Insights from simulations on the consequences of uncertainties in estimating the masses of observed galaxies

CAMPBELL, DAVID,JAMES,ROWNEY

Insights from simulations on the consequences of uncertainties in estimating the masses of observed galaxies

David James Rowney Campbell

Abstract

We make use of cutting-edge simulations of galaxy formation in a Λ cold dark matter (ΛCDM) Universe to investigate the impact of the uncertainties inherent to certain observational techniques for estimating the masses of galaxies on the conclusions that are drawn from studies using such methods. By performing virtual 'observations' of simulated galaxies, we estimate their stellar and dynamical masses in the same way as in particular observational studies. The satellite galaxies of the Milky Way are highly attractive candidates for dynamical studies, due to their proximity; and in general, satellite galaxies dominate the clustering of galaxies on small scales. The total dynamical masses and internal mass distributions of individual galaxies, along with the clustering of galaxies as a function of intrinsic properties such as stellar mass, each reflect the structure and evolution history of the underlying invisible dark matter that forms the structural spine of the Universe and incubates the formation and evolution of galaxies over cosmic time.

The observed stellar kinematics of dispersion-supported galaxies are often used to measure dynamical masses. Recently, several analytical relationships between the stellar line-of-sight velocity dispersion, the projected (2D) or deprojected (3D) half-light radius, and the total mass enclosed within the half-light radius, relying on the spherical Jeans equation, have been proposed. Here, we make use of the APOSTLE cosmological hydrodynamical simulations of the Local Group to test the validity and accuracy of such mass estimators for both dispersion and rotation-supported galaxies, for field and satellite galaxies, and for galaxies of varying masses, shapes, and velocity dispersion anisotropies. We find that the mass estimators of Walker et al. and Wolf et al. are able to recover the masses of dispersion-dominated systems with little systematic bias, but with a 1σ scatter of 25 and 23 percent, respectively. The error on the estimated mass is dominated by the impact of

the 3D shape of the stellar mass distribution, which is difficult to constrain observationally. This intrinsic scatter becomes the dominant source of uncertainty in the masses estimated for galaxies like the dwarf spheroidal (dSph) satellites of the Milky Way, where the observational errors in their sizes and velocity dispersions are small. Such scatter may also affect the inner density profile slopes of dSphs as derived from multiple stellar populations, relaxing the significance with which Navarro-Frenk-White profiles may be excluded, depending on the degree to which the relevant properties of the different stellar populations are correlated. Additionally, we derive a new optimal mass estimator that removes the residual biases and achieves a statistically significant reduction in the scatter to 20 percent overall for dispersion-dominated galaxies, allowing more precise and accurate mass estimates.

We present predictions for the two-point correlation function of galaxy clustering as a function of stellar mass, computed using two new versions of the GALFORM semi-analytic galaxy formation model. One model uses a universal stellar initial mass function (IMF), while the other assumes different IMFs for quiescent star formation and bursts. Particular consideration is given to how the assumptions required to estimate the stellar masses of observed galaxies (such as the choice of IMF, stellar population synthesis model, and dust extinction) influence the perceived dependence of galaxy clustering on stellar mass. Broad-band spectral energy distribution fitting is carried out to estimate stellar masses for the model galaxies in the same manner as in observational studies. We show clear differences between the clustering signals computed using the true and estimated model stellar masses. As such, we highlight the importance of applying our methodology to compare theoretical models to observations. We introduce an alternative scheme for the calculation of the merger time-scales for satellite galaxies in GALFORM, which takes into account the dark matter subhalo information from the underlying dark matter only simulation. This reduces the amplitude of small-scale clustering. The new merger scheme offers improved or similar agreement with observational clustering measurements, over the redshift range $0 < z < 0.7$. We find reasonable agreement with clustering measurements from GAMA, but find larger discrepancies for some stellar mass ranges and separation scales with respect to measurements from SDSS and VIPERS, depending on the GALFORM model used.

Insights from simulations on the consequences of uncertainties in estimating the masses of observed galaxies

David James Rowney Campbell

A thesis submitted to Durham University
in accordance with the regulations
for admittance to the degree of
Doctor of Philosophy

Department of Physics
Durham University
United Kingdom

2017

Contents

List of Tables

List of Figures

List of Abbreviations

ΛCDM Λ Cold Dark Matter [Λ = cosmological constant]

2D Two-dimensional

3D Three-dimensional

AGN Active Galactic Nucleus/Nuclei

AMR Adaptive Mesh Refinement

BAO Baryon Acoustic Oscillation

CDM Cold Dark Matter

CFHTLS Canada-France-Hawaii Telescope Legacy Survey

CMB Cosmic Microwave Background

COSMOS COSMic evolution Survey

DMO Dark Matter Only [simulation]

dSph dwarf Spheroidal [galaxy]

FOF Friends Of Friends [algorithm/halo]

GALEX GALaxy Evolution Explorer

GAMA Galaxy And Mass Assembly survey

Gon14 Gonzalez-Perez et al. (2014) [GALFORM model]

HDM Hot Dark Matter

HOD Halo Occupation Distribution

IMF Initial Mass Function

IR Infra-Red

Lac14 Lacey et al. (2016) [GALFORM model]

LMC Large Magellanic Cloud

M31 Andromeda Galaxy

MCMC Markov Chain Monte Carlo [algorithm]

MW Milky Way

NFW Navarro, Frenk & White (1996b, 1997) [density profile]

NIR Near Infra-Red

Planck2015 Full data from Planck satellite mission

SDF Subhalo Dynamical Friction [new GALFORM merger scheme]

SDSS Sloan Digital Sky Survey

SED Spectral Energy Distribution

SFH Star Formation History

SIDM Self-Interacting Dark Matter

SMC Small Magellanic Cloud

SNIa Type Ia Supernova(e)

SPH Smoothed Particle Hydrodynamics

SPS Stellar Population Synthesis [model]

SSP Simple Stellar Population

SWIRE Spitzer Wide-area Infra-Red Extragalactic survey

TP-AGB Thermally Pulsing Asymptotic Giant Branch

UKIDSS UKIRT Deep Sky Survey

UKIRT United Kingdom Infra-Red Telescope

UV Ultra-Violet

VIMOS VIsible Multi-Object Spectrograph

VIPERS VIMOS Public Extragalactic Redshift Survey

WDM Warm Dark Matter

WIMP Weakly Interacting Massive Particle

WIRCAM Wide-field Infra-Red CAMera

WMAP Wilkinson Microwave Anisotropy Probe

WMAP7 WMAP seven-year data

zCOSMOS COSMOS redshift survey with VIMOS

Declaration

This thesis is based on work carried out while the author was a postgraduate research student under the supervision of Carlos Frenk, Adrian Jenkins, Vincent Eke, and Carlton Baugh, at the Institute for Computational Cosmology, Department of Physics, Durham University, from 2013 to 2016.

