMAS/MACSYS shows a ~ 0.3 dex offset towards lower stellar masses when only the central 30 kpc are considered, and a near-perfect agreement when total masses are used. BAHAMAS/MACSYS BCGs are significantly less compact, driving most of the discrepancy seen in panel (a).

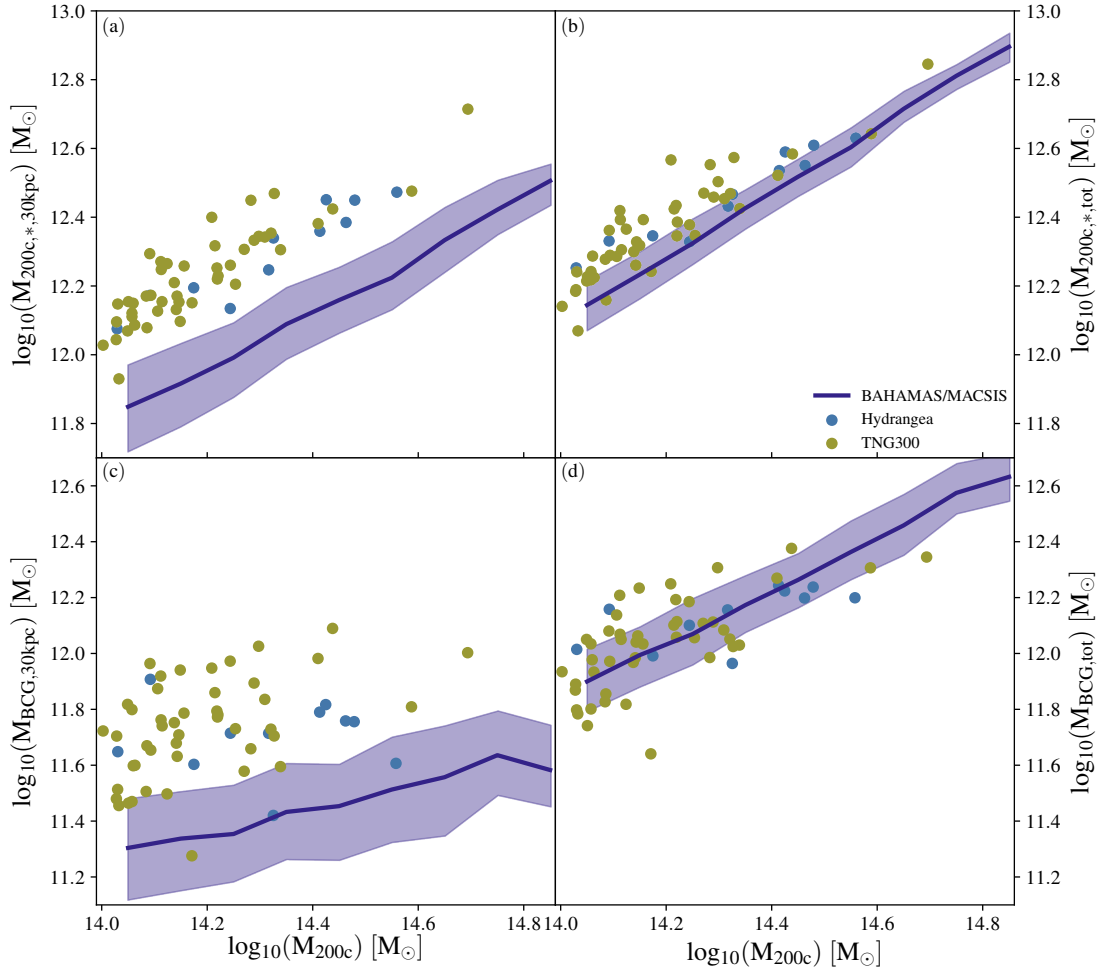


FIGURE 3.3: *Top row*: total stellar mass within R_{200c} using the 30 kpc spherical aperture estimate as a function of halo M_{200c} (left) and corresponding total stellar mass estimate without using an aperture (right). *Bottom row*: Stellar mass estimates of central galaxies with (left) and without (right) the 30 kpc aperture. All total stellar masses are higher than 30 kpc estimates but BAHAMAS/MAC SIS is particularly affected by it. Aperture effect on the central galaxies reveals that BAHAMAS/MAC SIS centrals are significantly less compact than TNG300 and Hydrangea.

3.4.2 Galaxy distribution as a function of radius

Numerous observational studies have found that the number and stellar mass density distributions in local massive clusters can be reasonably well described with a Navarro-Frenk-White (NFW) profile (Navarro et al., 1996), using a concentration parameter that is a factor of $\sim 2 - 3$ lower than the underlying mass density profile (Carlberg et al., 1997; Lin et al., 2004; Budzynski et al., 2012; van der Burg et al., 2015). BAHAMAS was shown to reproduce this observation (at $z \sim 0$) rather well (McCarthy et al., 2017). Once data is available, it will be interesting to find out whether this holds for higher

redshifts. Meanwhile, I can test whether the three simulations agree with one another, as well as investigate the profiles of star-forming and quiescent galaxies separately.

Since this test involves stacking clusters, it is necessary to ensure that the mean halo mass is consistent between the simulations. I choose to match the mean halo mass of Hydrangea because it has the fewest haloes with $\log_{10}(M_{200c}/M_{\odot}) \geq 14.0$ and excluding any would severely affect the number statistics. The upper or lower limits on halo mass selection for TNG300 and BAHAMAS/MACSYS are adjusted as necessary to obtain a mean $\log_{10}(M_{200c}/M_{\odot}) \approx 14.34$ for all three simulations. Given the results of Section 3.4.1, I choose to use total stellar mass estimates rather than aperture measurements, which means that they contain a substantial ICL component. In the case of radial distributions relative to cluster centre, however, galaxies are considered to be point sources. While it is customary to exclude the central galaxy in such plots in the literature, I include the BCGs in order to stay consistent throughout the chapter and with vdB20 (who include centrals in their GSMF estimates).

Figure 3.4 shows the number density (left column) and stellar mass density (right column) profiles in rows for all, star-forming, and quiescent galaxies, respectively. Examining panels (a) and (b) more closely shows that all three simulations produce similar profile gradients in both number and stellar mass. However, BAHAMAS/MACSYS shows a slight deficit of galaxies at all radii, while showing a much better agreement with TNG300 and Hydrangea in ρ_* . This is likely a manifestation of the same ICL component that was present in Fig. 3.3; the number of galaxies in BAHAMAS/MACSYS clusters is lower (especially at low radii) but the addition of ICL brings it almost in line with Hydrangea. All three simulations show an upturn in stellar mass density but not number density, which is likely due to the inclusion of central galaxies. TNG300 shows a slight excess of galaxies/stellar mass at intermediate radii of $R = 0.3r_{200c}$.

While observed galaxies are often divided into quiescent and star-forming by their location on the colour-colour diagram (for GOGREEN it is the U-V and V-J colour space) or spectral features, simulated galaxies are split up based on their specific star formation rates (sSFR). These quantities do not all result in identical classification as they represent somewhat different measures of star-formation activity. Computing galaxy colours is a considerable task, involving several models and assumptions. As this work is mainly a comparison between simulations, it is left as an exercise for future studies, when robust

data is available. In the meantime, sSFR is used to split the simulated galaxy populations. In all three simulations, there is a distinct population of star-forming galaxies (the Main Sequence, MS) and a long ‘tail’ of low-sSFR galaxies representing the quiescent population, many of which have $\text{sSFR} = 0$. The position of the MS need not be at the same sSFR for all three models. Furthermore, it need not be the same for cluster and field environments. Indeed, I show in Appendix B.1 that all three simulations have the MS peak at different sSFR positions, and that the peak also shifts to lower sSFR for cluster galaxies relative to the field. As a conservative solution, I construct the quiescent population from galaxies which are definitely not star-forming, i.e. their $\log_{10}(\text{sSFR}) \leq -12$. Galaxies with $\log_{10}(\text{sSFR}) \leq -12$ are assigned a value in the interval $[-13, -12]$ by sampling a uniform distribution. A value of $\log_{10}(\text{sSFR}) > -12$ is chosen as the line of demarcation between quiescent and star-forming galaxies in all three simulations.

I proceed to examine the star-forming and quiescent distributions separately. Panel (c) of Fig. 3.4 shows the star-forming number density. Hydrangea and TNG300 are in agreement and show constant gradient in log-space. BAHAMAS/MACSYS predicts a significantly lower abundance at all radii, except the centre where star-forming BCGs dominate, resulting in a sharp upturn. A similar trend can be seen in the stellar mass density in panel (d). The only exception is that TNG300 star-forming galaxies are more massive than Hydrangea at intermediate radii despite being almost identical in number.

Panels (e) and (f) show the quiescent galaxy number and stellar mass distributions. Contrary to previous panels, there are no significant differences between relative number and stellar mass density trends. There are, however, substantial differences between the three simulations. BAHAMAS/MACSYS has the most shallow of the three profiles, likely due to the abundance of star-forming centrals causing a down-turn in abundance of quiescent galaxies near the centre. Hydrangea has the steepest profiles. Indeed, there is a surprising absence of quiescent galaxies near the edges of Hydrangea clusters. In addition to this, the profiles are lower in magnitude overall (with the exception of very centres), consistent with significantly lower total quenched fractions in Fig. 3.6a. This may be an indication of AGN feedback not affecting the more distant galaxies at this stage. On the contrary, TNG300 shows the highest abundance of quiescent galaxies at all radii.

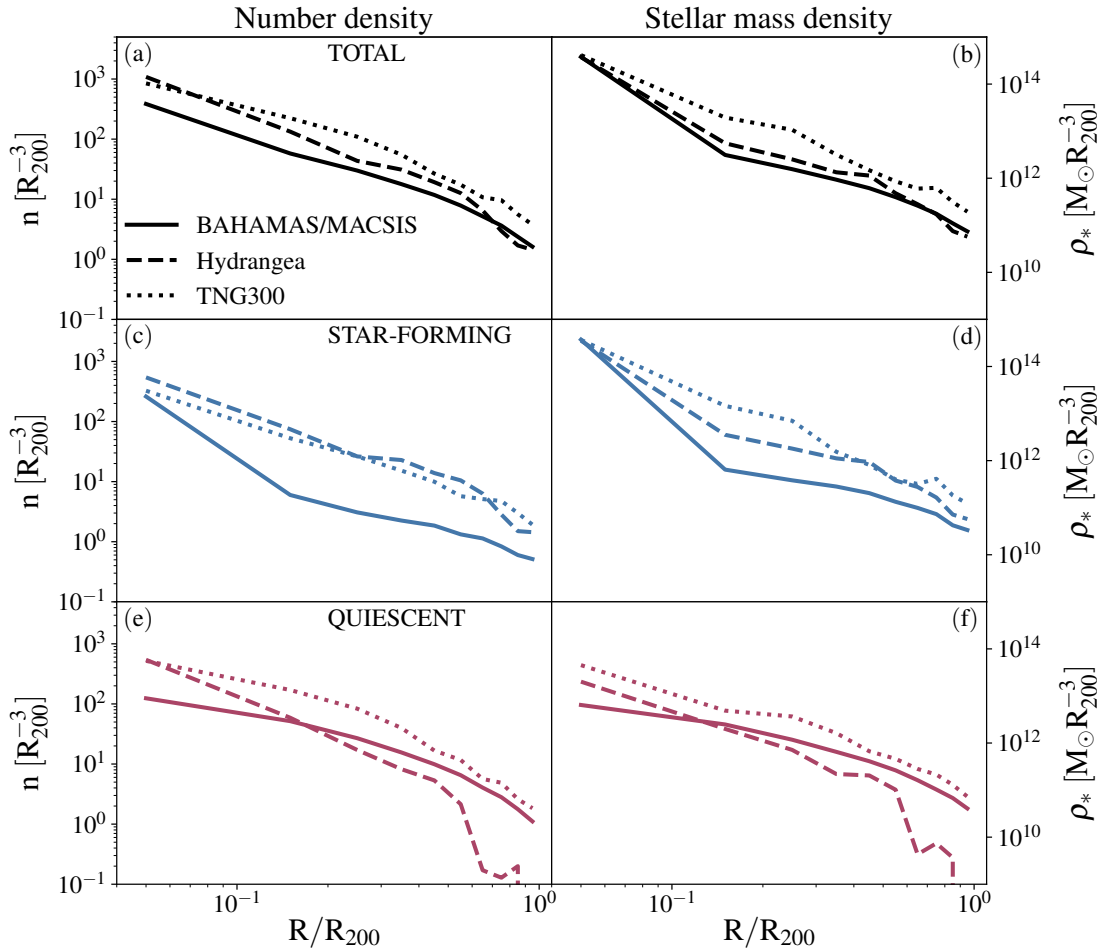


FIGURE 3.4: Galaxy number density (left column) and mass density (right column) as a function of cluster-centric radius, total as well as split up by galaxy type. Solid lines represent BAHAMAS/MACSYS, dashed lines represent Hydrangea, and dotted lines show TNG300 results. All three simulations show similar total radial distributions but differ substantially in star-forming and quiescent profiles. See text for more details.

3.4.3 Galaxy stellar mass function: field vs. clusters

The relative abundance of galaxies as a function of stellar mass (GSMF) is, perhaps, the most fundamental metric in studying galaxy formation and evolution. It is sensitive to both stellar and AGN feedback when studied as a function of redshift or time. Splitting the total GSMF into star-forming and quenched components can also differentiate the dominant processes driving the transition from star-forming to quiescent (Peng et al., 2010; van der Burg et al., 2013, 2018). Indeed, this is what van der Burg et al. (2020) have reported: the shapes of field and cluster star-forming/quiescent galaxy stellar mass functions are the same, indicating that whatever is quenching galaxies in the cluster environment is acting in the same fashion as mass-quenching in the field. This is a

striking difference from what one would naively expect; ram-pressure stripping quenches galaxies independently of their stellar mass and, thus, changes the shape of the GSMF. To perform this measurement, a clean separation of cluster and field galaxies is necessary, which is challenging at $z \geq 1$. This is tested in Chapter 4. In this test, I make a direct comparison to measurements from van der Burg et al. (2020), computing the simulated field and cluster stellar mass functions as well as quenched fractions with the full knowledge of membership.

As cluster richness is strongly correlated with halo mass, it is important to keep the halo mass distribution consistent between simulations and the observations they are compared to. This presents some possibilities and limitations depending on the abundance of GOGREEN-like haloes in the simulation (see Fig. 3.1). For BAHAMAS/MACSYS, there are numerous high-mass haloes that many different samples of 11 GOGREEN-like haloes can be drawn from. This presents an opportunity to explore the possible scatter in the SMF estimate due to shot noise. Therefore, in the case of BAHAMAS/MACSYS, 100 samples of 11 haloes are drawn (matching the GOGREEN distribution) and SMFs estimated. I then plot the median value in bins of stellar mass with 1σ scatter region represented by a shaded region. The same applies for the estimate of quenched fraction. For Hydrangea and TNG300, however, the most massive haloes are still not massive enough to properly match those present in GOGREEN. As the haloes are stacked, the important quantity to match is the mean halo mass. In the case of Hydrangea, there is no other option but to select the 11 most massive haloes shown in Fig. 3.1, giving a mean $\log_{10}M_{200c}/M_{\odot} = 14.34$. TNG300 suffers from the same lack of massive haloes at high-mass end but contains a much higher number of low-mass systems. I, therefore, draw 11 unique haloes matching the GOGREEN halo distribution as closely as possible giving a mean halo mass of $\log_{10}M_{200c}/M_{\odot} = 14.36$. For comparison, the 11 GOGREEN clusters have a mean $\log_{10}M_{200c}/M_{\odot} = 14.35$.

GOGREEN clusters have redshifts spanning the range $1.067 \leq z \leq 1.386$, with a mean of $z = 1.23$. The simulated clusters, however, are output in discrete redshifts in the form of snapshots. The closest common snapshot redshift for all five simulations is $z=1.0$. This is somewhat lower than the observations which will lead to higher galaxy abundances at any given stellar mass bin. This difference in GSMF is small ($< 0.1\text{dex}$), substantially smaller than either the error bars on data points or 1σ scatter due to halo

sampling¹. For all three physics models the field SMFs are computed by considering all galaxies above $\log_{10}(M_*/M_\odot) = 10$ in the simulation volume for BAHAMAS, EAGLE, and TNG300.

Figure 3.5 shows the measured GSMFs and f_q distributions for all three simulations. Field measurements are in the left-hand column and cluster quantities on the right. Each simulation occupies two rows: upper row for the GSMFs and lower for quenched fraction. GOGREEN observations are represented by data points with error bars, whereas simulations are shown by solid lines, switching to dashed lines when there are fewer than ten galaxies in a mass bin. Both observations and simulations show the total GSMFs in black, which are then split up into star-forming (blue) and quiescent (pink).

Examining BAHAMAS field GSMF in panel (a1) of Fig. 3.5 reveals that the observed total GSMF is reproduced rather well, with excess abundance at $\log_{10}(M_*/M_\odot) < 10.5$. The same feature can be seen in Figure 13 of McCarthy et al. (2017) and is due to resolution limitations of the simulation. The shapes of star-forming and quiescent SMFs are both similar to the resulting total but the absolute abundances do not match those observed. More importantly, the observed star-forming and quiescent SMFs cross at $\log_{10}(M_*/M_\odot) \approx 10.75$, with star-forming galaxies dominating the low-mass end and quiescent contributing the most at high masses. This is not seen in BAHAMAS, star-forming galaxies dominate the low-mass end but they also do so at high masses. This is reflected quite concisely in the corresponding quenched fraction in panel (a2): BAHAMAS recovers the correct low-mass f_q but, while the observed fraction of quiescent galaxies begins to match and exceed the star-forming abundance, the simulated abundances continue the trend independently of stellar mass. The field quenched fraction does not increase with stellar mass as it should. This suggests that AGN feedback or some other process is failing to halt star formation in massive galaxies.