All text and figures contained in this thesis have been produced by the author (except where explicitly stated otherwise), and the data analysis described in Chapters 2 and 3 (and Appendix B) is the direct work of the author, in collaboration with the supervisors named above, along with other researchers at Durham University and further afield. Chapter 2 (with Appendix B) of this thesis has been published in the form of a paper in the Monthly Notices of the Royal Astronomical Society (Campbell et al., 2017), and Chapter 3 (with Appendix C) has also been published as a paper in the same journal (Campbell et al., 2015). The author contributed to the development of the semi-analytic galaxy formation model described by Gonzalez-Perez et al. (2014), which is used extensively in Chapter 3.

The material contained in this thesis has not been submitted previously for a degree at Durham University, or any other institution.

Acknowledgements

I would like to thank each of my supervisors for their constant support and encouragement; their expert guidance and insight has made the work presented here possible. There are, of course, many other members of staff and students in Durham and throughout the Virgo Consortium who have helped me along the way, and made Durham a great place to be.

I gratefully acknowledge the support of a Science and Technology Facilities Council (STFC) studentship. This work was further supported by STFC (grant number ST/L00075X/1) and the European Research Council (grant number GA 267291). This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure.

In the beginning there was nothing, which exploded.

— Terry Pratchett

Chapter 1

Introduction

1.1 Towards a standard cosmological model

Since prehistoric times, human imagination has been captivated by the mysteries of the heavens, evoking countless explanations for the nature of the cosmos and our place within it. Our modern understanding of the Universe is rooted in the assumption of the *cosmological principle*: that on large scales the Universe is homogeneous and isotropic (such that observers located at different points in space will see more or less the same thing, in a statistical sense, at sufficiently large distances, in all directions).

Observational measurements of the rotation velocity profiles of stars and gas within disc galaxies, along with dynamical studies of hot gas and constituent galaxies within galaxy clusters, and measurements of the stellar velocity dispersions of spheroidal and elliptical galaxies, provide strong evidence that there is not enough luminous material within such systems to fully account for their gravitational dynamics. That is, significant contributions to the gravitational field are inferred to originate from matter that cannot be seen directly through electromagnetic radiation. This apparently invisible material is known as *dark matter*, and has been a topic of increasing interest in cosmology since the 1930s (Zwicky, 1933; Babcock, 1939; Kahn & Woltjer,

1959; Rubin & Ford, 1970; Ostriker & Peebles, 1973; Ostriker et al., 1974; Einasto et al., 1974; White & Rees, 1978). Dynamical studies suggest that galaxies (and clusters of galaxies) are embedded in overdense *haloes* of dark matter, which have approximately spherically symmetric mass distributions that increase in density towards the halo centre.

It is usually assumed that the cosmological dark matter is a fundamental particle that interacts only gravitationally with other particles (and possibly also via the weak interaction[1]). As such, the dark matter can be treated as a *collisionless fluid*; except in the (extremely hot) very early Universe, when even the dark matter particles will have been presumably in thermal equilibrium with other particles, ceasing at the time of *decoupling* of the dark matter.[2] For a collisionless fluid, density perturbations can be *erased* by random thermal motions; this effect, known as *free streaming* (or *collisionless damping*), competes with the self-gravity of the perturbations (where the self-gravity acts to increase perturbation amplitudes).

The mass of the hypothetical dark matter particle has important implications for the details of the formation of structure in the Universe, since the dark matter has presumably always dominated the matter density, and thus dominated the gravitational forces responsible for the evolution of structure (after very early times, when radiation, not matter, dominated the energy density of the Universe). A relatively less massive dark matter particle would have remained relativistic until later times as the infant Universe expanded and cooled, allowing such particles to travel (free stream) larger distances before slowing to low velocity, and to erase (smooth out) density variations on relatively large scales. Therefore, relatively 'light' dark matter particle candidates are said to have a relatively long *free streaming length*, and are referred to as *hot* dark matter (with rest mass energy $mc^2 \sim$ eV, where c is the speed of light; e.g. *massive neutrinos*). On the other hand, if the dark matter particle is relatively *cold*, i.e. it has a relatively high mass ($mc^2 \sim$ GeV; e.g. the *lightest supersymmetric particle*; Ellis et al. 1984), then such particles would have become non-relativistic much earlier in the expansion history of the Universe, thus allowing density fluctuations on smaller scales to survive

[1] The acronym WIMP (*weakly interacting massive particle*) is often used to refer to dark matter particle candidates of this kind.

[2] A notable exception is if the dark matter is composed of *axions* (Preskill et al., 1983). In this case, the dark matter was never in thermal equilibrium with other particles.

beyond very early times and to grow under gravity. Indeed, free streaming is suppressed sufficiently by the particle mass in the case of cold dark matter such that the influence of this effect on cosmic structure formation can be neglected. If the dark matter particles were somehow created with zero kinetic energy, then they would qualify as cold dark matter, even for low particle masses (this is the case for *axions*; Preskill et al. 1983).

Whether the dark matter is hot (HDM) or cold (CDM) has a direct impact on the size (mass) of structures that can form directly through the gravitational clustering of dark matter particles, assuming an initial density field with small perturbations around the cosmic mean. Free streaming in the case of HDM prevents the direct assembly of haloes smaller than those of large galaxy clusters, because density perturbations on smaller scales are erased at early times. However, dark haloes as small as an Earth mass are able to form directly in CDM. This distinction between CDM and HDM leads to very different pictures of how structure formation takes place in each scenario. For HDM, very massive haloes form first, which must then fragment to produce the haloes of lower masses required to host small clusters and individual galaxies. For CDM, structure formation proceeds through hierarchical clustering, where the smallest haloes form earliest, and merge sequentially to form progressively larger structures.