Panel (b1) of Fig. 3.5 shows the equivalent cluster GSMF for BAHAMAS/MACSYS. Starting with the total SMF, it is immediately apparent that the shape is not reproduced correctly: there is agreement at the lowest masses but there is a significant deficit in the $10.5 < \log_{10}(M_*/M_\odot) < 11.2$ and an excess at high masses ($11.5 < \log_{10}(M_*/M_\odot)$). This amounts to the simulated SMF missing out the ‘knee’ of the distribution. When split by

¹An approximate magnitude of this difference can be seen in Fig. 13 of McCarthy et al. (2017) where GSMFs for $z=1.0$ and $z=1.5$ are plotted in the top right panel.

galaxy type, the simulated SMF is dominated by quiescent galaxies at low M_* and star-forming at high masses. This is contrary to observations which show equal abundances at low masses and dominance of quiescent galaxies at high M_* . This is reflected in the quenched fraction (panel (b2)) which increases from 0.5 (elevated relative to the field) up to 1.0 with increasing M_* . BAHAMAS shows the exact opposite trend: f_q of almost 1.0 at low masses, decreasing to zero at the highest masses. Low-mass satellite galaxies are being quenched too efficiently upon infall, while high-mass galaxies, including centrals, continue forming stars when they should be quiescent.

Panel (c1) shows the field GSMF estimates using the 50 cMpc EAGLE AGNdT9 simulation. It is evident that the simulation box is too small to contain a representative number of galaxies above $\log_{10}(M_*/M_\odot) = 10.6$. The premature decline of the GSMF is most certainly driven by the box size. The low-mass end (below $\log_{10}(M_*/M_\odot) = 10.6$) contains sufficient galaxies to assess the model by considering just the amplitude alone. The AGNdT9 model appears to perform similarly to BAHAMAS: it over-predicts the abundance of low mass galaxies. Qualitatively, the vast majority of these galaxies are star-forming, which is in line with observations. However, the abundances do not agree quantitatively: there are substantially more low-mass, star-forming galaxies in EAGLE than are observed. The deficit of quiescent galaxies suggests that feedback does not efficiently regulate star formation at $z=1.0$.

High-mass galaxy abundance is less of an issue in Hydrangea in panel (d1), where it spans the full range of GOGREEN observations. The total GSMF function exhibits similar behaviour to what was seen for BAHAMAS: a reasonable match at lowest masses and $\log_{10}(M_*/M_\odot) \approx 10$, but an absence of the ‘knee’, as well as an excess of high mass centrals. When split up by galaxy type, the SMFs are quite different from those in panel (b1): Hydrangea has an abundance of low-mass, star-forming galaxies matching the GOGREEN estimate, however, the quiescent population is significantly less abundant relative to both GOGREEN and, especially so, BAHAMAS. Star-forming population dominates the total at high masses, a trend also seen in BAHAMAS. The resulting quenched fraction is, therefore, too low at all stellar masses. It displays a similarly negative correlation with stellar mass as BAHAMAS. Curiously, a more extreme case of such relationship was reported at $z=0$ by Bahé et al. (2017a) (see their Fig. 6) which suggests that the cluster-specific mechanism does not change between $z=1.0$ and $z=0$, it simply becomes more efficient.

Panel (e1) shows the field GSMF for TNG300. The simulation volume is sufficiently large to contain galaxies over the entire stellar mass range studied here. A reasonably good match is achieved for all but the highest stellar masses in the total SMF, with there being a clear excess of galaxies above $\log_{10}(M_*/M_\odot) \approx 11.5$ when compared to GOGREEN estimates. These very massive galaxies are extremely likely to be BCGs at the centres of clusters which suggests that there may be an over-cooling issue in the model. When split by type, the composition of GSMF is a more extreme version of what is seen in the case of BAHAMAS: the fraction of star-forming galaxies is substantially higher than what is observed, especially at high stellar masses. The quenched GSMF exhibits a sudden decline below $\log_{10}(M_*/M_\odot) \approx 10.6$, which may be indicative of a characteristic stellar mass at which feedback switches modes. This is reflected in quenched fraction being lower than observed at all stellar masses but, especially, at the extremes.

The excess number of massive galaxies carries over to the cluster GSMF in panel (f1). Here, the total SMF reproduces the correct abundances below $\log_{10}(M_*/M_\odot) \approx 11.2$, but shows an even more extreme tail of massive galaxies giving the overall function a more flattened shape. When split by galaxy type, TNG300 shows similar qualitative results to BAHAMAS: quiescent GSMF has an approximately correct shape and amplitude but the star-forming population is under-represented at low stellar masses and far too abundant at $\log_{10}(M_*/M_\odot) > 10.75$. The most massive galaxies are, again, star-forming. The absolute abundance differences are not quite so extreme and, therefore, combine to produce quenched fraction which is more-or-less correct at $\log_{10}(M_*/M_\odot) < 10.75$. At higher stellar masses f_q does not increase and even declines, contrary to observations. The unique feedback implementation in TNG300 appears to control satellite star formation adequately but does not sufficiently quench centrals, a feature of all three simulations.

All three models display some common features: (1.) more-or-less recovering the total field GSMF but failing to capture the details of quenching, with at best matching only a small portion of the $f_q - \log_{10}(M_*/M_\odot)$ relation; (2.) capturing the total cluster GSMF to within 1σ below $\log_{10}(M_*/M_\odot) = 11.25$ but showing a significant excess of high-mass galaxies, many of which are centrals in this case; (3.) high-mass galaxies are actively star-forming in all cases, which suggests that there may be issues with over-cooling or gas inflows. They also show a number of differences: (1.) both Hydrangea and TNG300 under-quench low-mass galaxies, with TNG300 showing a two-phase behaviour

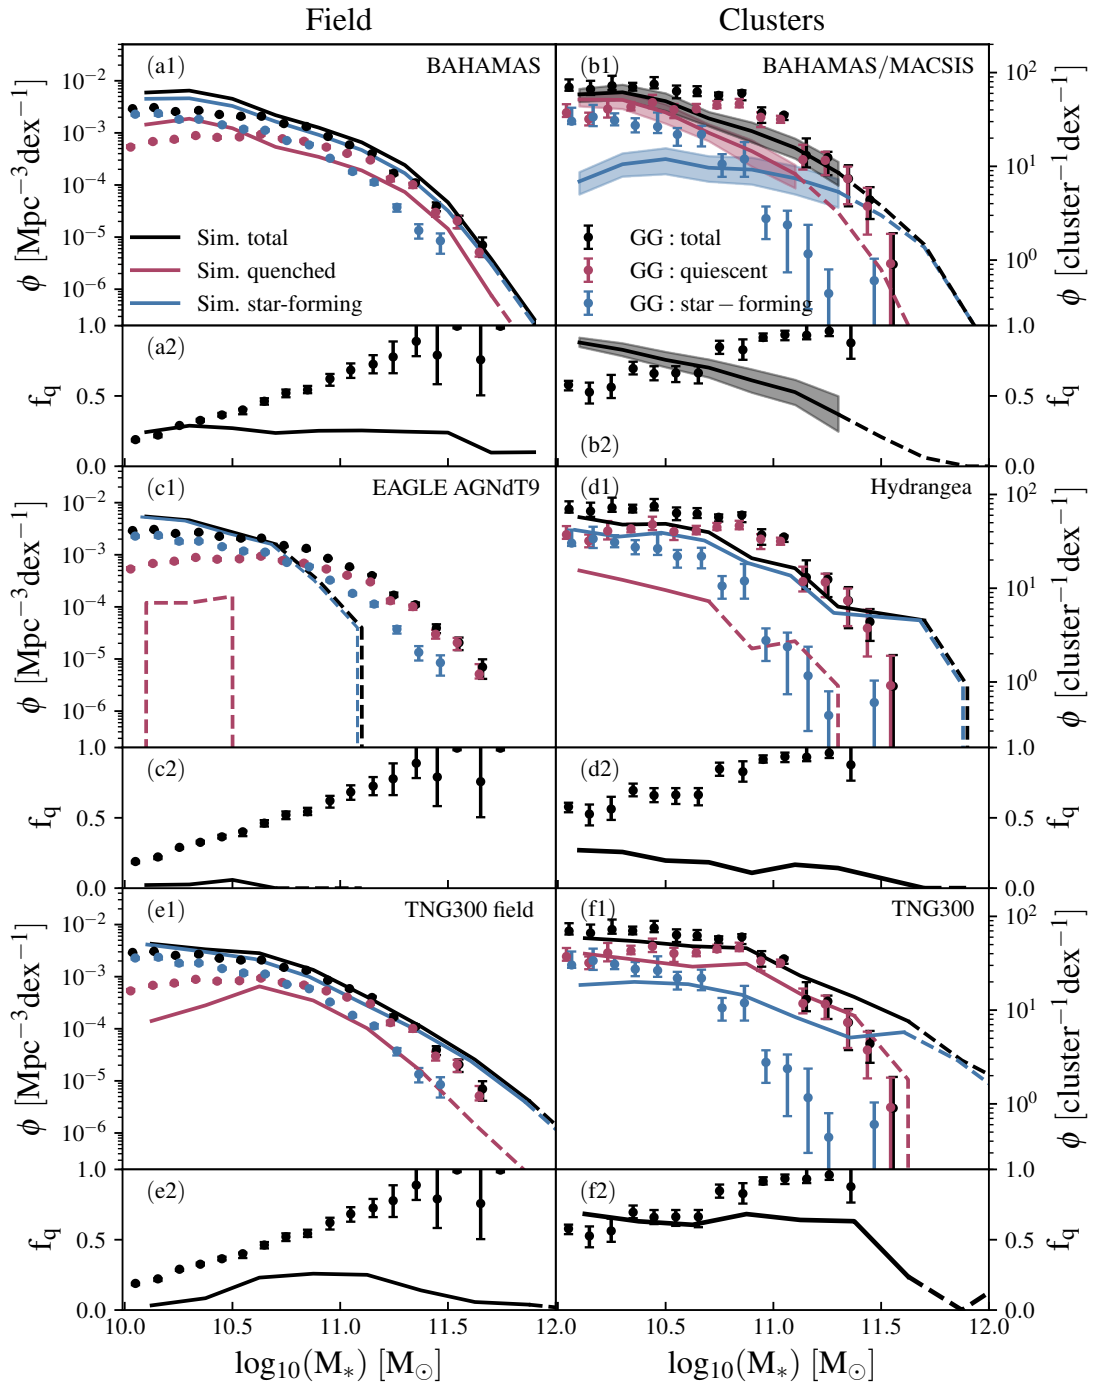


FIGURE 3.5: Cluster (left column) and field (right column) galaxy stellar mass functions and quenched fractions for the three simulation models (solid and dashed lines) while points with error bars show the observed GOGREEN SMFs (repeated for each panel). Black lines/points represent all galaxies, blue – star-forming, and pink–quenched. Dashed lines indicate bins with fewer than 10 galaxies in them. Shaded regions in the BAHAMAS/MACSIS panel represent a 1σ scatter region around the median computed from 100 GOGREEN-like 11 cluster samples. All simulations reproduce the total GSMF reasonably well but deviate from observations when split up by star-formation state. This is reflected in the quenched fraction, see main text for more details.

that switches modes at $\log_{10}(M_*/M_\odot) \approx 11.0$; (2.) Hydrangea is the only one where star-forming galaxies dominate (f_q below 0.5) in clusters regardless of stellar mass, both of which disagree with observations.

3.4.4 Halo quenched fraction

Figure 3.6a shows the quenched fraction of all FoF members within a radius of R_{200c} . The three simulations predict substantially different values: most galaxies in Hydrangea clusters are star-forming which results in mean f_q of ~ 0.2 , BAHAMAS model has the highest mean f_q of ~ 0.7 , while TNG300 contains clusters with quenched fractions that are consistent with either simulation but has a mean of $f_q \sim 0.5$. All three simulations appear to show a slight rise of f_q with halo mass, but this is best seen for BAHAMAS/MACSYS due to the abundance of massive haloes.

BAHAMAS/MACSYS cluster galaxies are more likely to be quenched than those in TNG300 or Hydrangea, while the BCGs are always star-forming. This is likely due to excessive fuelling by gas cooling in the central regions and depositing onto the central SMBH which drives AGN feedback. Since BAHAMAS/MACSYS has the lowest resolution, there are fewer particles available for energy injection. As a result, feedback events happen less frequently but more energy is injected per event in order to match the black hole scaling relations. This ‘bursty’ form of feedback may explain the slightly lower total stellar masses in panel (b) of Fig. 3.3; majority of BAHAMAS/MACSYS galaxies are no longer forming stars and so not increasing in stellar mass, which may have been the case for a significant period of time.

In order to isolate the effects of environment and represent the fraction of satellites which would be star-forming had they not fallen into a cluster I compute the quenched fraction excess (QFE): $(f_q^{\text{cluster}} - f_q^{\text{field}})/(1 - f_q^{\text{field}})$ as proposed by Wetzel et al. (2012). Figure 3.6b shows the QFE computed for GOGREEN and the three simulations of interest. Contrary to Wetzel et al. (2012), van der Burg et al. (2020) report a strongly increasing QFE with satellite mass leading to the conclusion that quenching mechanisms are not the same at $z \sim 1.0$ and $z \sim 0.0$. BAHAMAS returns a QFE that is strongly declining with stellar mass. This is not surprising given that f_q^{field} does not vary much with $\log_{10}(M_*)$ and f_q^{cluster} is anti-correlated with stellar mass. Both of these trends are incorrect, whether they are caused by the same mode of feedback is a question for future

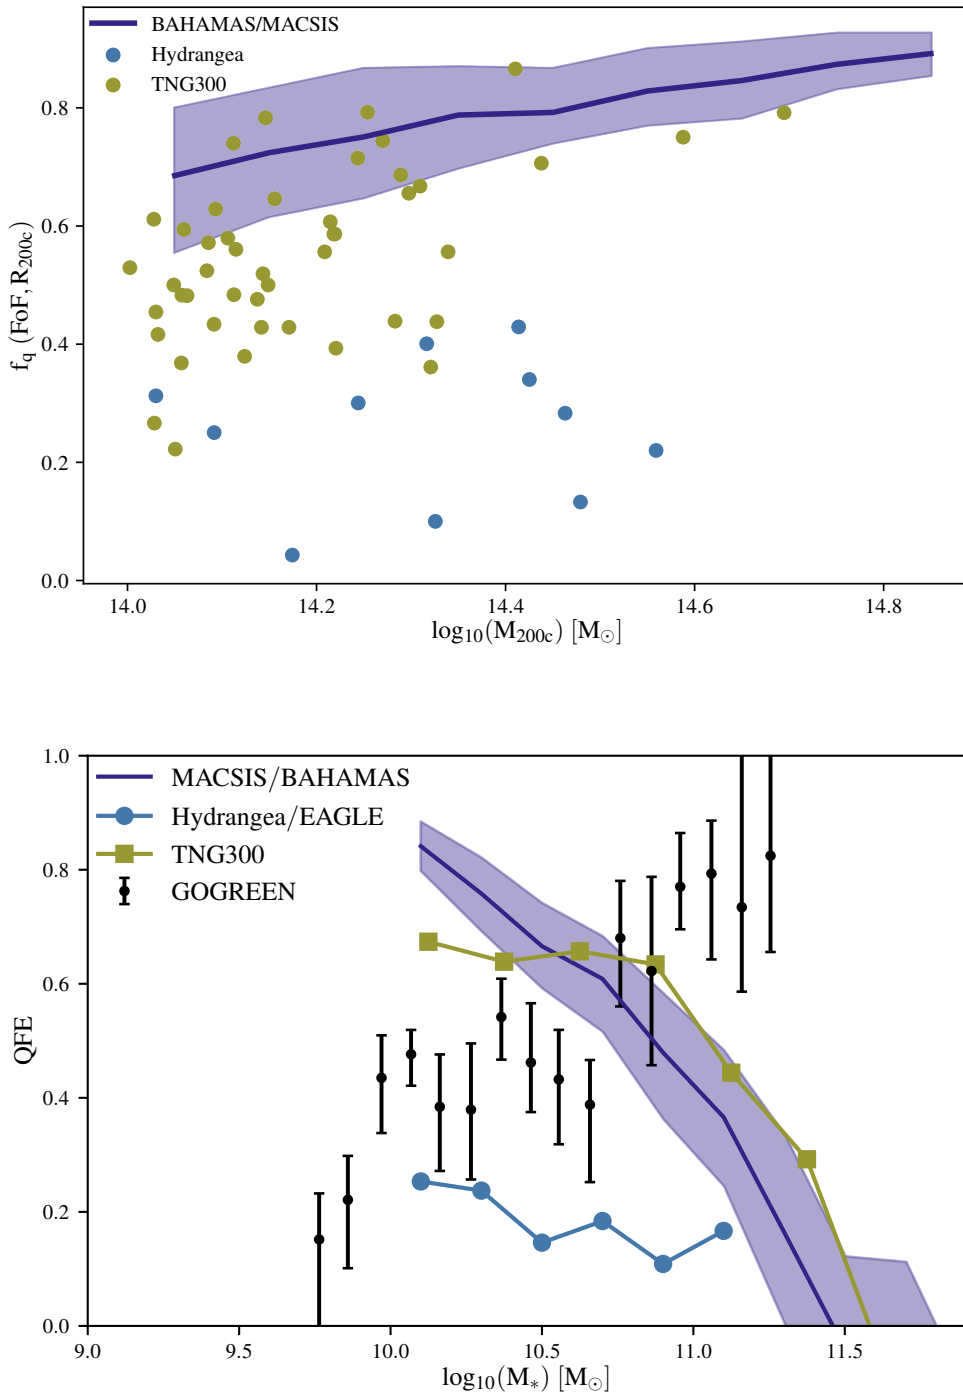


FIGURE 3.6: **Top:** Quenched fraction estimated from all FoF members within R_{200c} as a function of halo mass. A clear ranking of increasing $f_q(\text{FoF}, R_{200c})$ can be seen in the three simulations: hydrangea predicts the lowest f_q of ~ 0.2 , TNG300 predicts $f_q \approx 0.5$, and BAHAMAS/MACSIS predict median $f_q \approx 0.7$, although there is a slight trend with increasing halo mass.

Bottom: Quenched fraction excess for GOGREEN observations and the three simulations. BAHAMAS shows a strong anti-correlation with stellar mass, while Hydrangea predicts a stellar mass-independent relation. TNG300 shows both aspects: a flat curve at intermediate masses and an anti-correlation at high masses.

studies. Hydrangea shows at most a very slight anti-correlation but is consistent with no trend with stellar mass, assuming that scatter between samples would be similar to BAHAMAS. The overall amplitude is significantly lower than GOGREEN, especially at high stellar masses. This indicates that there is insufficient satellite quenching in the AG-NdT9 model. TNG300 shows a combination of the two trends: a mass-independent QFE in the range $10.0 < \log_{10}(M_*/M_\odot) < 11.0$ and a steep decline at $\log_{10}(M_*/M_\odot) > 11.0$.