Comparing the large-scale structure predicted by CDM and HDM cosmological simulations with observational maps of the distribution of galaxies clearly indicates that the observed clustering is inconsistent with the predictions assuming HDM, but closely resembles the predictions for a Universe whose matter density is dominated by CDM (White et al., 1983; Davis et al., 1985). Thus, CDM has emerged as the standard theoretical description to account for the invisible mass that we infer to pervade the Universe, where cosmic structure formation is understood to have been governed by the hierarchical clustering of cold dark matter particles into progressively larger haloes and filaments (the *cosmic web*), following the gravitational instability of an initial, almost smooth, density field with small perturbations. Baryonic gas is drawn into the gravitational potential wells of dark matter haloes, eventually condensing at sufficient densities in the halo centres to ignite star formation, and the birth of galaxies.

Recently there has been increasing interest in flavours of dark matter

that differ from the behaviour of CDM on the smallest scales of cosmological interest. If the dark matter is *warm* (WDM; $mc^2 \sim$ keV; e.g. *sterile neutrinos*; Shaposhnikov 2008), then the hierarchical structure formation scenario of CDM is preserved, but dark matter haloes smaller than those of dwarf galaxies do not form; and so the dark matter distribution is smooth on scales smaller than this, in contrast to the CDM case, where yet smaller dark matter haloes are expected to form in abundance. Another possibility is that the dark matter is significantly *self-interacting* (SIDM; e.g. Spergel & Steinhardt 2000), which could smooth out density peaks via scattering interactions.

The core motivation for the existence of cosmological dark matter is based on the need for a dominant, non-baryonic, dark mass component to reconcile the observed dynamics and structure of electromagnetically luminous material with the gravitational forces inferred on galactic and cosmic scales. However, if the dark matter is indeed a fundamental particle, as expected, it may be possible to detect its presence through techniques that are independent of gravitational arguments. Direct detections could be achieved on Earth from weak interactions with WIMPs, or conversion of axions into microwave photons within a magnetic resonant cavity (e.g. Asztalos et al. 2010). Alternatively, radiation from the decay or annihilation of distant dark matter particles may be detectable, particularly from high density regions, such as halo centres. However, currently there is no compelling evidence for the (non-gravitational) detection of dark matter, and its true identity remains a mystery. It is worth noting that the dark matter could be a mixture of several different non-baryonic particle species (e.g. neutrinos certainly *do exist*, and it seems that they are not quite massless, so they could contribute a *small* HDM component).

In the late 1990s, the High-z Supernova Search Team (Riess et al., 1998), and the Supernovae Cosmology Project (Perlmutter et al., 1999), announced discoveries that observations of the brightnesses of distant Type Ia supernovae (SNIa), out to a redshift of $z \sim 1$, indicate that not only is the Universe expanding, but the rate of expansion is currently *accelerating*. This suggests that the density of the Universe is currently dominated by a mass (energy) component with *negative pressure*. This mysterious density component is referred to as *dark energy*.

The simplest theoretical description of dark energy is Einstein's famous

cosmological constant, Λ (with units of inverse length squared); originally proposed as a means of reconciling a static Universe with general relativity. A positive value of Λ opposes gravity, acting to accelerate the rate of cosmic expansion. Considering a sufficiently large, positive, cosmological constant allows us to account for the current acceleration observed in the expansion of space, but admits another cosmological unknown: there is no theoretical consensus on the physical identity of the dark energy, and it may be that a perfect cosmological constant turns out to be an inadequate description of this phenomenon. Given its negative pressure, work is done on the cosmological constant fluid as space expands, allowing the energy density to remain fixed. In physical terms, the cosmological constant is often thought of as relating to the constant vacuum (zero-point) energy density of empty space. However, the vacuum energy density predicted on particle physics grounds is many orders of magnitude higher than the energy density due to Λ expected from cosmological arguments (this is known as the *cosmological constant problem*; e.g. Coles & Lucchin 2002).

Discovered by Penzias & Wilson (1965) and Dicke et al. (1965), the *cosmic microwave background* (CMB) is a relic of the early Universe, and provides a powerful observational window into the physical state of the Universe at the time that the constituent photons of the CMB were emitted. CMB photons are detected from all directions on the sky, and exhibit a black body spectrum with a temperature of approximately 3 K that is almost perfectly uniform in all directions (once the motion of the Earth, and emission from other sources are corrected for).

Shortly before the CMB was produced, there was thermal equilibrium between photons and baryons throughout the Universe; there were enough photons present with sufficient energy to immediately ionise any hydrogen atoms that formed from free protons and electrons in the primordial plasma. As the Universe expanded, cosmological redshifting reduced the number of high energy photons severely enough to allow hydrogen atoms to survive. This epoch is known (rather misleadingly) as *recombination*. Following recombination, the Universe became transparent to radiation, as the photon energies were low enough to prevent interaction with the newly formed hydrogen atoms. It is photons from this *decoupling* of radiation from baryonic matter that we observe today as the CMB. Thus the observed CMB photons

reveal the *surface of last scattering* of electromagnetic radiation.

Since the CMB was produced following a state of thermal equilibrium with baryonic matter, angular variations (*anisotropies*) in the observed temperature of the CMB reflect variations in the matter density at the time of CMB decoupling. Such anisotropies are indeed detected, at a level of around 10^{-5} relative to the mean temperature averaged over the whole sky. The angular power spectrum of CMB temperature anisotropies provides extensive insight into the state of the matter density field at the time of CMB decoupling: the temperature fluctuation amplitudes are proportional to the matter density perturbation amplitudes at decoupling (this proportionality is preserved at later times as long as the density perturbations, $\Delta\rho$, are still *linear*, i.e. $|\Delta\rho|/\rho \ll 1$). Before decoupling, the pressure in the photon-baryon fluid due to interactions between electrons and photons competed with the gravitational forces, driving waves referred to as *baryon acoustic oscillations* (BAOs). The characteristic BAO length-scale is an important standard ruler, not only when applied to the CMB radiation, but also in the clustering of galaxies (which formed from the matter distribution that underwent the oscillations prior to CMB decoupling). See Mo et al. (2010) for a detailed discussion of the physical origins of different features observed in the CMB radiation field.