3.4.5 Quenched fraction profile

The effectiveness of galaxy quenching in clusters for each simulation can be best seen in Figure 3.7 where I plot the quenched fraction as a function of radius. Here one can see the effect of the substantial deficit in star-forming galaxies in BAHAMAS/MACSIS clusters: quenched fraction is above 0.8 at most radii and the decline in the centre is most likely caused by star-forming centrals. BAHAMAS model is very efficient at quenching satellite galaxies even at the edges of clusters. TNG300 has a similar behaviour but is lower in magnitude. Quenched fraction peaks at ~ 0.8 at intermediate radii, but declines towards the cluster centre and outskirts. Again, centrals are more likely to be star-forming, which can explain the former deficit. Hydrangea has the most different f_q profile of all three. Not only does it have quenched fraction below 0.5 at any radius, it reaches $f_q \approx 0$ at cluster outskirts. Cluster quenching mechanisms are relatively very inefficient in the EAGLE model at $z = 1$.

3.5 Summary and Conclusions

There is growing evidence that environmental quenching mechanisms may be different at $z \sim 1$ from those observed in the local Universe. Satellite galaxies appear to be transitioning from star-forming to quiescent in a similar way that isolated centrals in the ‘field’ undergo the same transformation, just faster. Cosmological hydrodynamical simulations are often calibrated on observed field quantities. They have been shown to be lacking when subjected to tests of environmental quenching at $z \sim 0$. With environmental processes potentially being more like the field at higher redshifts, it is possible that simulations are capable of reproducing such observational trends. Understanding why simulations are in agreement with observations (or not) at high redshift may inform

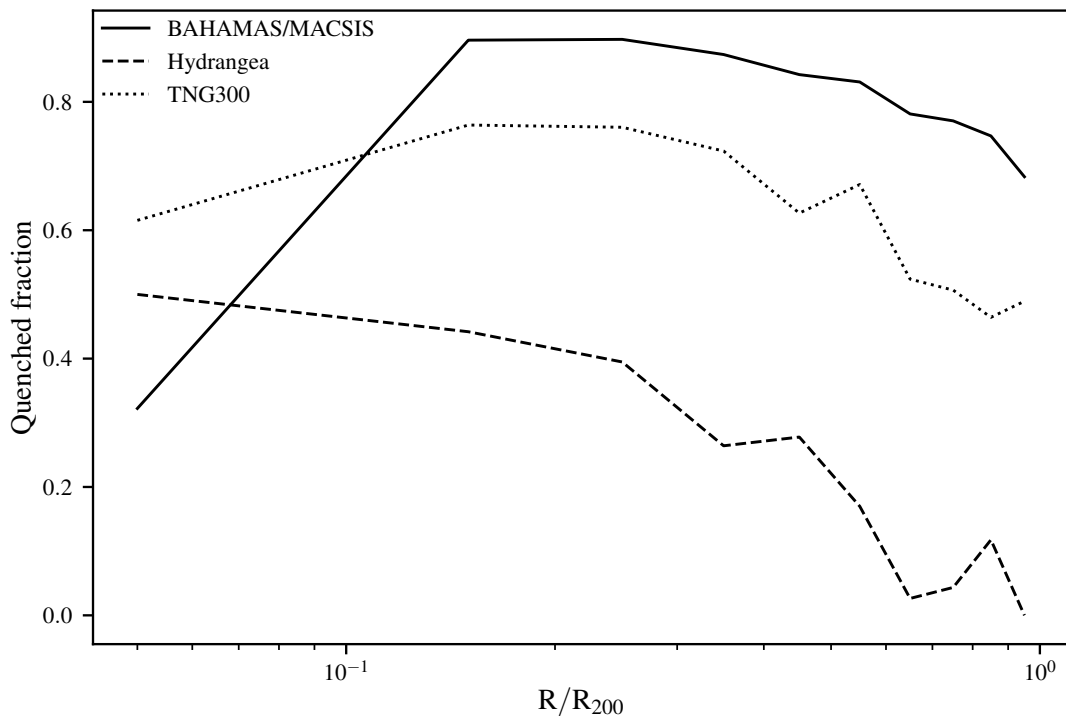


FIGURE 3.7: Quenched fraction for galaxies with $\log_{10}(M_*/M_\odot) > 10.0$ as a function of radius. Solid line represents BAHAMAS/MAC SIS, dashed line - Hydrangea, and dotted line shows the TNG300 prediction. BAHAMAS/MAC SIS and TNG300 predict $f_q \gtrsim 0.7$ until the inner regions of $0.15R_{200c}$ where they both decline. Likely due to star-forming centrals.

and guide the physical modelling towards a better agreement at low redshifts. With large, high quality, spectroscopic observations at redshifts $z > 1$ such as GOGREEN, it is becoming increasingly possible to subject simulations to rigorous tests of environmental processes at earlier stages of evolution of the Universe. In this chapter, I have taken three different implementations of physical models (BAHAMAS, EAGLE, and TNG300 as periodic boxes and MAC SIS, Hydrangea as zoom-in simulations) and produced a number of observationally-motivated relations that will test their validity once data becomes available.

I began by comparing the stellar content of haloes, both in the field and clusters in the same halo mass regime as GOGREEN. This revealed that centrals in BAHAMAS/MAC SIS are significantly less compact when compared to TNG300/Hydrangea. There is a significantly higher Intra-Cluster Light (ICL) component associated with BAHAMAS/MAC SIS centrals. All three simulation codes were found to reproduce the total galaxy

stellar mass function (field and cluster) reasonably well. Their predictions deviate in different ways when galaxies are split up by type into star-forming and quiescent, however. In all cases, centrals were found to be star-forming which is quite contrary to available observations.

Next, I focussed on the quenched fraction measures in clusters and the field. All three models predict elevated quenched fractions in clusters relative to the field but do not reproduce the observed amplitude nor trend with stellar mass of galaxies. The models give three different predictions for the $f_q - \log_{10}(M_{200c}/M_{\odot})$ relation, as well as substantially different distributions of quenched fraction as a function of cluster-centric radius. Once observational data for these quantities becomes available, they will be able to rule out models which do not work and also inform future choices.

All of the above comparisons to GOGREEN have been made under the assumption that the derived halo masses are accurate and galaxies have been assigned reliable ‘cluster’ or ‘field’ status. The same requirements will apply to future surveys. In the next chapter I perform a test on the accuracy of membership assignment in van der Burg et al. (2020) and its effect on the measured star-forming/quiescent galaxy stellar mass functions.

Chapter 4

Testing the accuracy of observational galaxy selection with simulations

4.1 Introduction

Studying dense environments as a function of time (or redshift) can be a very effective way of understanding the underlying mechanisms of galaxy quenching. As gas accretion rates, relative gas masses, and star-formation rates of galaxies are known to peak at $z \sim 2$ (Madau & Dickinson, 2014) and decline to the present day, there is a possibility that the dominant quenching mechanism changes with time. By studying the differences between high and low redshift dense environments one may be able to gain further understanding of processes at each epoch.

Large spectroscopic samples of galaxies in groups and clusters are being compiled at ever increasing redshifts. Some examples of such surveys include: EDisCS (White et al., 2005), MeNEACS (Sand et al., 2012), CCCP (Hoekstra et al., 2012), CNOC (Yee et al., 1996), GEEC (Wilman et al., 2005), CLASH (Postman et al., 2012), GEEC2 (Balogh et al., 2014), GCLASS (Muzzin et al., 2012), and GOGREEN (Balogh et al., 2017). When performing such surveys a trade-off must be made between redshift and stellar mass range. Deep spectroscopic observations are extremely expensive to make,

requiring as much as multiple hours of integration time per cluster. Despite great efforts, spectroscopic completeness for cluster galaxies is never 100%. In the case of GOGREEN ($\bar{z} = 1.23$), only about 1/3 of cluster member galaxies with stellar masses above $\log_{10}(M_*/M_\odot) = 10.0$ have reliable spectra. For this reason, most surveys supplement their samples with multi-band photometric data in order to obtain redshifts (photo-z), stellar masses, and star formation rates from spectral energy distribution (SED) fits to models of stellar populations.

The implications of such incompleteness depend on what the surveys are being used for. If the goal is to study galaxy dynamics for the purposes of measuring cluster centres, masses, radii, etc., incomplete spectroscopic samples can be used quite successfully. Such approaches are mainly used to study the gravitational potential and the nature of dark matter (e.g. Biviano et al. 2016). Algorithms working in the projected phase-space, such as ‘shifting gapper’ (Fadda et al., 1996) and ‘Clean’ (Mamon et al., 2013), are commonly used to assign membership. For studies focussed on galaxy environment and its effect on member galaxies, however, stellar mass-completeness is a strong requirement. Sample purity is often sacrificed in favour of completeness by introducing a line-of-sight (LOS) velocity cut centred on the cluster. Cluster centre and velocity are often taken to be equal to the most massive cluster galaxy (the BCG) or obtained from one of the algorithms mentioned above (Muzzin et al., 2012; van der Burg et al., 2014). Such an approach has recently been used in van der Burg et al. (2020) (henceforth vdB20).

In vdB20, spectroscopically observed galaxies were assigned membership by introducing a circular aperture of $R = 1$ Mpc, centred on the cluster BCG, and a cut in velocity relative to the cluster corresponding to a redshift difference of $|\Delta z| = 0.02$, giving observed velocity of $c \times |\Delta z| = 6000$ km/s and a rest-frame velocity of $c \times |\Delta z| / (1 + z) \simeq 2700$ km/s. This choice corresponds to $2 - 3\sigma_{LOS}$ of the most massive GOGREEN clusters. Galaxies with only photo-z estimates (non-targets, because they were not targeted for spectroscopic observation) were given a more generous $|\Delta z| = 0.08$. By assuming that spectroscopic galaxies are a representative sub-sample of the entire sample, vdB20 were able to introduce correction factors in order to assign non-targets as being members/non-members in a probabilistic fashion (see Sec. 3.5 of vdB20 for a more detailed description). The effect of such an exercise is that effective membership selection applied to *all* galaxies is that of spectroscopically-targeted ones, i.e. $R = 1$ Mpc and $|\Delta z| = 0.02$. Several assumptions are made in this process and vdB20 have performed a number of

tests to assess its accuracy. They have not, however, used a mock catalogue to apply such a selection and compare against ‘true’ membership. Hydrodynamical simulations are an ideal tool for testing such methods as they not only contain realistic clusters with known membership but also reproduce realistic galaxies. These can then be used to assess the effect of impurities on a variety of measurements, such as the galaxy stellar mass function - the key result of vdB20. There are two parts of the method to be tested: the circular aperture combined with a spectroscopic cut of $|\Delta z| = 0.02$ and corrections applied to non-target galaxies to account for photo-z uncertainties. In this chapter, I test the accuracy of the first part of this method, leaving the latter part for later studies.

This chapter is structured in the following way: Section 4.2 describes the hydrodynamical simulation used to perform the tests; Section 4.3.1 describes how observational redshifts were computed; Sec. 4.4 presents the results in the form of phase-space diagrams of different samples, effect on GSMF and quenched fraction estimates, as well as the radial distribution of galaxies. Finally, in Sec. 4.5, I summarise the findings and suggest improvements on the existing methods.

4.2 Simulation description

To perform this test I use the BAHAMAS (BARYONS and HALOES OF MASSIVE SYSTEMS, McCarthy et al. 2017, 2018) set of cosmological, hydrodynamic simulations. With a periodic box size of 400 cMpc/h, it is sufficiently large to enable the introduction of $\Delta z = 0.02$ in the line-of-sight direction. BAHAMAS includes subgrid treatments for components which cannot be resolved using simulation particles. Radiative cooling/heating rates are computed following (Wiersma et al., 2009a); star-formation follows the implementation of (Schaye & Dalla Vecchia, 2008); stellar evolution and chemical synthesis follows the model of Wiersma et al. (2009b); stellar feedback is implemented following the model of Dalla Vecchia & Schaye (2008b); finally, AGN feedback follows the model of Booth & Schaye (2009). Parameters governing the efficiencies of AGN and stellar feedback were adjusted so that the simulation reproduces the observed galaxy stellar mass function for $M_* \geq 10^{10} M_\odot$ and baryon content of groups and clusters, as dictated by the gas mass fraction–halo mass relation from high-resolution X-ray observations. Here I use the run that adopts a WMAP9 (Hinshaw et al., 2009) cosmology with massless neutrinos.

Haloes are identified by using a standard friends-of-friends (FoF) percolation algorithm applied on the dark matter (DM) particles in three dimensions. Each particle is connected to its neighbours with a fixed ‘linking length’, b , which is a fraction of the mean inter-particle separation:

$$s = \frac{L_{\text{box}}}{N^{\frac{1}{3}}}, \quad (4.1)$$

where L_{box} is the size of the simulation box and N is the number of particles in it. For BAHAMAS these values are 400 cMpc/h and 1024, respectively. The linking length used in this study is $b = 0.2s$. Star and gas particles are assigned to haloes by locating the nearest DM particle and associating the baryons to the same halo it belongs to. Galaxies and their properties are then identified by the SUBFIND algorithm (Springel et al., 2001; Dolag et al., 2009). Central galaxies are defined to be baryons belonging to the most massive subhalo in a FoF group. All other galaxies in the same FoF group are then deemed to be satellites. Whenever I refer to a FoF group in the following sections, I mean the FoF members obtained through this method.

4.3 Computing observable quantities

4.3.1 Observed redshift

There are two main components contributing to an observed recession velocity of a distant galaxy: (i) Hubble flow associated with the expansion of the Universe and (ii) galaxy’s peculiar velocity along the line of sight. A table of comoving distances for given redshifts can be computed following Hogg (1999):

$$d_c(z) = d_H \int_0^z \frac{dz'}{E(z')}, \quad (4.2)$$

where $d_C(z)$ is the comoving distance at redshift z , $d_H = c/H_0$ is the Hubble distance (where c is the speed of light and H_0 the Hubble constant at the present time), and

$$E(z') = \sqrt{\Omega_r(1+z')^4 + \Omega_m(1+z')^3 + \Omega_k(1+z')^2 + \Omega_\Lambda}, \quad (4.3)$$

with $\Omega_r, \Omega_m, \Omega_k, \Omega_\Lambda$ representing radiation, matter, curvature, and cosmological constant densities, respectively. For this test I choose a snapshot which is closest to GOGREEN mean redshift ($\bar{z} = 1.23$). The comoving distance to the centre of the simulation box at $z = 1$ is $d_C(z = 1) = 3363.07$ Mpc. The comoving co-ordinates of galaxies in the simulation box are known and can be added/subtracted to/from the ‘snapshot redshift’, z_{hub} , (in the chosen line-of-sight direction) to account for their position relative to the centre. The value of z_{hub} can then be obtained from the previously computed table of $d_C(z)$ and z .

The line-of-sight component of (proper) velocity, v_{pec} , can be easily computed for all galaxies in the simulation box. Since $v_{pec} \ll c$, redshift and velocity are related by: $z_{pec} \equiv v_{pec}/c$. Finally, observed redshift, z_{obs} , can be computed via:

$$(1 + z_{obs}) = (1 + z_{hub})(1 + z_{pec}). \quad (4.4)$$

With observed redshifts computed for every galaxy in the simulation box,

$$|\Delta z| = z_{galaxy} - z_{cluster} \quad (4.5)$$

can be computed for every cluster of interest, and members selected in a way that is consistent with vdB20.

4.3.2 Cluster velocity dispersion

To obtain the velocity dispersion from observed cluster member line-of-sight velocities I use the ‘gapper’ algorithm (Beers et al., 1990), which has been successfully used on observed clusters by Eke et al. (2004) and Robotham et al. (2011). Under this scheme, galaxy velocities are sorted in increasing order, then velocity dispersion is estimated by:

$$\sigma = \frac{\sqrt{\pi}}{N(N-1)} \sum_{i=1}^{N-1} \omega_i g_i, \quad (4.6)$$

where $\omega_i = i(N-i)$ and $g_i = v_{i+1} - v_i$, where N is the number of galaxies in the group or cluster, and v_i is the i^{th} velocity from a list of the galaxy velocities (in one dimension), which has been sorted in ascending order.

4.4 Results

Clusters are not cylindrical in shape and so applying a fixed selection like that chosen by vdB20 risks including field galaxies in the cluster sample while also excluding some cluster members. If contamination is sufficiently large, this risks affecting the measured quenched fractions and changing the shapes of GSMFs. In the following tests, I choose clusters similar to those used by vdB20 but, since no comparison to observations is made, I make no effort to match the selection function or the mean halo mass. Instead, I select all clusters with $\log_{10}(M_{200c}/M_{\odot}) \geq 14.0$ so as to maximise the galaxy number counts. All galaxies with $\log_{10}(M_{*}/M_{\odot}) \geq 10.0$ are considered for analysis. I first assess the distribution of members under each selection in phase-space, showing that peculiar velocities are not the dominant cause of contaminants. I then plot GSMFs for the two main selections and demonstrate that the ‘observational’ selection contains too many star-forming galaxies. Finally, I investigate the cluster-centric distribution of galaxies through radial profiles, showing where the contaminants originate.

4.4.1 Phase-space diagrams

Figure 4.1 shows the distribution of member galaxies under three different selection criteria, divided into three bins of host halo mass (rows), and plotted showing two types of cluster-centric radii (columns). The x-axis is semi-logarithmic to best show the full range in radii, with linear scale in the inner regions of the cluster (within r_{200c}) and logarithmic scale outside. The y-axis shows the spread in cluster-centric velocity, taking the BCG as the centre and normalising by line-of-sight velocity dispersion. $\sigma_{\text{LOS,obs}}$ is computed using the observed velocities of FoF member galaxies using the ‘gapper’ algorithm (Beers et al., 1990) and used in all three cases.