Theoretical predictions for the CMB angular power spectrum can be generated using sophisticated programs such as CMBFAST (Seljak & Zaldarriaga, 1996) and CAMB (Lewis & Challinor, 2011), for a given set of cosmological parameters. In this way, modelling of the temperature anisotropies of the CMB can provide powerful constraints on cosmological models. Analysis of the CMB power spectrum indicates that the Universe is at least very close to spatially flat, if not perfectly flat. It is perhaps surprising that the spatial geometry should happen to (almost) possess the fortuitously simple special case of Euclidean flatness. The apparent happenstance of spatial flatness is known as the *flatness problem* of the Hot Big Bang model. Additionally, it is not immediately clear how the CMB temperature manages to be so close to being perfectly isotropic, since only small regions of the sky could ever have been in contact and reached thermal equilibrium before CMB decoupling (the so-called *horizon problem*). These apparent issues with the standard cosmological model are resolved if the Universe underwent a period of *inflation*

following the Big Bang (Guth, 1981; Albrecht & Steinhardt, 1982; Guth & Pi, 1982; Hawking, 1982; Linde, 1982a,b; Starobinsky, 1982; Linde, 1983; Linde et al., 1994; Lidsey et al., 1997; Lyth & Riotto, 1999; Guth, 2000). Inflation also resolves theoretical issues with the apparently sparse abundances of hypothetical *relic particles*, such as magnetic monopoles.

Perhaps the most compelling aspect of inflationary theory is that it provides a mechanism to generate the perturbations in the density field that are revealed in the CMB; perturbations which were the seeds of the formation of cosmic structure through gravitational instability. The quantum mechanical *uncertainty principle* (in particular the energy-time uncertainty) allows for the creation and annihilation of particles in otherwise empty space on extremely short time-scales. Inflation caught hold of these quantum fluctuations, and amplified them to scales far beyond the reach of the quantum mechanical realm; imprinting small irregularities throughout the cosmic density field, which remained at the end of the inflationary period.

We now have a theoretical model in which quantum fluctuations following the Big Bang give rise to all observed structures in the present Universe. Within this framework, the canonical model that has risen to prominence over the past few decades is *Λ cold dark matter* (ΛCDM). In ΛCDM, the spatial geometry is perfectly flat, the energy density of the Universe is currently dominated by dark energy in the form of a perfect cosmological constant, and the matter density is dominated by non-baryonic cold dark matter. Six parameters are required to specify the ΛCDM model:[3] H_0, $\Omega_{\mathrm{m},0}$, $\Omega_{\mathrm{b},0}$, n_{s}, σ_8, and τ (see Table 1.1 for their definitions); although the parameters may be constrained and quoted in alternative forms, e.g. $\Omega_{\mathrm{m},0}h^2$. Other important quantities can be derived from this set of six parameters, e.g. $\Omega_{\Lambda,0} = 1 - \Omega_{\mathrm{m},0}$ and $\Omega_{\mathrm{dm},0} = \Omega_{\mathrm{m},0} - \Omega_{\mathrm{b},0}$. The current best constraints on the ΛCDM parameters come from the Planck mission, combining CMB temperature anisotropy data with CMB polarisation and gravitational lensing data, along with independent constraints from SNIa luminosities, measurements of the characteristic BAO length-scale as imprinted on the observed clustering of galaxies, and the value of H_0. Table 1.1 lists the values of the

[3]Note that τ is much less fundamental than the other five ΛCDM model parameters. However, since a theoretical prediction of τ would require detailed and precise knowledge of the physics of galaxy formation, τ is generally treated as a free parameter.

Table 1.1 – Various important cosmological parameters, derived assuming a spatially flat ΛCDM cosmology. h is defined by $H_0 = 100h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, where H_0 is the Hubble parameter at the present time. $\Omega_{\mathrm{b},0}$, $\Omega_{\mathrm{dm},0}$, $\Omega_{\mathrm{m},0}$, and $\Omega_{\Lambda,0}$ are the current density parameters for baryons, dark matter, all non-relativistic matter, and the cosmological constant, respectively. n_s is the spectral index of primordial fluctuations, and σ_8 normalises the linear power spectrum at redshift zero. τ is the optical depth to reionisation, and z_re is the redshift of reionisation. t_0 is the age of the Universe, expressed in billions of years (Gyr). The first column gives the parameter values from WMAP7 (taken from the final column of table 1 of Komatsu et al. 2011), and the second column gives the values from Planck2015 (taken from the final column of table 4 of Planck Collaboration 2016). In each case, CMB data is combined with additional constraints when determining the values of the cosmological parameters.

Parameter	WMAP7	Planck2015
h	0.702 ± 0.014	0.6774 ± 0.0046
$\Omega_{\mathrm{b},0}h^2$	0.02255 ± 0.00054	0.02230 ± 0.00014
$\Omega_{\mathrm{dm},0}h^2$	0.1126 ± 0.0036	0.1188 ± 0.0010
$\Omega_{\mathrm{m},0}h^2$	0.1352 ± 0.0036	0.14170 ± 0.00097
$\Omega_{\Lambda,0}$	0.725 ± 0.016	0.6911 ± 0.0062
n_s	0.968 ± 0.012	0.9667 ± 0.0040
σ_8	0.816 ± 0.024	0.8159 ± 0.0086
τ	0.088 ± 0.014	0.066 ± 0.012
z_re	10.6 ± 1.2	$8.8^{+1.2}_{-1.1}$
t_0 [Gyr]	13.76 ± 0.11	13.799 ± 0.021

ΛCDM parameters derived using the full Planck data (Planck2015), and the slightly older WMAP seven-year data (WMAP7). The preferred values of the cosmological parameters from Planck2015 and WMAP7 are very similar. The Planck analysis finds excellent agreement between the predictions of ΛCDM and many relevant observational constraints (Planck Collaboration, 2016). The simulations used in Chapters 2 and 3 assume a ΛCDM cosmology with parameters consistent with WMAP7.

For reviews of the standard cosmological model see, for example, Liddle (2003), Coles & Lucchin (2002), Mo et al. (2010), and Frenk & White (2012).

1.2 Simulating galaxy formation

1.2.1 The matter power spectrum

Inflationary models predict that the Universe was seeded with initial density fluctuations that follow Gaussian statistics. The temperature anisotropies observed in the CMB are also found to be consistent with Gaussian statistics. If the primordial density perturbations in the cosmic matter distribution are indeed a homogeneous and isotropic *Gaussian random field*, then all the relevant statistical information can be described using only the *power spectrum*, $P(k)$, for wavenumber k.

Let the density at position $\boldsymbol{x}$ be $\rho(\boldsymbol{x})$. We can write

$$\rho(\boldsymbol{x}) = \overline{\rho}\,[1 + \delta(\boldsymbol{x})]\ , \tag{1.1}$$

where $\overline{\rho}$ is the mean density at the time of interest, and $\delta(\boldsymbol{x})$ is the *overdensity*. Thus defined, $\delta(\boldsymbol{x})$ contains all information on the density fluctuations, throughout all space. It is convenient to work with the Fourier transform of the overdensity field, rather than with $\delta(\boldsymbol{x})$ itself. Expanding the perturbation amplitude at any point as a sum of plane waves,

$$\delta(\boldsymbol{x}) = \sum_{\boldsymbol{k}} \hat{\delta}(\boldsymbol{k}) \exp(i\boldsymbol{k}\cdot\boldsymbol{x})\ , \tag{1.2}$$

where $\boldsymbol{k}$ is the wavevector, with magnitude $k = |\boldsymbol{k}|$. The (complex) amplitudes of the Fourier modes are then given by

$$\hat{\delta}(\boldsymbol{k}) = \frac{1}{L^3}\int_V \delta(\boldsymbol{x}) \exp(-i\boldsymbol{k}\cdot\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}\ , \tag{1.3}$$

where V is a finite cubic volume of space, with side length L, sufficiently large that the region may be approximated as periodic in all directions, and much larger than the scale of any manifest density structure (e.g. Coles & Lucchin 2002). The requirement of periodicity sets a discretisation condition for the allowed wavevectors:

$$\boldsymbol{k} = \frac{2\pi}{L}(i_x, i_y, i_z)\ , \tag{1.4}$$

for integers i_x, i_y, and i_z (since $k = 2\pi/\lambda$ for wavelength λ).

The *power spectrum* of the density perturbations is defined as

$$P(k) = L^3 \langle |\hat{\delta}(\boldsymbol{k})|^2 \rangle \,. \tag{1.5}$$

Note that $P(k)$ depends only on the *magnitude* of $\boldsymbol{k}$, otherwise our assumption of isotropy would be violated (cosmological principle). $P(k)$ is proportional to the variance of the Fourier mode amplitudes for wavenumber k, and has dimensions of volume. For a Gaussian random field that is homogeneous and isotropic, $P(k)$ provides a complete statistical description of the density perturbations.

If $\hat{\delta}(\boldsymbol{k})$ is a Gaussian random field, then its real, $\hat{\delta}_{\rm R}(\boldsymbol{k})$, and imaginary, $\hat{\delta}_{\rm I}(\boldsymbol{k})$, parts are each individually Gaussian distributed around zero (with equal dispersion), and thus $\delta(\boldsymbol{x})$ and $\rho(\boldsymbol{x})$ must also be Gaussian random fields (see Coles & Lucchin 2002 for more details). In order for $\delta(\boldsymbol{x})$ to be real, we require $\hat{\delta}(\boldsymbol{k}) = \hat{\delta}^*(-\boldsymbol{k})$.[4] In the linear regime ($|\delta(\boldsymbol{x})| \ll 1$), each Fourier mode grows independently, i.e. $\hat{\delta}(\boldsymbol{k})$ evolves separately with time for each $\boldsymbol{k}$. However, once the density perturbations become non-linear, the assumption of Gaussianity no longer holds, and the different Fourier modes can no longer be considered to evolve independently of each other. Since $\hat{\delta}(\boldsymbol{k})$ is complex, we can parameterise it as having an amplitude, $|\hat{\delta}(\boldsymbol{k})|$, and a phase angle, $\hat{\theta}(\boldsymbol{k})$, in the complex plane, such that

$$\hat{\delta}(\boldsymbol{k}) = \hat{\delta}_{\rm R}(\boldsymbol{k}) + i\hat{\delta}_{\rm I}(\boldsymbol{k}) = |\hat{\delta}(\boldsymbol{k})| \exp(i\hat{\theta}(\boldsymbol{k})) \,, \tag{1.6}$$

where $\hat{\delta}_{\rm R}(\boldsymbol{k})$, $\hat{\delta}_{\rm I}(\boldsymbol{k})$, $|\hat{\delta}(\boldsymbol{k})| = [\hat{\delta}^2_{\rm R}(\boldsymbol{k}) + \hat{\delta}^2_{\rm I}(\boldsymbol{k})]^{1/2}$, and $\hat{\theta}(\boldsymbol{k}) = \arctan 2[\hat{\delta}_{\rm I}(\boldsymbol{k}), \hat{\delta}_{\rm R}(\boldsymbol{k})]$ are all real quantities. Since $\hat{\delta}_{\rm R}(\boldsymbol{k})$ and $\hat{\delta}_{\rm I}(\boldsymbol{k})$ are each Gaussian distributed, with the same dispersion, there is equal probability for any $\hat{\theta}$ in the range $[0, 2\pi)$. Using the amplitude and phase of $\hat{\delta}(\boldsymbol{k})$, equation (1.2) becomes

$$\delta(\boldsymbol{x}) = \sum_{\boldsymbol{k}} |\hat{\delta}(\boldsymbol{k})| \exp(i[\boldsymbol{k} \cdot \boldsymbol{x} + \hat{\theta}(\boldsymbol{k})]) \,. \tag{1.7}$$

Inflationary models typically predict that the *initial power spectrum*, $P_{\rm i}(k)$

[4] Thus $\hat{\delta}_{\rm R}(\boldsymbol{k}) = \hat{\delta}_{\rm R}(-\boldsymbol{k})$, and $\hat{\delta}_{\rm I}(\boldsymbol{k}) = -\hat{\delta}_{\rm I}(-\boldsymbol{k})$. The fact that this condition forces $\delta(\boldsymbol{x})$ to be real follows from equation (1.2) and the identity $\exp(i\alpha) = \cos(\alpha) + i\sin(\alpha)$.