Navy circles represent FoF members (linking length $b = 0.2$) of the main (simulated) halo unconstrained by any additional selection criteria. Examining the left-hand column of Fig. 4.1 reveals that member galaxies extend well beyond r_{200c} under this selection. The most distant galaxies can be found as far as $4r_{200c}$ from the cluster centre in projected space. This does not appear to scale with halo mass. There is relatively little scatter in LOS velocity relative to the BCG, with members confined within $\pm 3\sigma_{\text{LOS,obs}}$. Again, there does not appear to be any changes with halo mass, other than different sampling.

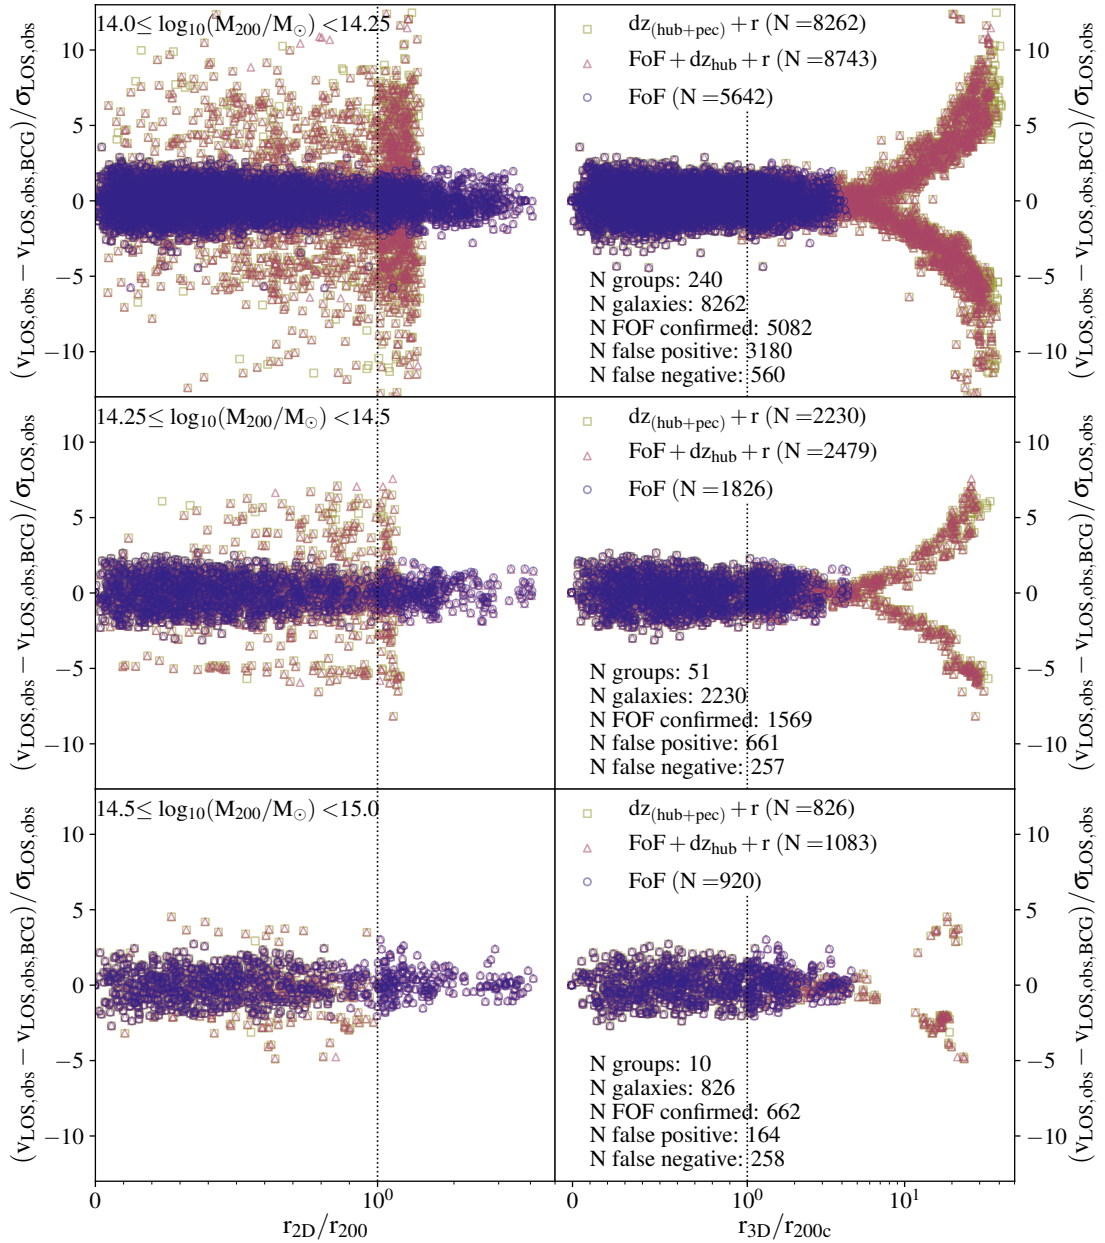


FIGURE 4.1: Phase-space diagrams for three galaxy selections in bins of halo mass. Navy circles show the FoF selection, pink triangles represent FoF-members together with additional galaxies included due to their proximity to the cluster centre, yellow squares show the full ‘observational’ selection of vdB20. Left column shows radial distances in two transverse dimensions; right columns show the three-dimensional radius, normalised by r_{200c} . Rows display samples in three different halo mass bins. Note that x-axis is semi-logarithmic; it is linear in the inner regions of the cluster (within r_{200c}) and logarithmic outside. Numbers displayed in the legends represent the number of galaxies under each selection, whereas numbers near the bottom show diagnostic information between FoF and observational samples.

A similar distribution is seen in the right-hand side column, where the three-dimensional radius is plotted: FoF member galaxies extend out to $\sim 4r_{200c}$ in any direction.

How does imposing the geometry of observational selection affect the sample? In this instance, I exclude peculiar velocities and investigate the contaminants added due to their positions in space (FoF + $dz_{\text{hub}} + r$, i.e. a combination of FoF members, $r_{2D} < 1$ Mpc, and $|\Delta z_{\text{hub}}| = 0.02$ selections). Pink triangles represent a selection of all FoF members *and* galaxies included by the observational selection applied to z_{hub} only. It is immediately clear that the transverse $R = 1$ Mpc cut of the selection is too small to capture all cluster members. Beyond $R_{2D} \sim 1.5r_{200c}$, scatter decreases abruptly to FoF level because only FoF members are represented. This feature can be seen to move closer to (and below) $R_{2D} \sim r_{200c}$ with increasing halo mass, as r_{200c} approaches (and exceeds) 1 Mpc. For comparison, the three halo mass bins have mean $\bar{r}_{200c} = [0.72, 0.86, 1.14]$ Mpc, in order of increasing halo mass. While scatter in the LOS BCG-centric velocity for FoF members is $\sim 6\sigma_{\text{LOS,obs}}$, for the FoF + $dz_{\text{hub}} + r$ sample it increases to $\sim 22\sigma_{\text{LOS,obs}}$. For comparison, mean velocity dispersion for the three halo mass bins, as estimated using FoF members, is $\bar{\sigma}_{\text{LOS,Obs}} = [920, 1020, 1348]$ km/s, whereas the line-of-sight velocity cut is $|\Delta v| = c|\Delta z| \sim 6000$ km/s. This clearly indicates a substantial line-of-sight contamination as galaxies with such velocities cannot possibly be bound to the cluster. By examining the 3D radial distribution it becomes apparent that all of these high-velocity galaxies are at radii far beyond the most distant FoF members. Line-of-sight projection effects are the biggest contributors of non-member galaxies in the sample under this selection.

Peculiar velocities cannot be easily separated from the Hubble flow in real observations. They can both add and subtract to the observed redshift estimate, possibly removing member galaxies from a selection. Such a selection ($dz_{\text{hub+pec}} + r$) is plotted in yellow squares and is the closest representation of the observational selection used in vdB20. Compared to FoF + $dz_{\text{hub}} + r$, this selection shows similar scatter in velocity space, with most galaxies identified as members by both selections. In three dimensions, it becomes apparent that only the most distant member galaxies in this phase-space are affected by peculiar velocity. This suggests that peculiar velocities do not play a major role¹ in membership assignment at $z = 1$.

¹It must be noted that this is likely due to the generous line-of-sight velocity cut. If a smaller value was used, redshift-space distortions would begin to play a more major role.

The important comparison, however, is between this selection and FoF members. In this case, the $d_{\text{hub+pec}} + r$ selection excludes cluster members with $r_{2D} > 1$ Mpc but includes a large number of contaminants in the line-of-sight direction. This can be seen very clearly by comparing the left and right columns in Fig. 4.1. Additionally, although gapper-computed velocity dispersion is quite robust against outliers, mean velocity dispersion computed using the ‘observational’ sample of galaxies are $\bar{\sigma}_{\text{LOS,Obs}} = [1359, 1422, 1558]$ km/s for the three halo mass bins, respectively. These are, on average, 25% higher than those derived using just FoF members. Contamination is sufficiently large that it affects the derived velocity dispersions.

For additional analysis, some diagnostic information is displayed at the bottom of each panel in the right column. From this, one can see that the observational selection achieves sample purity of $\sim 60 - 80\%$, a false positive fraction of $\sim 25 - 60\%$, and a false negative fraction of $\sim 11 - 39\%$ relative to FoF selection. However, it should be noted that FoF selection is not necessarily the ‘true’ membership, either. It varies depending on the choice of linking length. Sample purity and the fraction of false negative members increase with increasing halo mass, while the fraction of false positives decreases quite rapidly. This is consistent with observational selection being too conservative and introducing a large number of false positive members for low-mass haloes, while for the most massive haloes it is more likely to exclude member galaxies rather than add contaminants (although the numbers are quite close and sampling relatively poor). Field galaxies are predominantly star-forming, while cluster galaxies are more likely to be quenched. Including a large number of field galaxies in the cluster sample, while excluding some of the cluster galaxies may lead to a change in the measured GSMF and, subsequently, quenched fraction. However, many of the nearby field galaxies may actually belong to neighbouring groups and be undergoing pre-processing, making them quenched. This would undo some of the effects of field contaminants, negating the bias. This is investigated in the next section.

4.4.2 Effects on GSMF and quenched fraction

Figure 4.2 shows the measured GSMFs (left column) and quenched fractions (right column) under FoF and observational selections. Rows of main panels show estimates in bins of halo mass. Each panel is accompanied by a ratio of observational to FoF

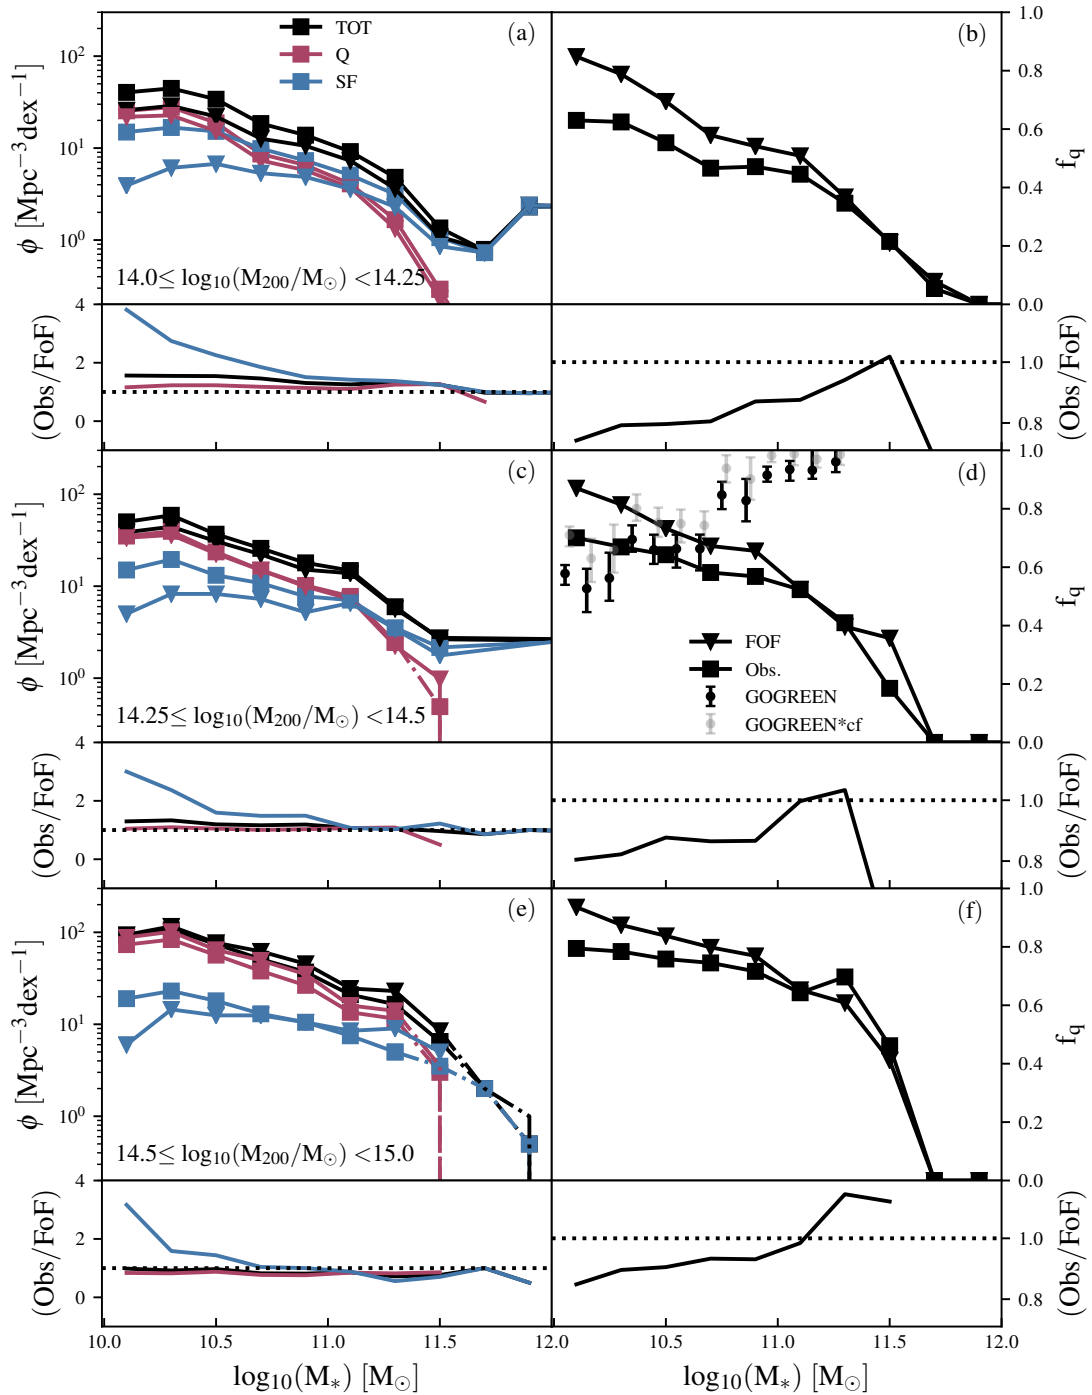


FIGURE 4.2: Cluster GSMF measurements for observational and FoF selections, split by galaxy type (left column) and quenched fraction as a function of stellar mass estimates for both selections (right column). Lines with square markers represent the ‘observational’ selection ($dz_{\text{hub}+\text{pec}} + r$) and lines with triangles show FoF selection measurements. For GSMFs, blue lines represent star-forming galaxies, pink lines show quiescent, and black lines represent all types of galaxies. Dash-dotted lines represent bins with fewer than 10 galaxies in them. Each row represents a different halo mass bin as indicated. Each panel is accompanied by a ratio of observational to FoF selection for each quantity. Observational selection over-estimates the number of low-mass, star-forming galaxies in clusters, leading to suppressed quenched fractions by up to 25%. Data points in panel (d) show GOGREEN quenched fractions (black) from Ch. 3 and those same points multiplied by a correction factor to account for selection bias (grey). Grey points have also been shifted by 0.02 dex on the x-axis for visibility.

selection for each quantity. By examining GSMFs in the left column, it becomes apparent that quiescent population is captured similarly well by both selections. Comparing observational to FoF selections, it is over-estimated by $\sim 10\%$ at all stellar masses in the lowest halo mass bin, likely due to the inclusion of galaxies in neighbouring groups. Meanwhile, in the highest halo mass bin, it is under-estimated by $\sim 10\%$, likely as a consequence of high false negative fraction. However, these differences are very small compared to the star-forming population. In all halo mass bins, the number of low-mass, star-forming galaxies is over-estimated by up to a factor of 4. The lowest halo mass bin is most affected by this. Here, the excess extends from 4 at the lowest masses to eventually reach unity at $\log_{10}(M_*/M_\odot) = 11.5$. The excess is less severe in higher halo mass bins: at intermediate halo masses the excess peaks at 3 and declines more rapidly to zero at $\log_{10}(M_*/M_\odot) = 11.2$; at the highest halo masses, the excess declines more rapidly, from 3 to zero by $\log_{10}(M_*/M_\odot) = 11.2$. The total GSMF is biased accordingly, as it is simply the sum of quiescent and star-forming populations. Both selections agree on the number of overly-massive, star-forming central galaxies - a feature of BAHAMAS and other simulations as shown in the previous chapter.

The excess of low-mass, star-forming galaxies has a significant effect on the measured quenched fractions (see Fig. 4.2(b)). In the lowest halo mass bin, the observational selection results in a quenched fraction that is $\sim 25\%$ lower than FoF sample. This difference decreases with increasing stellar mass, reaching zero at $\log_{10}(M_*/M_\odot) = 11.5$. A similar pattern is seen in higher halo mass bins, with quenched fraction deficit decreasing only slightly and agreement reached at lower stellar masses, in line with what was seen for GSMFs. A disagreement between the two selections is still very much present. The effect of membership bias on GOGREEN quenched fractions can be seen in panel (d). Here, GOGREEN measurements (black points) have been multiplied by a correction factor (cf) to demonstrate the magnitude that such a correction would have. The correction factor was estimated from haloes in all three halo mass bins.

Low-mass, star-forming galaxies from the nearby field are being included in the observational selection. The number of such galaxies is high enough to change the measured quenched fraction in a significant way. Lowest halo mass clusters are affected by this the most, which is mainly due to their velocity dispersion being significantly lower than the line-of-sight cut of $|\Delta z| = 0.02$ being applied. There is some indication that nearby group galaxies are also being included but the number of these galaxies is not nearly

enough to offset the contamination by field galaxies. A slight deficit in quiescent galaxies can be seen for the highest halo mass bin. This is indicative of the observational selection being too restrictive and excluding genuine cluster galaxies, mostly by imposing a radial cut.