(at some arbitrarily early time), has a simple power law form, such that

$$P_{\rm i}(k) \propto k^{n_{\rm s}} \,, \tag{1.8}$$

where $n_{\rm s}$ is the *spectral index of primordial fluctuations* (see Table 1.1). For 'late' times, i.e. after recombination, the *linear theory* power spectrum is

$$P(k,z) = D^2(k,z)T^2(k)P_{\rm i}(k) \,, \tag{1.9}$$

where $D(k,z)$ is the *linear growth factor*, and $T(k)$ is the *linear transfer function* (e.g. Coles & Lucchin 2002). $T(k)$ accounts for the detailed changes in the shape of the power spectrum between inflation and recombination (and the subsequent decoupling of the CMB). $D(k,z)$ encodes the evolution of the amplitude and shape of the power spectrum at later times, according to linear theory. The dependence of $D(k,z)$ on k is significant only on very small scales, such that on larger scales (lower k) the shape of $P(k,z)$ is essentially fixed at recombination (and thus only the power spectrum amplitude evolves with time on large scales), but the shape of $P(k,z)$ does indeed evolve with time on very small scales (i.e. for very high k). The (large-scale) linear growth factor can be computed given that the Hubble parameter is $H(z) = H_0 E(z)$, such that

$$D(z) \propto H(z) \int_z^\infty \frac{(1+z')}{E^3(z')} \, {\rm d}z' \,, \tag{1.10}$$

where the normalisation is arbitrary (e.g. Mo et al. 2010).[5]

The variance, $\sigma^2(R)$, of the density fluctuations within spheres of radius R can be found by integrating the power spectrum:

$$\sigma^2(R) = \frac{1}{2\pi^2} \int_0^\infty P(k) \widehat{W}_R^2(k) k^2 \, {\rm d}k \,, \tag{1.11}$$

where the function

$$\widehat{W}_R(k) = \frac{3}{(kR)^3} \left[\sin(kR) - kR\cos(kR) \right] \,, \tag{1.12}$$

is the Fourier transform of a spherical 'top hat' window function of radius

[5] Note that equation (1.10) is valid for the special dark energy cases of ΛCDM (i.e. $w = -1$), $w = -1/3$, or matter domination ($w = 0$), but is not fully general; where for pressure, p, and density, ρ, the dark energy equation of state is $p = w\rho c^2$.

R (e.g. Coles & Lucchin 2002). By convention, the normalisation of $P(k)$ is usually specified in terms of

$$\sigma_8 = \sigma(8\,h^{-1}\,\mathrm{Mpc})\,. \tag{1.13}$$

Thus, σ_8 is the root-mean-square density fluctuation amplitude within spheres of radius $8\,h^{-1}$ Mpc (see Table 1.1). For the purposes of power spectrum normalisation, σ_8 is evaluated at redshift zero, where the linear power spectrum is extrapolated from recombination according to linear theory (via equation 1.9).

It is worth noting that the power spectrum is the Fourier transform of the *two-point autocorrelation function*, $\xi(r)$, such that

$$P(k) = 4\pi \int_0^\infty \xi(r) \frac{\sin(kr)}{kr} r^2 \,\mathrm{d}r\,, \tag{1.14}$$

and in the reverse case

$$\xi(r) = \frac{1}{2\pi^2} \int_0^\infty P(k) \frac{\sin(kr)}{kr} k^2 \,\mathrm{d}k\,, \tag{1.15}$$

(e.g. Coles & Lucchin 2002). $\xi(r)$ gives the excess probability of finding two particles with separation r, with respect to the expectation for a uniformly random distribution. The statistical properties of a homogeneous and isotropic Gaussian random field are therefore fully specified if either its $P(k)$ or $\xi(r)$ are known, and either of these functions can be computed as the Fourier transform of the other.

Given an appropriate set of cosmological parameters, e.g. the relevant ΛCDM parameters listed in Table 1.1 (h, $\Omega_{\mathrm{b},0}$, $\Omega_{\mathrm{m},0}$, n_s, σ_8), the suitably normalised linear power spectrum at late times (or the transfer function) can be computed using software such as CAMB (Lewis & Challinor, 2011). The predicted $P(k)$ can then be used to construct initial conditions for cosmological simulations, such that density perturbations in the simulated matter distribution at high redshift are statistically consistent with those of the real Universe, as inferred through observations of the CMB. In this way, the initial conditions can be constructed while the density perturbations are comfortably within the linear regime, long before structure formation becomes non-linear.

1.2.2 Dark matter simulations

Initial conditions for cosmological simulations can be constructed by using the *Zel'dovich approximation* (Zel'dovich, 1970; Efstathiou et al., 1985), given the desired linear power spectrum of matter density perturbations, $P(k)$. For a cubic volume of fixed comoving side, L, we start by specifying the *unperturbed* positions of the N sampling particles that will trace the evolution of the cosmic matter distribution in six-dimensional position-velocity space (*phase space*), having mean density equal to that implied by the cosmological parameters, with periodic boundary conditions. In the simplest case, the particles all have the same mass, but this is not required. The unperturbed positions could form a regularly spaced Cartesian grid, or an irregular *glass* (random positions evolved with *repulsive* 'gravity' to reach a stable configuration; White 1996). The effective resolution is determined by the number of particles, or equivalently, the particle mass(es). The peculiar velocities of the unperturbed particles are set to zero.

Given the unperturbed particle positions, the gravitational potential field is perturbed by slightly displacing the particles, such that Gaussian fluctuations are introduced to the density field, consistent with the specified $P(k)$. In the Zel'dovich approximation, the particle displacements increase with time in proportion to the linear growth factor, $D(z)$, such that the particles move along fixed straight-line trajectories, where the directions are set by the initial displacements. The overdensity field, $\delta(\boldsymbol{x})$, grows in amplitude in the same way as the particle displacements (and so $P(k,z) \propto D^2(z)$ in the linear regime, as in equation 1.9). The initial particle velocities are derived from the gradient of the gravitational potential field (acceleration) at the time of the initial displacement (or derived from the displacements themselves, since velocity is proportional to displacement in the Zel'dovich approximation). The finite box size, L, sets a limit on the longest wavelength (smallest k) modes that can be represented ($k \sim 2\pi/L$), and on small scales the particle mass (particle spacing, d) implies a spatial Nyquist frequency of $k \sim 2\pi/d$, below which any fluctuation power cannot be resolved. The redshift chosen for the initial displacements (i.e. for the *initial conditions*) should be low enough to ensure that the displacements are much larger than any numerical noise, but high enough that the density perturbations are still linear, and

there are no shell crossings (typically we use $z = 127$, so the cosmological scale factor is 128 times smaller than at the present time). The Zel'dovich approximation corresponds to *first-order Lagrangian perturbation theory*. Including second order terms (e.g. Jenkins 2010) increases the accuracy of the initial conditions (and allows a lower initial redshift).

Once the initial conditions have been constructed, i.e. the masses, positions, and velocities of all particles have been specified, at some initial redshift, the evolution of the particle positions and velocities over time can be followed by repeatedly computing the gravitational forces and integrating over suitable time steps. In the case where all particles have the same mass, and only gravitational forces are considered, so the particles are *collisionless*, this is a homogeneously sampled *dark matter only* (DMO) cosmological simulation. DMO simulations are of course a class of *N-body simulation*, following the evolution of a system of N particles that are subject to physical (i.e. gravitational) forces. The baryons are still accounted for in DMO simulations (and the mean density is still correct), in that the baryons contribute to the masses of the collisionless 'dark matter' particles, but they are assumed to be negligibly interacting, except through gravitational forces.

To compute the gravitational forces acting on all particles throughout the simulation volume at any given time, the most accurate approach is to use *direct summation*, such that the force on some particle, j, due to another particle, i, is

$$\boldsymbol{F}_{ij} = -\frac{G m_i m_j}{|\boldsymbol{r}_{ij}|^3} \boldsymbol{r}_{ij} \,, \tag{1.16}$$

where G is the universal gravitational constant, $\boldsymbol{r}_{ij}$ is the separation vector from particle i to particle j, and m_n is the mass of particle n (n.b. $\boldsymbol{F}_{ji} = -\boldsymbol{F}_{ij}$). It is legitimate to use Newtonian gravity in cosmological simulations, since general relativity predicts the same linear structure growth as in the Newtonian approximation, and the relative velocities induced in the non-linear regime are much less than the speed of light (Chisari & Zaldarriaga, 2011; Frenk & White, 2012). In practice, direct summation via equation (1.16) becomes prohibitively slow for large particle numbers. More sophisticated techniques are required, in order to achieve significant speed improvements, at the cost of small (and tunable) force errors. For example, the forces can be computed using *Fourier methods* on a periodic grid that fills the simulation

volume (Hockney & Eastwood, 1988), or a geometric *tree* can be implemented, so that the forces due to distant particles can be approximated as the force due to the centre of mass of a tree cell that contains many particles (Barnes & Hut, 1986). The popular simulation code GADGET-2 combines these two approaches, implementing a tree on small scales, and a Fourier grid to compute large-scale force contributions (Springel, 2005).

In the case of Fourier techniques, the equation that is solved in order to evaluate the gravitational forces is not actually equation (1.16), but rather the Poisson equation,

$$\nabla^2\Phi = 4\pi G\rho \,, \tag{1.17}$$

where Φ is the gravitational potential, and ρ is the density. Given a regular mesh that fills the periodic volume, a density value is assigned to each grid cell by convolving the particle mass density field with a suitable weighting kernel function. The Fourier transform of the density grid is then the product of the Fourier transform of the kernel function (which is known) and the Fourier transform of the true density field ($\hat{\rho}$). The Fourier transform of equation (1.17) is then $-k^2\hat{\Phi} = 4\pi G\hat{\rho}$, where $\hat{\Phi}$ is the Fourier transform of Φ. Similarly, since the force per unit mass (gravitational acceleration) is $\boldsymbol{g} = -\boldsymbol{\nabla}\Phi$, we have the Fourier transform $\hat{\boldsymbol{g}} = -i\hat{\Phi}\boldsymbol{k}$. Thus $\hat{\rho}$ can be used to compute $\hat{\boldsymbol{g}}$ (via $\hat{\Phi}$), and inverse Fourier transforming, to compute $\boldsymbol{g}$ itself. Once $\boldsymbol{g}$ is computed in this way for all points on the grid, the gravitational forces acting on the individual particles are evaluated using interpolation. It should be noted that for simulation volumes with periodic boundary conditions, the interparticle forces are distorted for particle separations comparable to the periodic scale.

Since two-body scattering interactions between representative dark matter 'particles' (whose masses may be many orders of magnitude greater than that of the Sun) should not be considered to be physically meaningful, it is necessary to *soften* the gravitational force on small scales. Force softening also mitigates the need for the very small timesteps that such close interactions would require in order to maintain accurate time integration. In GADGET, this is achieved by using a *softening kernel* function that gradually suppresses the gravitational force on scales less than 2.8ϵ, where ϵ is the *Plummer-equivalent softening length*. The gravitational force remains Newtonian on scales larger than 2.8ϵ, in this case. The choice of gravitational softening length can have

important implications for structure growth, the dynamical realism of the simulation, and the convergence of simulation predictions with increasing mass resolution (e.g. Power et al. 2003).

Extensive insight into the formation and evolution of structure in a ΛCDM Universe can be gained through performing DMO simulations, since the gravitational clustering of dark matter dominates the growth of haloes and the larger-scale cosmic web. Simulations of large cosmologically representative volumes resolved at fixed particle mass, such as the Millennium Simulation (Springel et al., 2005), facilitate the study of large-scale structure and the statistical properties of the population of dark matter haloes, and can be used as the basis of semi-analytic models of galaxy formation (to be discussed in §1.2.3). It is also possible to select individual haloes, or larger regions of interest, from such homogeneously resolved DMO simulations and to *resimulate* the target object(s) at very high resolution (using particles of relatively low mass), while reproducing the original density field on larger scales at relatively low resolution (with higher mass particles). This *zoom* technique can be used to study the detailed substructure of haloes identified in the original (parent) simulation, revealing at high resolution the structure and kinematics of *subhaloes* embedded in the main halo, starting from cosmological initial conditions consistent with the Gaussian random field realisation and $P(k)$ of the parent simulation, as in the Aquarius simulation project of Springel et al. (2008).