4.4.3 Radial profiles

Figure 4.3 shows the 3D radial distributions of galaxy number density (left column) and quenched fraction (right column). By examining the number density profiles (regardless of halo mass bin), one can see that both member selections agree in the inner regions of clusters (within r_{200c}), however they diverge quite dramatically outside. Whereas FoF selection declines quite steeply, terminating by $R_{3D} = 6r_{200c}$, observational selection extends out as far as $R_{3D} = 30r_{200c}$. More importantly, the dominant type of galaxy changes around r_{200c} ; quiescent galaxies dominate in the inner regions but most galaxies in the outer regions are star-forming. The same pattern is present in all halo mass bins.

Examining the quenched fraction distribution shows this very clearly. Both selections yield $f_q \sim 0.8$ inside clusters at all halo masses (it scales with halo mass very slightly, see Fig. 3.6a). Under FoF selection, f_q reaches a plateau at $f_q \sim 0.8$ before terminating abruptly. This is likely the quenched fraction of neighbouring groups. Observational selection, on the other hand, decreases to a lower level and continues to drop. It eventually approaches the global average quenched fraction for the simulation box, further showing that large numbers of field galaxies are being introduced into the sample.

4.5 Summary and conclusions

When comparing properties of cluster galaxies to the field, establishing an unbiased membership sample is essential. In this chapter, I have replicated a spectroscopic member selection in a hydrodynamical simulation by computing observed redshift measurements for every galaxy, then introducing cuts of a circular aperture ($R_{2D} = 1$ Mpc) combined with cluster-centric redshift limit ($|\Delta z| = 0.08$), as used by van der Burg et al. (2020). This ‘observational’ galaxy selection was then compared to members identified by the Friends-of-Friends algorithm. I first examined the distribution of member galaxies in phase-space, using two- and three-dimensional BCG-centric radius measures. For

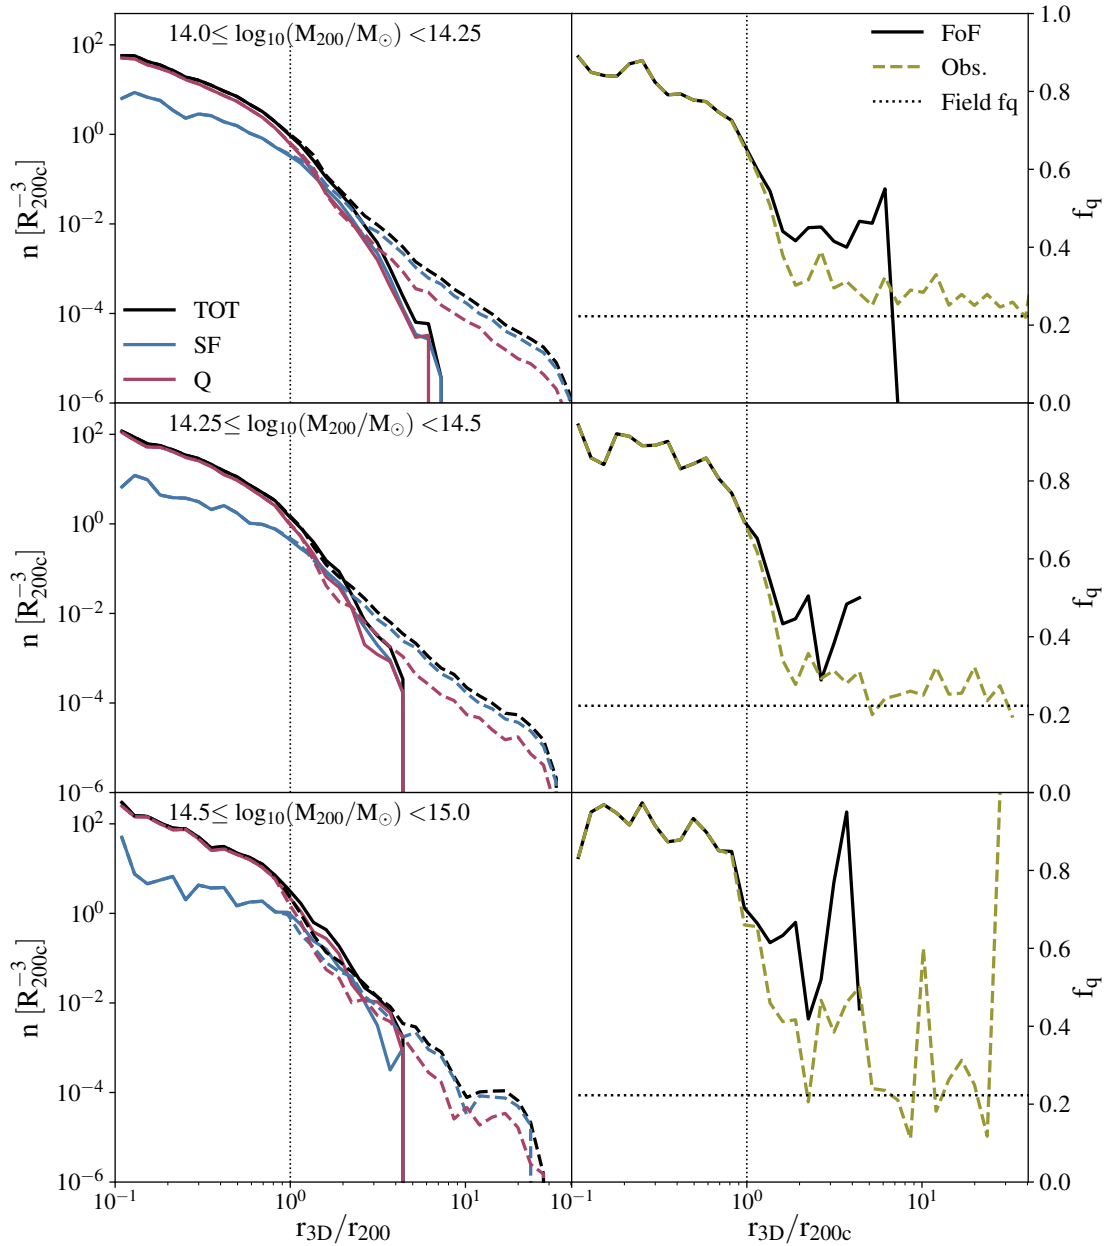


FIGURE 4.3: Galaxy number density (left column) and quenched fraction (right column) radial distribution in clusters. Solid lines represent FoF members, whereas dashed lines show the observational selection of galaxies. Dotted vertical lines indicate the position of r_{200c} , and dotted horizontal lines show the mean quenched fraction of all galaxies in the simulation box. Number density profiles are split up by galaxy type: blue lines for star-forming and pink lines for quiescent galaxies. Both selections agree within $R_{3D} \leq r_{200c}$ of the cluster but deviate significantly outside, with observational selection introducing mostly star-forming galaxies from up to $R_{3D} = 30r_{200c}$. These galaxies also have quenched fraction equivalent to that of the field.

this instance only, I also made a selection excluding peculiar velocities in order to gauge their significance. I then computed cluster galaxy stellar mass functions as done by vdB20 and in Chapter 3, showing differences in the types of galaxies selected by each sample. Finally, I looked at detailed radial distribution of galaxies in each sample, split by type.

The following findings have come about from this exercise:

- FoF-selected galaxies extend out to $R_{3D} \sim 6r_{200c}$. It likely includes galaxies from neighbouring groups.
- The circular, transverse aperture of $R_{2D} = 1$ Mpc aims at matching the r_{200c} of clusters at $z \sim 1$. It is too large for the lower half of halo masses considered here (and in vdB20). It is a good compromise between maximising members and minimising contaminants but, since r_{200c} values are known for observed clusters in vdB20, using the actual values would be more appropriate.
- Line-of-sight interlopers are the main source of contaminants. This is due to a fixed line-of-sight velocity cut of $|\Delta v| = c|\Delta z| \sim 6000$ km/s. Since estimates of σ_{LOS} exist for GOGREEN, a selection which scales linearly with σ_{LOS} of each cluster would yield better results.
- Peculiar velocities are not the dominant driver of contaminants or membership errors. Most contaminants are introduced by the geometry of selection volume, mainly in the line-of-sight direction.
- Large LOS contamination numbers lead to the inclusion of low-mass, star-forming galaxies. This is most severe for low-mass haloes, with up to a factor of 4 more star-forming galaxies when compared to FoF selection. Number of quiescent galaxies is more-or-less equal between selections.
- All halo masses are severely affected by interlopers, however, contamination decreases somewhat with halo mass.
- As a result of increased star-forming galaxy number, measured quenched fraction is lower by up to 25% in the observational sample relative to FoF selection.

- Observationally-selected galaxies can be found up to $30r_{200c}$ away from the cluster centre. Most galaxies at these radii are star-forming, with quenched fractions consistent with those found in the wider field.
- Both selections agree on galaxy distribution within the central r_{200c} of clusters.

Assigning membership is a difficult task at redshifts as high as $z \sim 1$. The pronounced overdensity of clusters considered in vdB20 means that, even with a liberal selection, measurements within the inner r_{200c} are well reproduced. However, the measurements of GSMF and f_q can be severely biased. Such biases need to be taken into consideration, especially when measuring such sensitive features as the shapes of GSMFs. Fine-tuning the selection on mock catalogues (rather than using a one-size-fits-all approach) such as used in this study would lead to a more accurate membership selection and, as a result, better measurements of environmental quenching and a better comparison to simulations.

Chapter 5

Summary, conclusions, and future work

The aim of this thesis has been to make contributions towards understanding the role that environment plays in galaxy evolution. This involved both improving on established methods of member identification and also introducing a novel map-based approach, which allows for connections between galaxies and the underlying environmental components responsible for quenching.

I began this thesis with Chapter 1, introducing the concept of a cosmic web, the way it is characterised, and effects it has on galaxies. I described all three constituent components: dark matter, gas, and stars, and emphasised the abundance of each; it has not necessarily correlated well with the amount of observational focus through history. I then gave a brief description of galaxy transformation and observational evidence in support of the current framework. Part of this framework is the environment a galaxy occupies; there are several ways of studying this and I gave an overview of each.

There are various limitations to how environmental influence on galaxy evolution has been studied in the past. In order to make further progress in understanding what the underlying cause of environmental transformation is, I proposed a different method of studying it in Chapter 2. This map-based approach involves spatially cross-correlating a measure of cluster components thought to be responsible for quenching (such as hot gas) and an ensemble property of galaxies (e.g. quenched fraction). The resulting cross-power spectra quantify the direct link between the quantity causing quenching and its effect on

galaxies - something which has not been possible with conventional methods to date. As a demonstration, I applied this method to currently available data from SDSS, Planck, and ROSAT. Even with relatively low-quality data, strong cross-correlation signals can be measured at low redshifts. I also demonstrated that the same method can be applied to data from cosmological, hydrodynamical simulations reproducing the measured signal. Simulations are instrumental in understanding the signals and the underlying physical processes behind them.

Hydrodynamical simulations fall somewhat short of fully capturing the correct physics necessary for environmental quenching. At low redshift, satellites appear to be quenched too efficiently. There is some evidence suggesting that quenching may be different at higher redshifts of $z \geq 1$. In Chapter 3, I performed a comparison of three different simulation codes at $z = 1$ with a focus on quenching in massive clusters. Comparing to available data for some measurements, none of the simulations fully agree. Some large differences were seen in other measurements for which data is not available yet, suggesting that the implementation of physics varies dramatically at this stage of structure formation. Comparison to observations, once they become available, will be able to inform future simulations about stellar and AGN feedback at intermediate and low redshifts.

One of the benefits of map-based analysis presented in Chapter 2 is that no membership identification is necessary. This task is challenging at low redshifts where data is relatively good, but it is even more difficult at $z \geq 1$, where deep spectroscopy is prohibitive. In this regime, studies often resort to simple volume selections. In Chapter 4, I explored the effects of an imperfect membership selection using a hydrodynamical simulation. This analysis revealed that crude aperture and line-of-sight velocity cuts can introduce large numbers (up to 20% more) of low-mass, star-forming, field galaxies into the cluster sample. Such contamination can change the shapes of measured galaxy stellar mass functions and quenched fractions - measurements used to infer the effects of environment.

The question of how galaxies are quenched in dense parts of the Universe still awaits a conclusive answer. This is the case not due to a lack of effort in trying to understand the problem but because it is just so complex. Environmental quenching is sensitive to all aspects of structure formation: gravitational collapse, gas flows, and astrophysical

processes such as star-formation and feedback. The latter processes are particularly poorly understood, but appear to play a major role. It is further complicated by the dynamic range over which the processes can take place: a galaxy can begin to quench in a relatively low-mass group which falls into a supermassive cluster, completing the quenching process. All structures between the two extremes need to be observed and simulated in good numbers, which is far from an easy task. Therefore, the situation is such: any one (or several) of ten or so processes, driven by three mechanisms can be the reason for quenching at any given time during structure formation. This can change over time and location in the Universe. Deeper, more complete galaxy surveys, combined with observations of dark matter and hot gas will lead to a bigger picture of the processes. Hydrodynamical simulations will help to understand the observables and connect them to the mechanisms at play.

The work presented here has multiple possible directions for further advancement. Chapter 2 introduced the method and demonstrated that it can be applied to existing data. Future high sensitivity X-ray (e.g., with eROSITA) and tSZ effect data (e.g., Simons Observatory, CMB-S4) will allow for measurements to be made at significantly smaller scales and higher depths. Larger, deeper spectroscopic surveys (e.g., WAVES, DESI) are also expected to improve these detections dramatically. More work can be done on the theoretical side, too. Changes to simulation parameters (such as AGN and stellar feedback schemes) could lead to better understanding of the underlying mechanisms and developments of improved feedback schemes. The obvious extension to Chapter 3 is addition of observational data in order to serve as a differentiator between simulations which reproduce realistic observables and those which fall short. It is very possible that none of the three simulation codes used fully match all observations. In this case further adjustments of parameters and models would need to be made in order to achieve a better match at $z \sim 1$. These could then be tested on low redshift data. Finally, Chapter 4 needs to be extended by implementing photometric redshift errors to allow for the full observational selection as done in van der Burg et al. (2020). Other methods of member selection could then be also implemented, leading to a project assessing various different methods. Such a comparison would inform observational studies and allow for better samples to be chosen.

Appendix A

A.1 SPH smoothing and shot noise

In order to map a point object, such as a galaxy from a catalogue or a gas particle from a simulation, onto an extended grid, I make use of a kernel interpolation technique commonly used in smoothed particle hydrodynamic simulations. The kernel, $W(r, h)$, is defined as a spline with a smoothing scale h :

$$W(r, h) = \frac{8}{\pi h^3} \begin{cases} 1 - 6 \left(\frac{r}{h}\right)^2 + 6 \left(\frac{r}{h}\right)^3, & 0 \leq \frac{r}{h} \leq \frac{1}{2}, \\ 2 \left(1 - \frac{r}{h}\right)^3, & \frac{1}{2} < \frac{r}{h} \leq 1, \\ 0, & \frac{r}{h} > 1. \end{cases} \quad (\text{A.1})$$

In the case of simulations, h is determined by the size of a gas particle's 3D smoothing length, whereas when making galaxy maps h is determined by:

$$h = \sqrt{\frac{N_{\text{sph}}}{\pi n}}, \quad (\text{A.2})$$

where N_{sph} is the number of nearest neighbours to be smoothed over and n is the mean surface density of galaxies for the given map.

This spline can be approximated with a Gaussian function, with one notable exception that the spline has a well-defined extent. The distributions of h for each simulation/map can be seen in Figures A.1a and A.1b. All samples show similar distributions in h , with 'characteristic scales' similar between corresponding samples in simulations and observations. The difference in number of galaxies is reflected in the size of each kernel,


with spectroscopic sample demanding a kernel as large as 2 degrees and photometric sample requiring one no larger than 0.4 deg.

The overall effect of smoothing on a power spectrum can be seen in Figure A.2 as solid lines, where two versions of $\check{N}_{\text{tot}} - \check{N}_{\text{tot}}$ (total galaxy overdensity) power spectrum are presented; before smoothing (cyan) and after (navy). The discrete (i.e., unsmoothed) map used to compute this power spectrum was constructed by assigning a value of zero to all pixels with no galaxies within the main SDSS footprint, rather than using SPH smoothing to interpolate. (Note that area outside the footprint is masked as shown in Fig. 2.6 [panel *d*].) The unsmoothed power spectrum rises until it reaches the pixel scale and drops abruptly to zero, which occurs at angular scales significantly smaller than smoothing scales discussed in this paper.


In order to indicate a characteristic scale of SPH smoothing effects, I plot a vertical, black solid line at the scale which marks a 50% deviation of SPH-smoothed power spectrum from the discrete case. For SPH smoothing, these scales are: $\ell = 670$ and $\ell = 1480$ for the ‘spectroscopic’ and ‘photometric’ samples, respectively. This is relative to $\ell = 440$ and $\ell = 3700$ for *Planck* and ROSAT beam effects, respectively. Therefore, the power spectra are limited by SPH smoothing scale in all cases except tSZ cross-correlations with the ‘photometric’ sample, where *Planck* beam effects dominate. The vertical line indicates this accordingly. The effects of SPH smoothing are the same for both simulations, within cosmic variance uncertainty.

Figure A.2 also shows the effects of shot noise on the galaxy overdensity power spectrum. The dash-dotted lines show the level of shot-noise present in the maps. These were computed following Feldman et al. (1994), by randomising galaxy position coordinates in order to remove any structure present in the maps. This yields a power spectrum of constant C_ℓ over the scales presented here. The shot noise power spectrum responds to SPH smoothing in the same way as the signal (cyan and navy dash-dotted lines). I also checked that the effect is the same in the simulations but do not show this explicitly. All \check{N}_{tot} power spectra presented in the main text are shot noise subtracted.

How to rigorously estimate the shot noise contribution to the $\check{f}_q - \check{f}_q$ correlation (quenched fraction power spectrum) is less obvious, however. Since the mean quenched fraction is not conserved when making randomised maps (there is no information about the *number*



(A) Distribution of smoothing parameter (h) for the EAGLE, BAHAMAS, and SDSS photometric low-redshift samples. Vertical, dash-dotted lines indicate the mean smoothing scale for each case. While the distributions are similar in their shapes and mean smoothing scale, differences will translate into small-scale (high ℓ) power-spectrum deviations.



(B) The same as Fig. A.1a but for the high-redshift sample. Note the shift to smaller smoothing scales.

FIGURE A.1

of galaxies used to compute quenched fraction), it is unclear how to calculate a normalisation for the shot noise contribution. I therefore leave the quenched fraction power spectrum affected by shot noise. However, I do not expect shot noise to be a dominant component for this power spectrum for two reasons. First, shot noise is subdominant in both the total galaxy density and quenched galaxy density power spectra and (ignoring the differences of slightly different smoothing kernels) if \check{f}_q is just a division of the two, then its shot noise properties would also be sub-dominant. Secondly, as demonstrated by computing the null-tests, the relative orientation of the two maps is very important. The measured signal vanishes if the maps are mis-aligned, this would not be the case if they were shot noise-dominated.

Lastly, I highlight that shot noise does not affect any cross-spectra, as noise properties are different and certainly not spatially correlated between two different maps.

A.2 Contamination from AGN

Here I explore the possible bias introduced into the X-ray-based cross-spectra by AGN. I test this by introducing an additional mask component which covers the regions with confirmed AGN sources.

The second ROSAT all-sky survey point source catalogue¹ (2RXS, Boller et al. 2016) contains the most complete list of point-like sources observed with the PSPC in the full energy range of RASS. This catalogue also cross-matches the observed sources against a catalogue of known AGNs by Véron-Cetty & Véron (2010). Some ~ 8000 sources are confirmed as AGN to within 1 arcmin of the original source. I use these sources to construct two new masks covering different number of pixels around the AGN.

Note that the 0.1-2.4 keV flux limit of the 2RXS RASS point source catalog is $\approx 10^{-13}$ ergs/s/cm² (Boller et al., 2016). Taking the maximum redshift of our deeper SDSS photometric sample ($z = 0.15$), this corresponds to a conservative soft X-ray luminosity limit of $\approx 6.1 \times 10^{42}$ ergs/s. For the redshift range 0.015–0.2, Miyaji et al. (2001) find that $L_{X,*}$ (the characteristic AGN soft X-ray luminosity) in the same band is $\approx 3.6_{-2.0}^{+4.4} \times 10^{43}$ ergs/s (assuming $H_0 = 70$ km/s/Mpc). Hasinger et al. (2005) combined various ROSAT surveys with deeper Chandra and XMM-Newton observations to derive a more precise

¹<https://heasarc.gsfc.nasa.gov/W3Browse/rosat/rass2rxs.html>




FIGURE A.2: A demonstration of SPH smoothing and shot noise effects on the power spectra for the photometric galaxy sample. Solid lines represent \check{N}_{tot} -auto power spectra for discrete (cyan) and SPH-smoothed (navy) cases. Solid, black, vertical line indicates the ℓ -scale where they begin to differ by more than 50%. Dash-dotted lines show the level of shot noise present in our galaxy maps in discrete (cyan) and SPH-smoothed (navy) forms. Dashed lines show the shot noise-subtracted power spectra in both cases.

(but consistent) constraint of $L_{X,*} = 2.82_{-1.30}^{+3.07} \times 10^{43}$ ergs/s over the same redshift range. Thus, the 2XRS catalog typically probes about a factor of 5 below the characteristic AGN soft X-ray luminosity and should therefore capture most of the X-ray AGN signal in this low redshift regime. Using the luminosity function data in Table 3 of Hasinger et al. (2005), I estimate that $\approx 90\%$ of the X-ray AGN signal lies above the 2XRS point source limit of 6×10^{42} ergs/s when integrating from 10^{42} ergs/s up (the lower limit of the Hasinger luminosity functions). However, I cannot exclude a possible non-negligible contribution from X-ray AGN with luminosities fainter than 10^{42} ergs/s.

In terms of our masking approach, I note that the combined resolution of ROSAT with the PSPC camera is ≈ 1.8 arcmin. This is approximately equal to the resolution of the HEALPix maps. It would be reasonable to assume that a point source can be approximated as 1.8 arcmin in this case and, thus, assigned one pixel - this is our first mask. A more aggressive masking technique is to include the 8 neighbouring pixels as well, masking each source with an area of $\approx 5 \times 5$ arcmin². These masks are then combined with the total mask composed of SDSS footprint, the *Planck* tSZ Milky Way 40% and point-source masks, and the RASS X-ray zero exposure mask.

Figure A.3 shows the resulting $X - \check{N}_{\text{tot}}$ cross-spectra for the AGN-unmasked case shown in Fig. 2.8 in navy, AGN-1px mask in pink, and AGN-9px mask in green. I observe a small difference in the measured cross-spectrum for both of the AGN masks relative to the unmasked case, implying that AGN do slightly contaminate our measured cross-spectrum. However, the magnitude of the effect is not large enough to question the overall nature of the detection (i.e., it is dominated by hot diffuse gas) or to alter the main conclusions of our study. Note that all three cross-spectra are still consistent over most angular scales with BAHAMAS, which has been calibrated to contain observed gas fractions inside groups and clusters and reproduces their X-ray scaling relations (see McCarthy et al. 2017). These two pieces of evidence lead us to conclude that AGN point sources are sub-dominant in the cross-correlations measured in this study, though with the caveat that the faint end of the AGN soft X-ray luminosity function is not well constrained below 10^{42} ergs/s.

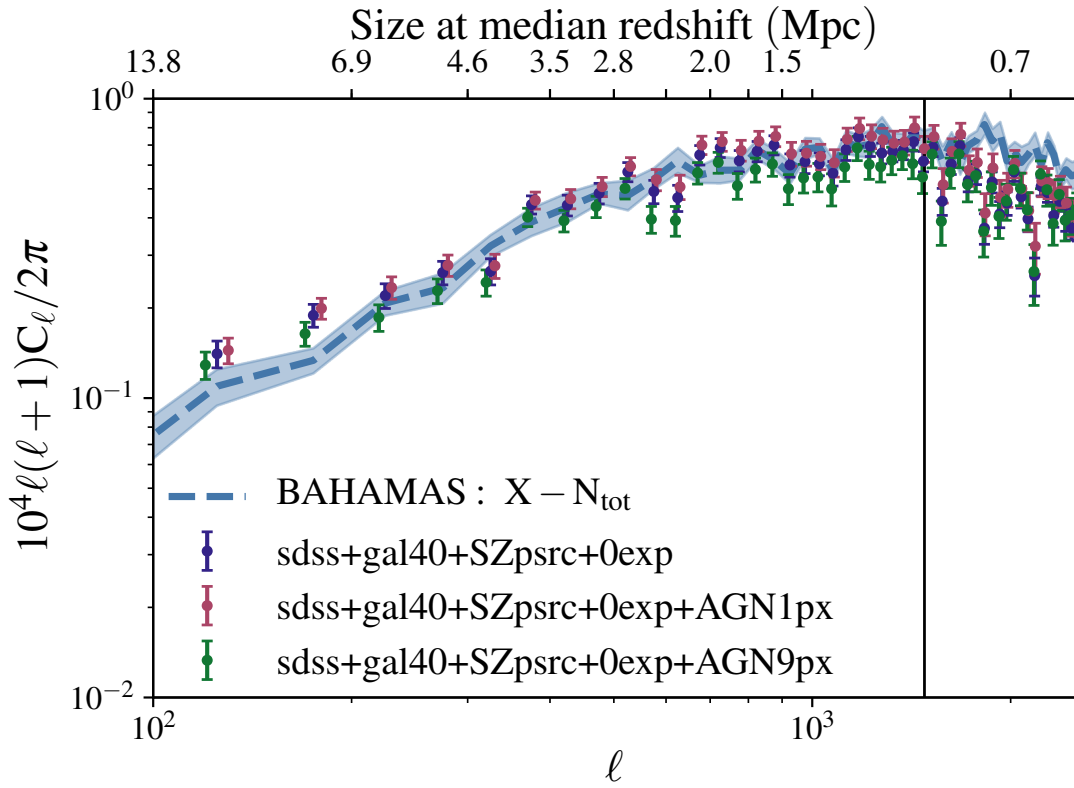


FIGURE A.3: $X - \check{N}_{\text{tot}}$ cross-power spectra computed using different masks for AGN sources in the RASS x-ray map. Navy data points show the standard case from Figure 2.8 where AGN sources are not masked. Pink and navy data points represent 1 pixel and 9 pixel masks for AGN sources, respectively. Blue dashed line shows the equivalent power spectrum from BAHAMAS, which has been calibrated to reproduce hot gas fractions but does not contain AGN sources in X-ray emission. Data points have been artificially offset along the x-axis to make them more visible. A small shift in the cross-power spectrum is visible when masking the AGN, particularly when the larger (9-pixel) mask is employed, suggesting AGN contribute at a sub-dominant level to the total (unmasked) cross-power spectrum.

Appendix B

B.1 Choice of star-forming–quiescent division

The division between star-forming and quenched galaxies is normally located by finding the minimum between the two observed populations in sSFR. However, due to the knowledge of precise SFR, simulated galaxies do not form a distinct ‘red cloud’. A long distribution in sSFR is formed instead, extending all the way to $\text{sSFR} = 0$. Furthermore, the position of the star-forming sequence changes with feedback implementation. The location of a clean division between star-forming and quenched galaxies is, therefore, not obvious. Figure B.1 shows the distributions of field galaxies in both $\text{sSFR} - \log_{10}(M_*)$ space and sSFR histograms for all three simulations. All galaxies with $\text{sSFR} \leq -12$ (definitely not star-forming) have been assigned a random value in the interval $[-13, -12]$. It can be seen from sSFR histograms (right column) that the main sequence is not always at the same location. This is especially true for high-mass galaxies. The situation is further complicated by cluster galaxies, Fig. B.2 shows equivalent distributions for this subset of galaxies. The position of main sequence has shifted to lower values of sSFR for Hydrangea and TNG300, but has stayed approximately the same for BAHAMAS/-MACSIS. This is a feature of feedback and cannot be ignored. A division chosen for field galaxies is not applicable to cluster galaxies, and it does not apply to all simulations equally. As a conservative approach, which avoids labelling star-forming galaxies as quenched, $\text{sSFR} > -12$ is chosen as the division. All three simulations can, therefore, use a single division for field and cluster galaxies, making the comparison consistent.

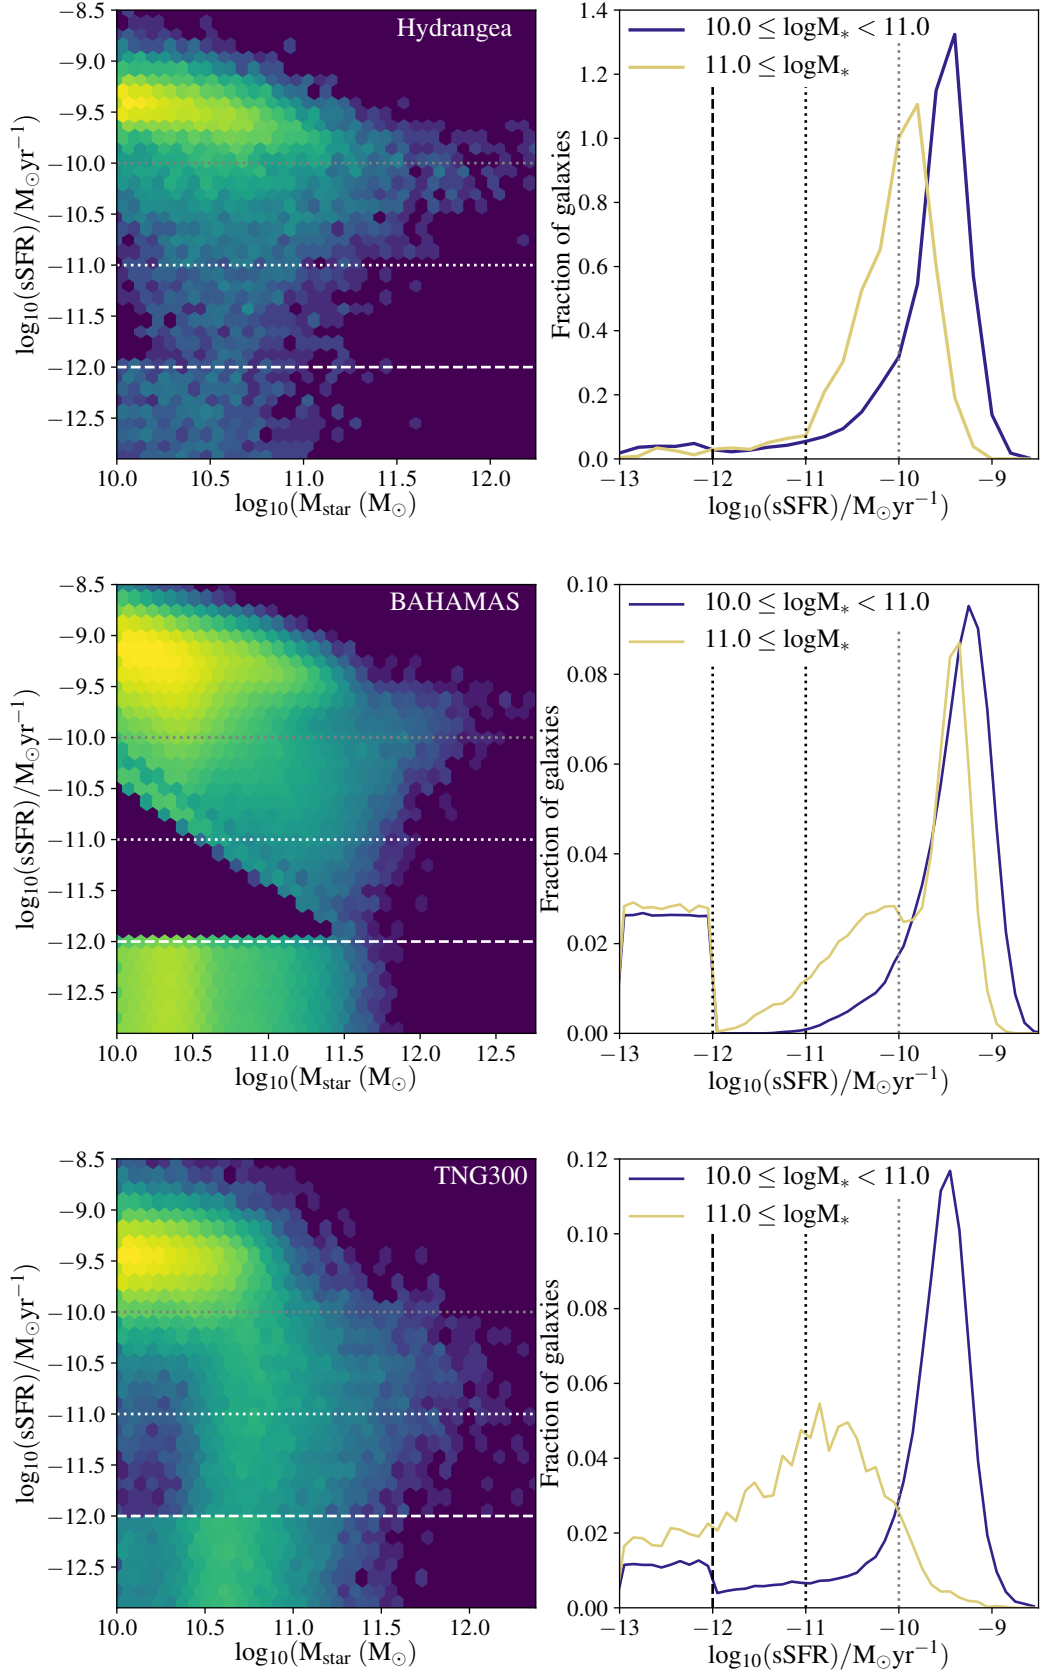


FIGURE B.1: **Left column:** field galaxy density distributions in the $\log_{10}(\text{sSFR}) - \log_{10}(M_{\text{star}})$ space. Low resolution leads to a sharp edge in BAHAMAS distribution.

Right column: Probability distribution functions for the corresponding $\log_{10}(\text{sSFR})$ values. Galaxies with $\log_{10}(\text{sSFR}) \leq -12$ have been assigned a random value in the interval $[-13.0, -12.0]$. The position of main sequence is different for each simulation. As a result, a conservative value of $\log_{10}(\text{sSFR}) > -12$ is chosen to be the division between star-forming and quiescent galaxies.

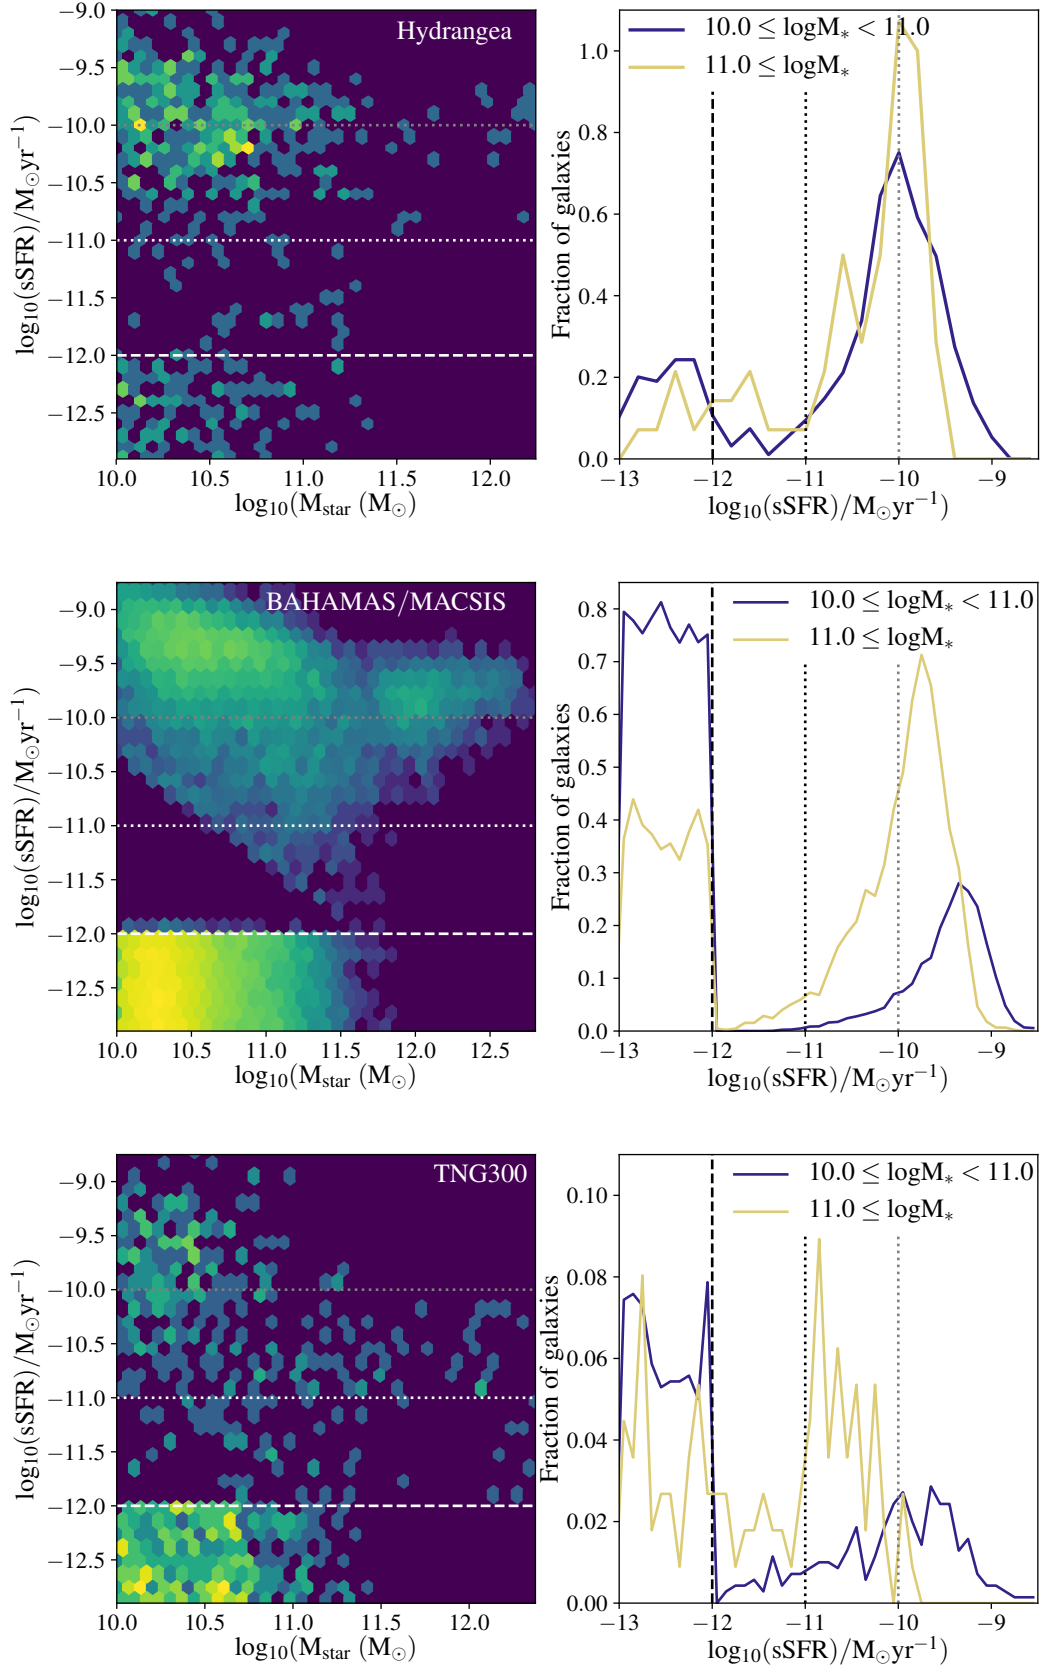


FIGURE B.2: **Left column:** cluster galaxy density distributions in the $\log_{10}(\text{sSFR}) - \log_{10}(M_{\text{star}})$ space. **Right column:** probability distribution functions for the corresponding $\log_{10}(\text{sSFR})$ values. Galaxies with $\log_{10}(\text{sSFR}) \leq -12$ have been assigned a random value in the interval $[-13.0, -12.0]$. The position of main sequence is different for each simulation. As a result, a conservative value of $\log_{10}(\text{sSFR}) > -12$ is chosen to be the division between star-forming and quiescent galaxies.

Bibliography

- Abadi M. G., Moore B., Bower R. G., 1999, MNRAS, 308, 947
- Abazajian K. N., et al., 2009, ApJS, 182, 543
- Abbott T. M. C., et al., 2018, ApJS, 239, 18
- Abell G. O., 1958, ApJS, 3, 211
- Abell G. O., Corwin Harold G. J., Olowin R. P., 1989, ApJS, 70, 1
- Abolfathi B., et al., 2018, ApJS, 235, 42
- Adam R., et al., 2017, A&A, 606, A64
- Agertz O., et al., 2007, MNRAS, 380, 963
- Alam S., et al., 2017, MNRAS, 470, 2617
- Alonso D., Sanchez J., Slosar A., 2018, arXiv e-prints,
- Anderson M. E., Churazov E., Bregman J. N., 2016, MNRAS, 455, 227
- Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, MNRAS, 426, 2046
- Aragón-Calvo M. A., Jones B. J. T., van de Weygaert R., van der Hulst J. M., 2007, A&A, 474, 315
- Aragón-Calvo M. A., van de Weygaert R., Jones B. J. T., 2010, MNRAS, 408, 2163
- Astropy Collaboration et al., 2013, A&A, 558, A33
- Baertschiger T., Joyce M., Sylos Labini F., 2002, ApJ, 581, L63
- Bahé Y. M., McCarthy I. G., 2015, MNRAS, 447, 969

- Bahé Y. M., McCarthy I. G., Balogh M. L., Font A. S., 2013, MNRAS, 430, 3017
- Bahé Y. M., Schaye J., Crain R. A., McCarthy I. G., Bower R. G., Theuns T., McGee S. L., Trayford J. W., 2017a, MNRAS, 464, 508
- Bahé Y. M., et al., 2017b, MNRAS, 470, 4186
- Bahé Y. M., et al., 2019, MNRAS, 485, 2287
- Baldry I. K., 2008, Astronomy and Geophysics, 49, 5.25
- Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Ž., Lupton R. H., Nichol R. C., Szalay A. S., 2004, ApJ, 600, 681
- Baldry I. K., Balogh M. L., Bower R. G., Glazebrook K., Nichol R. C., Bamford S. P., Budavari T., 2006, MNRAS, 373, 469
- Baldry I. K., et al., 2018, MNRAS, 474, 3875
- Balogh M. L., Navarro J. F., Morris S. L., 2000, ApJ, 540, 113
- Balogh M. L., Baldry I. K., Nichol R., Miller C., Bower R., Glazebrook K., 2004, ApJ, 615, L101
- Balogh M. L., et al., 2014, MNRAS, 443, 2679
- Balogh M. L., et al., 2016, MNRAS, 456, 4364
- Balogh M. L., et al., 2017, MNRAS, 470, 4168
- Bamford S. P., et al., 2009, MNRAS, 393, 1324
- Barnes D. J., Kay S. T., Henson M. A., McCarthy I. G., Schaye J., Jenkins A., 2017, MNRAS, 465, 213
- Barnes D. J., et al., 2019, MNRAS, 488, 3003
- Barro G., et al., 2017, ApJ, 851, L40
- Barsanti S., et al., 2018, ApJ, 857, 71
- Baugh C. M., 2006, Reports on Progress in Physics, 69, 3101
- Beck R., Dobos L., Budavári T., Szalay A. S., Csabai I., 2016, MNRAS, 460, 1371

- Beers T. C., Flynn K., Gebhardt K., 1990, *AJ*, 100, 32
- Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
- Bekki K., Couch W. J., Shioya Y., 2002, *ApJ*, 577, 651
- Bell E. F., et al., 2004, *ApJ*, 608, 752
- Bell E. F., Zheng X. Z., Papovich C., Borch A., Wolf C., Meisenheimer K., 2007, *ApJ*, 663, 834
- Berger M. J., Olinger J., 1984, *Journal of Computational Physics*, 53, 484
- Berrier J. C., Bullock J. S., Barton E. J., Guenther H. D., Zentner A. R., Wechsler R. H., 2006, *ApJ*, 652, 56
- Bertschinger E., Gelb J. M., 1991, *Computers in Physics*, 5, 164
- Bigiel F., et al., 2011, *ApJ*, 730, L13
- Biviano A., van der Burg R. F. J., Muzzin A., Sartoris B., Wilson G., Yee H. K. C., 2016, *A&A*, 594, A51
- Blanton M. R., Berlind A. A., 2007, *ApJ*, 664, 791
- Blanton M. R., Moustakas J., 2009, *ARA&A*, 47, 159
- Blanton M. R., et al., 2003, *ApJ*, 594, 186
- Blanton M. R., Eisenstein D., Hogg D. W., Schlegel D. J., Brinkmann J., 2005, *ApJ*, 629, 143
- Blumenthal G. R., Faber S. M., Flores R., Primack J. R., 1986, *ApJ*, 301, 27
- Bode P., Ostriker J. P., 2003, *ApJS*, 145, 1
- Boller T., Freyberg M. J., Trümper J., Haberl F., Voges W., Nandra K., 2016, *A&A*, 588, A103
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Boselli A., Gavazzi G., 2006, *PASP*, 118, 517
- Brammer G. B., et al., 2009, *ApJ*, 706, L173

- Brammer G. B., et al., 2011, *ApJ*, 739, 24
- Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151
- Brough S., et al., 2011, *MNRAS*, 413, 1236
- Brown T., et al., 2017, *MNRAS*, 466, 1275
- Bryant J. J., et al., 2015, *MNRAS*, 447, 2857
- Budzynski J. M., Kuposov S. E., McCarthy I. G., McGee S. L., Belokurov V., 2012, *MNRAS*, 423, 104
- Bykov A. M., Paerels F. B. S., Petrosian V., 2008, *Space Sci. Rev.*, 134, 141
- Byrd G., Valtonen M., 1990, *ApJ*, 350, 89
- Carlberg R. G., et al., 1997, *ApJ*, 485, L13
- Cautun M., van de Weygaert R., Jones B. J. T., 2013, *MNRAS*, 429, 1286
- Cen R., Ostriker J. P., 1999, *ApJ*, 514, 1
- Cervantes Sodi B., Li C., Park C., Wang L., 2016, in Napolitano N. R., Longo G., Marconi M., Paolillo M., Iodice E., eds, Vol. 42, *The Universe of Digital Sky Surveys*. p. 85 ([arXiv:1507.07954](https://arxiv.org/abs/1507.07954)), doi:10.1007/978-3-319-19330-4_12
- Chiang L.-Y., Chen F.-F., 2011, *ApJ*, 738, 188
- Chisari N. E., et al., 2019, *The Open Journal of Astrophysics*, 2, 4
- Chung A., van Gorkom J. H., Kenney J. D. P., Vollmer B., 2007, *ApJ*, 659, L115
- Cimatti A., Daddi E., Renzini A., 2006, *A&A*, 453, L29
- Colella P., Woodward P. R., 1984, *Journal of Computational Physics*, 54, 174
- Colless M., et al., 2001, *MNRAS*, 328, 1039
- Conroy C., 2013, *ARA&A*, 51, 393
- Conroy C., Wechsler R. H., 2009, *ApJ*, 696, 620
- Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201

- Cooper M. C., Newman J. A., Madgwick D. S., Gerke B. F., Yan R., Davis M., 2005, *ApJ*, 634, 833
- Cowan N. B., Ivezić Ž., 2008, *ApJ*, 674, L13
- Cowie L. L., Songaila A., 1977, *Nature*, 266, 501
- Crain R. A., et al., 2015, *MNRAS*, 450, 1937
- Croton D. J., et al., 2005, *MNRAS*, 356, 1155
- Croton D. J., et al., 2006, *MNRAS*, 365, 11
- Cullen L., Dehnen W., 2010, *MNRAS*, 408, 669
- Dalla Vecchia C., Schaye J., 2008a, *MNRAS*, 387, 1431
- Dalla Vecchia C., Schaye J., 2008b, *MNRAS*, 387, 1431
- Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
- Darvish B., Mobasher B., Sobral D., Rettura A., Scoville N., Faisst A., Capak P., 2016, *ApJ*, 825, 113
- Davé R., et al., 2001, *ApJ*, 552, 473
- Davies J. J., Crain R. A., McCarthy I. G., Oppenheimer B. D., Schaye J., Schaller M., McAlpine S., 2019, *MNRAS*, 485, 3783
- Davis M., Geller M. J., 1976, *ApJ*, 208, 13
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
- De Lucia G., Springel V., White S. D. M., Croton D., Kauffmann G., 2006, *MNRAS*, 366, 499
- Dekel A., Birnboim Y., 2006, *MNRAS*, 368, 2
- Dekel A., Silk J., 1986, *ApJ*, 303, 39
- Del Popolo A., Pace F., Lima J. A. S., 2013, *MNRAS*, 430, 628
- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604

- Diemand J., Moore B., Stadel J., 2005, *Nature*, 433, 389
- Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
- Dressler A., 1980, *ApJ*, 236, 351
- Dressler A., Smail I., Poggianti B. M., Butcher H., Couch W. J., Ellis R. S., Oemler Jr. A., 1999, *ApJS*, 122, 51
- Dressler A., Oemler Augustus J., Poggianti B. M., Smail I., Trager S., Shectman S. A., Couch W. J., Ellis R. S., 2004, *ApJ*, 617, 867
- Driver S. P., et al., 2006, *MNRAS*, 368, 414
- Duffell P. C., MacFadyen A. I., 2011, *ApJS*, 197, 15
- Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C., Battye R. A., Booth C. M., 2010, *MNRAS*, 405, 2161
- Durier F., Dalla Vecchia C., 2012, *MNRAS*, 419, 465
- Eke V. R., et al., 2004, *MNRAS*, 355, 769
- Elmegreen B. G., Efremov Y. N., 1997, *ApJ*, 480, 235
- Evans F. A., Parker L. C., Roberts I. D., 2018, *MNRAS*, 476, 5284
- Fabello S., Kauffmann G., Catinella B., Li C., Giovanelli R., Haynes M. P., 2012, *MNRAS*, 427, 2841
- Faber S. M., et al., 2007, *ApJ*, 665, 265
- Fabian A. C., 1994, *ARA&A*, 32, 277
- Fabian A. C., 2012, *ARA&A*, 50, 455
- Fadda D., Girardi M., Giuricin G., Mardirossian F., Mezzetti M., 1996, *ApJ*, 473, 670
- Farouki R., Shapiro S. L., 1981, *ApJ*, 243, 32
- Feldman H. A., Kaiser N., Peacock J. A., 1994, *ApJ*, 426, 23
- Fillingham S. P., Cooper M. C., Wheeler C., Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., 2015, *MNRAS*, 454, 2039

- Foltz R., et al., 2018, *ApJ*, 866, 136
- Forman R., 1998, *Advances in Mathematics*, 134, 90
- Foster A. R., Ji L., Smith R. K., Brickhouse N. S., 2012, *ApJ*, 756, 128
- Fryxell B., et al., 2000, *ApJS*, 131, 273
- Fujita Y., 1998, *ApJ*, 509, 587
- Fujita Y., 2004, *PASJ*, 56, 29
- Furlong M., et al., 2015, *MNRAS*, 450, 4486
- Gallazzi A., Brinchmann J., Charlot S., White S. D. M., 2008, *MNRAS*, 383, 1439
- Gehren T., Fried J., Wehinger P. A., Wyckoff S., 1984, *ApJ*, 278, 11
- Giannantonio T., et al., 2016, *MNRAS*, 456, 3213
- Górski K. M., Hivon E., 2011, HEALPix: Hierarchical Equal Area isoLatitude Pixelization of a sphere, *Astrophysics Source Code Library* (ascl:1107.018)
- Gunn J. E., Gott III J. R., 1972, *ApJ*, 176, 1
- Guo Q., et al., 2011, *MNRAS*, 413, 101
- Haines C. P., et al., 2015, *ApJ*, 806, 101
- Harnois-Déraps J., et al., 2017, *MNRAS*, 471, 1619
- Hashimoto D., Nishizawa A. J., Shirasaki M., Macias O., Horiuchi S., Tashiro H., Oguri M., 2019, *MNRAS*, 484, 5256
- Hasinger G., Miyaji T., Schmidt M., 2005, *A&A*, 441, 417
- Hatfield P. W., Jarvis M. J., 2017, *MNRAS*, 472, 3570
- Heckman T. M., Best P. N., 2014, *ARA&A*, 52, 589
- Hill J. C., Spergel D. N., 2014, *Journal of Cosmology and Astro-Particle Physics*, 2014, 030
- Hinshaw G., et al., 2009, *ApJS*, 180, 225
- Hinshaw G., et al., 2013, *ApJS*, 208, 19

- Hirschmann M., De Lucia G., Wilman D., Weinmann S., Iovino A., Cucciati O., Zibetti S., Villalobos Á., 2014, *MNRAS*, 444, 2938
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
- Hockney R. W., Eastwood J. W., 1988, *Computer simulation using particles*
- Hoekstra H., Mahdavi A., Babul A., Bildfell C., 2012, *MNRAS*, 427, 1298
- Hogg D. W., 1999, arXiv e-prints, pp astro-ph/9905116
- Hogg D. W., et al., 2004, *ApJ*, 601, L29
- Hojjati A., et al., 2016, preprint, ([arXiv:1608.07581](https://arxiv.org/abs/1608.07581))
- Holmberg E., 1941, *ApJ*, 94, 385
- Hopkins P. F., 2013, *MNRAS*, 428, 2840
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950
- Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, 421, 3522
- Hubble E. P., 1925, *ApJ*, 62, 409
- Hubble E. P., 1926, *ApJ*, 64, 321
- Huchra J. P., Geller M. J., 1982, *ApJ*, 257, 423
- Hunter J. D., 2007, *Computing in Science and Engineering*, 9, 90
- Hurier G., Douspis M., Aghanim N., Pointecouteau E., Diego J. M., Macias-Perez J. F., 2015, *A&A*, 576, A90
- Hurier G., Singh P., Hernández-Monteagudo C., 2019, *A&A*, 625, L4
- Jaffé Y. L., et al., 2016, *MNRAS*, 461, 1202
- Jones E., Oliphant T., Peterson P., 2001, *SciPy: Open source scientific tools for Python*, <http://www.scipy.org/>
- Jones B. J. T., van de Weygaert R., Aragón-Calvo M. A., 2010, *MNRAS*, 408, 897

- Kannan R., Vogelsberger M., Stinson G. S., Hennawi J. F., Marinacci F., Springel V., Macciò A. V., 2016, *MNRAS*, 458, 2516
- Kauffmann G., White S. D. M., Guiderdoni B., 1993, *MNRAS*, 264, 201
- Kauffmann G., et al., 2003, *MNRAS*, 341, 54
- Kauffmann G., White S. D. M., Heckman T. M., Ménard B., Brinchmann J., Charlot S., Tremonti C., Brinkmann J., 2004, *MNRAS*, 353, 713
- Kawata D., Mulchaey J. S., 2008, *ApJ*, 672, L103
- Kawinwanichakij L., et al., 2017, *ApJ*, 847, 134
- Kenney J. D. P., Abramson A., Bravo-Alfaro H., 2015, *AJ*, 150, 59
- Kennicutt Robert C. J., 1998, *ARA&A*, 36, 189
- Kereš D., Katz N., Weinberg D. H., Davé R., 2005, *MNRAS*, 363, 2
- Kimm T., et al., 2009, *MNRAS*, 394, 1131
- Kodama T., Smail I., Nakata F., Okamura S., Bower R. G., 2001, *ApJ*, 562, L9
- Kormendy J., Ho L. C., 2013, *ARA&A*, 51, 511
- Kormendy J., Richstone D., 1995, *ARA&A*, 33, 581
- Koukoufilippas N., Alonso D., Bilicki M., Peacock J. A., 2020, *MNRAS*, 491, 5464
- Kraft R. P., et al., 2017, *ApJ*, 848, 27
- Krolik J. H., 1999, *Active galactic nuclei : from the central black hole to the galactic environment*
- Krumholz M. R., Dekel A., McKee C. F., 2012, *ApJ*, 745, 69
- Krumpe M., Miyaji T., Coil A. L., Aceves H., 2012, *ApJ*, 746, 1
- Kukstas E., McCarthy I. G., Baldry I. K., Font A. S., 2019, *arXiv e-prints*, p. arXiv:1907.06645
- Lacey C., Cole S., 1993, *MNRAS*, 262, 627
- Lahav O., Lilje P. B., Primack J. R., Rees M. J., 1991, *MNRAS*, 251, 128

- Larson R. B., Tinsley B. M., Caldwell C. N., 1980, *ApJ*, 237, 692
- Li I. H., et al., 2011, *MNRAS*, 411, 1869
- Libeskind N. I., et al., 2018, *MNRAS*, 473, 1195
- Lin Y.-T., Mohr J. J., Stanford S. A., 2004, *ApJ*, 610, 745
- Liu C., Hao L., Wang H., Yang X., 2019, *ApJ*, 878, 69
- Lotz M., Remus R.-S., Dolag K., Biviano A., Burkert A., 2019, *MNRAS*, 488, 5370
- Lu T., Gilbank D. G., McGee S. L., Balogh M. L., Gallagher S., 2012, *MNRAS*, 420, 126
- Madau P., Dickinson M., 2014, *ARA&A*, 52, 415
- Maier C., Ziegler B. L., Haines C. P., Smith G. P., 2019, *A&A*, 621, A131
- Makiya R., Ando S., Komatsu E., 2018, *MNRAS*, 480, 3928
- Mamon G. A., Biviano A., Boué G., 2013, *MNRAS*, 429, 3079
- Marchesini D., van Dokkum P. G., Förster Schreiber N. M., Franx M., Labbé I., Wuyts S., 2009, *ApJ*, 701, 1765
- Marinacci F., et al., 2018, *MNRAS*, 480, 5113
- Marinoni C., Davis M., Newman J. A., Coil A. L., 2002, *ApJ*, 580, 122
- Mayer L., Mastropietro C., Wadsley J., Stadel J., Moore B., 2006, *MNRAS*, 369, 1021
- Mayer L., Governato F., Kaufmann T., 2008, *Advanced Science Letters*, 1, 7
- McAlpine S., et al., 2016, *Astronomy and Computing*, 15, 72
- McCarthy I. G., Frenk C. S., Font A. S., Lacey C. G., Bower R. G., Mitchell N. L., Balogh M. L., Theuns T., 2008, *MNRAS*, 383, 593
- McCarthy I. G., Le Brun A. M. C., Schaye J., Holder G. P., 2014, *MNRAS*, 440, 3645
- McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, 465, 2936
- McCarthy I. G., Bird S., Schaye J., Harnois-Deraps J., Font A. S., van Waerbeke L., 2018, *MNRAS*, 476, 2999

- McGee S. L., Balogh M. L., Bower R. G., Font A. S., McCarthy I. G., 2009, MNRAS, 400, 937
- McGee S. L., Balogh M. L., Wilman D. J., Bower R. G., Mulchaey J. S., Parker L. C., Oemler A., 2011, MNRAS, 413, 996
- McGee S. L., Bower R. G., Balogh M. L., 2014, MNRAS, 442, L105
- Merritt D., 1983, ApJ, 264, 24
- Merritt D., 1984, ApJ, 276, 26
- Mihos J. C., 2004, Clusters of Galaxies: Probes of Cosmological Structure and Galaxy Evolution, p. 277
- Mihos J. C., Hernquist L., 1994a, ApJ, 425, L13
- Mihos J. C., Hernquist L., 1994b, ApJ, 431, L9
- Miller R. H., 1986, A&A, 167, 41
- Miyaji T., Hasinger G., Schmidt M., 2001, A&A, 369, 49
- Mo H. J., Mao S., White S. D. M., 1998, MNRAS, 295, 319
- Monaghan J. J., Lattanzio J. C., 1985, A&A, 149, 135
- Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, Nature, 379, 613
- Moore B., Lake G., Katz N., 1998, ApJ, 495, 139
- Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJ, 524, L19
- Morse M., 1934, Proceedings of the National Academy of Science, 20, 46
- Moster B. P., Naab T., White S. D. M., 2018, MNRAS, 477, 1822
- Muldrew S. I., et al., 2012, MNRAS, 419, 2670
- Muzzin A., et al., 2012, ApJ, 746, 188
- Muzzin A., et al., 2013, ApJ, 777, 18
- Muzzin A., et al., 2014, ApJ, 796, 65

- Naiman J. P., et al., 2018, MNRAS, 477, 1206
- Nandra K., et al., 2007, ApJ, 660, L11
- Navarro J. F., Frenk C. S., White S. D. M., 1995, MNRAS, 275, 720
- Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
- Neistein E., Li C., Khochfar S., Weinmann S. M., Shankar F., Boylan-Kolchin M., 2011, MNRAS, 416, 1486
- Nelson D., et al., 2018, MNRAS, 475, 624
- Nicastro F., Mathur S., Elvis M., 2008, Science, 319, 55
- Noeske K. G., et al., 2007, ApJ, 660, L47
- Nulsen P. E. J., 1982, MNRAS, 198, 1007
- O'Shea B. W., Bryan G., Bordner J., Norman M. L., Abel T., Harkness R., Kritsuk A., 2004, arXiv e-prints, pp astro-ph/0403044
- Oemler Jr. A., 1974, ApJ, 194, 1
- Oman K. A., Hudson M. J., 2016, MNRAS, 463, 3083
- Pakmor R., Springel V., 2013, MNRAS, 432, 176
- Pakmor R., Bauer A., Springel V., 2011, MNRAS, 418, 1392
- Pandey S., Baxter E. J., Hill J. C., 2020, Phys. Rev. D, 101, 043525
- Papovich C., et al., 2018, ApJ, 854, 30
- Paradis D., Dobashi K., Shimoikura T., Kawamura A., Onishi T., Fukui Y., Bernard J.-P., 2012, A&A, 543, A103
- Park C., Choi Y.-Y., Vogeley M. S., Gott J. Richard I., Blanton M. R., SDSS Collaboration 2007, ApJ, 658, 898
- Pasquali A., Gallazzi A., van den Bosch F. C., 2012, MNRAS, 425, 273
- Peebles P. J. E., 1975, ApJ, 196, 647
- Peebles P. J. E., 1980, The large-scale structure of the universe

- Peebles P. J., Ratra B., 2003, *Reviews of Modern Physics*, 75, 559
- Peng Y.-j., et al., 2010, *ApJ*, 721, 193
- Peng Y.-j., Lilly S. J., Renzini A., Carollo M., 2012, *ApJ*, 757, 4
- Peng Y., Maiolino R., Cochrane R., 2015, *Nature*, 521, 192
- Perez F., Granger B. E., 2007, *Computing in Science and Engineering*, 9, 21
- Perez F., Granger B. E., Hunter J. D., 2011, *Computing in Science and Engineering*, 13, 13
- Pillepich A., et al., 2018a, *MNRAS*, 473, 4077
- Pillepich A., et al., 2018b, *MNRAS*, 475, 648
- Pimblet K. A., Smail I., Edge A. C., O'Hely E., Couch W. J., Zabludoff A. I., 2006, *MNRAS*, 366, 645
- Pintos-Castro I., Yee H. K. C., Muzzin A., Old L., Wilson G., 2019, *ApJ*, 876, 40
- Planck Collaboration et al., 2014a, *A&A*, 571, A15
- Planck Collaboration et al., 2014b, *A&A*, 571, A16
- Planck Collaboration et al., 2014c, *A&A*, 571, A18
- Planck Collaboration et al., 2014d, *A&A*, 571, A21
- Planck Collaboration et al., 2016a, *A&A*, 594, A13
- Planck Collaboration et al., 2016b, *A&A*, 594, A22
- Poggianti B. M., et al., 2016, *AJ*, 151, 78
- Poggianti B. M., et al., 2017, *ApJ*, 844, 48
- Postman M., Geller M. J., 1984, *ApJ*, 281, 95
- Postman M., et al., 2012, *ApJS*, 199, 25
- Prat J., et al., 2018, *Phys. Rev. D*, 98, 042005
- Price D. J., 2008, *Journal of Computational Physics*, 227, 10040

- Quilis V., Moore B., Bower R., 2000, *Science*, 288, 1617
- Ramatsoku M., et al., 2019, *MNRAS*, 487, 4580
- Rasmussen J., Mulchaey J. S., Bai L., Ponman T. J., Raychaudhury S., Dariush A., 2012, *ApJ*, 757, 122
- Reimers D., 2002, *Space Sci. Rev.*, 100, 89
- Roberts I. D., Parker L. C., 2017, *MNRAS*, 467, 3268
- Roberts I. D., Parker L. C., Brown T., Joshi G. D., Hlavacek-Larrondo J., Wadsley J., 2019, *ApJ*, 873, 42
- Robotham A. S. G., et al., 2011, *MNRAS*, 416, 2640
- Salim S., et al., 2007, *ApJS*, 173, 267
- Sand D. J., et al., 2012, *ApJ*, 746, 163
- Scannapieco C., et al., 2012, *MNRAS*, 423, 1726
- Schaap W. E., van de Weygaert R., 2000, *A&A*, 363, L29
- Schawinski K., et al., 2014, *MNRAS*, 440, 889
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J., et al., 2010, *MNRAS*, 402, 1536
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Schechter P., 1976, *ApJ*, 203, 297
- Schiminovich D., et al., 2007, *ApJS*, 173, 315
- Schmidt M., 1959, *ApJ*, 129, 243
- Sedgwick T. M., Baldry I. K., James P. A., Kelvin L. S., 2019, *MNRAS*,
- Shapley H., Ames A., 1926, *Harvard College Observatory Circular*, 294, 1
- Silk J., Rees M. J., 1998, *A&A*, 331, L1
- Simard L., Mendel J. T., Patton D. R., Ellison S. L., McConnachie A. W., 2011, *ApJS*, 196, 11

- Simionescu A., et al., 2011, *Science*, 331, 1576
- Smith S., 1936, *ApJ*, 83, 23
- Smith R. K., Brickhouse N. S., Liedahl D. A., Raymond J. C., 2001, *ApJ*, 556, L91
- Somerville R. S., Davé R., 2015, *ARA&A*, 53, 51
- Soucail G., Mellier Y., Fort B., Hammer F., Mathez G., 1987, *A&A*, 184, L7
- Sousbie T., 2011, *MNRAS*, 414, 350
- Sousbie T., Pichon C., Kawahara H., 2011, *MNRAS*, 414, 384
- Sparke L. S., Gallagher John S. I., 2007, *Galaxies in the Universe*
- Spitzer Lyman J., Baade W., 1951, *ApJ*, 113, 413
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., Yoshida N., White S. D. M., 2001, *New Astron.*, 6, 79
- Springel V., Di Matteo T., Hernquist L., 2005, *MNRAS*, 361, 776
- Springel V., et al., 2018, *MNRAS*, 475, 676
- Staniszewski Z., et al., 2009, *ApJ*, 701, 32
- Stoughton C., et al., 2002, *AJ*, 123, 485
- Strateva I., et al., 2001, *AJ*, 122, 1861
- Sunyaev R. A., Zeldovich Y. B., 1969, *Nature*, 223, 721
- Szapudi I., Prunet S., Pogosyan D., Szalay A. S., Bond J. R., 2001, *ApJ*, 548, L115
- Tanimura H., et al., 2019, *MNRAS*, 483, 223
- Taranu D. S., Dubinski J., Yee H. K. C., 2013, *ApJ*, 778, 61
- Taylor E. N., et al., 2011, *MNRAS*, 418, 1587
- Taylor E. N., et al., 2015, *MNRAS*, 446, 2144

- Teyssier R., 2002, *A&A*, 385, 337
- Tonnesen S., Bryan G. L., 2012, *MNRAS*, 422, 1609
- Tonnesen S., Stone J., 2014, *ApJ*, 795, 148
- Treu T., Ellis R. S., Kneib J.-P., Dressler A., Smail I., Czoske O., Oemler A., Natarajan P., 2003, *ApJ*, 591, 53
- Tristram M., Macías-Pérez J. F., Renault C., Santos D., 2005, *MNRAS*, 358, 833
- Tyson J. A., Valdes F., Wenk R. A., 1990, *ApJ*, 349, L1
- Umehata H., et al., 2019, *Science*, 366, 97
- Vale A., Ostriker J. P., 2006, *MNRAS*, 371, 1173
- Valluri M., 1994, *ApJ*, 430, 101
- Vandenbroucke B., De Rijcke S., 2016, *Astronomy and Computing*, 16, 109
- Véron-Cetty M. P., Véron P., 2010, *A&A*, 518, A10
- Vikram V., Lidz A., Jain B., 2017, *MNRAS*, 467, 2315
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, *Nature Reviews Physics*, 2, 42
- Voges W., 1993, *Advances in Space Research*, 13, 391
- Voges W., et al., 1999, *VizieR Online Data Catalog*, 9010
- Voit G. M., Kay S. T., Bryan G. L., 2005, *MNRAS*, 364, 909
- Wadsley J. W., Stadel J., Quinn T., 2004, *New Astron.*, 9, 137
- Walcher J., Groves B., Budavári T., Dale D., 2011, *Ap&SS*, 331, 1
- Weinberg S., Bond J. R., 2009, *Cosmology*. Vol. 62, doi:10.1063/1.3156334,
- Weinberger R., et al., 2017, *MNRAS*, 465, 3291
- 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
- Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2014, MNRAS, 439, 2687
- Whitaker K. E., et al., 2011, ApJ, 735, 86
- White S. D. M., Frenk C. S., 1991, ApJ, 379, 52
- White S. D. M., Rees M. J., 1978, MNRAS, 183, 341
- White S. D. M., et al., 2005, A&A, 444, 365
- Wiersma R. P. C., Schaye J., Smith B. D., 2009a, MNRAS, 393, 99
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, MNRAS, 399, 574
- Wilman D. J., Balogh M. L., Bower R. G., Mulchaey J. S., Oemler A. J., Carlberg R. G., 2005, arXiv e-prints, pp astro-ph/0510077
- Wilman D. J., Zibetti S., Budavári T., 2010, MNRAS, 406, 1701
- Wuyts S., et al., 2011, ApJ, 742, 96
- Yang X., Mo H. J., van den Bosch F. C., Jing Y. P., 2005, MNRAS, 356, 1293
- Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, ApJ, 671, 153
- Yang X., Mo H. J., van den Bosch F. C., 2009, ApJ, 695, 900
- Yang X., Mo H. J., van den Bosch F. C., Zhang Y., Han J., 2012, ApJ, 752, 41
- Yee H. K. C., Ellingson E., Carlberg R. G., 1996, ApJS, 102, 269
- York D. G., et al., 2000, AJ, 120, 1579
- Zavala J., Balogh M. L., Afshordi N., Ro S., 2012, MNRAS, 426, 3464
- Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, Journal of Open Source Software, 4, 1298
- Zwicky F., 1933, Helvetica Physica Acta, 6, 110

- da Silva A. C., Barbosa D., Liddle A. R., Thomas P. A., 2000, *MNRAS*, 317, 37
- van Daalen M. P., Schaye J., Booth C. M., Dalla Vecchia C., 2011, *MNRAS*, 415, 3649
- van Daalen M. P., McCarthy I. G., Schaye J., 2020, *MNRAS*, 491, 2424
- van den Bosch F. C., 1998, *ApJ*, 507, 601
- van der Burg R. F. J., et al., 2013, *A&A*, 557, A15
- van der Burg R. F. J., Muzzin A., Hoekstra H., Wilson G., Lidman C., Yee H. K. C., 2014, *A&A*, 561, A79
- van der Burg R. F. J., Hoekstra H., Muzzin A., Sifón C., Balogh M. L., McGee S. L., 2015, *A&A*, 577, A19
- van der Burg R. F. J., McGee S., Aussel H., Dahle H., Arnaud M., Pratt G. W., Muzzin A., 2018, *A&A*, 618, A140
- van der Burg R. F. J., et al., 2020, arXiv e-prints, p. arXiv:2004.10757
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Computing in Science and Engineering*, 13, 22
- van der Wel A., et al., 2014, *ApJ*, 788, 28