A generic prediction of (Λ)CDM DMO simulations with cosmological initial conditions is that dark matter haloes have density profiles that are closely approximated by the *NFW profile*, as discovered by Navarro, Frenk & White (1996b, 1997):

$$\rho(r) = \frac{\rho_s}{(r/r_s)(1 + r/r_s)^2} \,, \tag{1.18}$$

where $\rho(r)$ is the spherically averaged density at radius r, and the parameters ρ_s and r_s are characteristic density and radius values, respectively. By convention, the mass of a halo is usually expressed in terms of the *virial mass*, M_{200}, which is the mass enclosed within one *virial radius*, r_{200}, of the halo centre. The virial radius is defined as that of the sphere that encloses a mean density equal to 200 times the critical density of the Universe (ρ_{crit}). The *concentration* of a halo is defined as r_{200}/r_s.

The NFW profile tends towards an inner density slope of $\rho \propto r^{-1}$ (known as a steep central density *cusp*), and a slope of $\rho \propto r^{-3}$ for large r. There has been a great deal of discussion as to whether or not observed galaxies really do have cusps in their inner dark matter density profiles, as predicted (*prima facie*) by DMO simulations set in a ΛCDM Universe, or whether the inner dark matter distributions of their host haloes actually tend to *cores* of constant density (e.g. Simon et al. 2003; Swaters et al. 2003; Simon et al. 2005; Koch et al. 2007; Valenzuela et al. 2007; de Blok et al. 2008; Oh et al. 2008, 2011; Walker & Peñarrubia 2011; Adams et al. 2012; Strigari et al. 2014). This *cusp-core problem* is often heralded as a symptom of some inherent failure of the ΛCDM model. However, confusion due to degeneracies in observational work may lead to identification of a dark matter core where there is in fact a cusp (e.g. Evans et al. 2009; Adams et al. 2012; Strigari et al. 2014), while hydrodynamical effects and energetic feedback may be able to produce central cores in haloes that would otherwise have cuspy inner density profiles, in a way that is entirely consistent with ΛCDM (e.g. Navarro et al. 1996a; Binney et al. 2001; Mashchenko et al. 2008; Governato et al. 2010; Pontzen & Governato 2012). Several additional well publicised small-scale 'problems' with ΛCDM seem to disappear once the important baryonic physics is taken into account (e.g. Moore et al. 1999; Klypin et al. 1999; Boylan-Kolchin et al. 2011; Brooks et al. 2013; Sawala et al. 2016; Zhu et al. 2016). The dark matter distributions of the satellite galaxies of the Milky Way have played a particularly prominent role in this debate.

1.2.3 Semi-analytic models

Given the particle data (masses, positions, velocities, etc.) output at various *snapshots* in time from a DMO simulation of cosmic structure formation in a ΛCDM Universe, we would like some way of relating this detailed prediction of the properties of the underlying dark matter structure of the Universe with the properties of the luminous galaxies that are expected to form in the halo centres, illuminating the dark matter density peaks. That is, we want to know how exactly the dark matter haloes are populated with galaxies. Only with a theoretical model that can make comprehensive predictions of observable galactic properties, based on the *ab initio* simulation of the

formation of structure within a ΛCDM Universe, can we begin to make fully robust comparisons to key observational constraints on the interwoven properties of the galactic population, and in doing so, draw meaningful conclusions on the validity of the ΛCDM paradigm and the precise details of the physical processes that govern the formation and evolution of galaxies.

A simple approach to relating simulated dark matter haloes to observable galaxies is to assume that the stellar mass of a galaxy is somehow precisely determined by the mass of its host halo. The *halo mass function* (number density of haloes as a function of mass) predicted by a DMO simulation can be compared to the *stellar mass function* (number density of galaxies as a function of stellar mass) as inferred from an observational survey. If the stellar mass is assumed to be an exact, monotonic, function of halo mass, then via *abundance matching* we can assign a stellar mass to each halo in the simulation volume, or a halo mass to any observed galaxy, thus defining the *stellar mass to halo mass relation*. In practice, more accurate results may be obtained by considering the maximum mass ever achieved by a halo, rather than its final mass (at $z = 0$), taking careful account of the evolution of satellite galaxies, and compensating for biases due to the scatter in observationally inferred stellar masses at fixed halo mass (e.g. Behroozi et al. 2010; Guo et al. 2010; Moster et al. 2010, 2013).

Another approach to populating haloes with galaxies is the *halo occupation distribution* (HOD) technique, which is based on specifying the probability distribution describing the numbers of different types of galaxies that reside in haloes of different masses (e.g. Neyman & Scott 1952; Peacock & Smith 2000; Berlind & Weinberg 2002; Cooray & Sheth 2002; Kravtsov et al. 2004; Zehavi et al. 2011). The HOD method is similar to abundance matching, when applied to a DMO simulation, but introduces some level of stochasticity. While abundance matching and HOD modelling can provide theoretical predictions for key statistics, such as the clustering (correlation function) of galaxies as a function of stellar mass in ΛCDM, they cannot follow the complex evolution history of individual galaxies in a realistic and self-consistent way, starting from the infall of primordial gas into the gravitational potential well of a protohalo. For a rigorous, self-consistent, method of predicting the properties of the galaxies that reside in the haloes produced in DMO cosmological simulations, we can turn to *semi-analytic*

galaxy formation models.

Building on the work of White & Rees (1978) and White & Frenk (1991), semi-analytic models use simple analytical prescriptions and approximations to implement the key, interlinked, physical processes that are involved in the formation and evolution of galaxies, including: gas heating and cooling, the formation of gas discs, star formation, chemical enrichment of gas and energy feedback from stellar evolution and supernovae, black hole formation and energy feedback from active galactic nuclei (AGN), galaxy mergers, and bursts of star formation due to instabilities in gas discs (e.g. Benson et al. 2000a; Cole et al. 2000; Hatton et al. 2003; Baugh et al. 2005; Bower et al. 2006; Croton et al. 2006; Monaco et al. 2007; Somerville et al. 2008).

The equations used to model the complex physical processes of galaxy formation in semi-analytic models are motivated by theoretical considerations, either from first principles or numerical simulations. In some cases, the equations may be fully specified given independent empirical constraints on the relevant parameters (e.g. Lagos et al. 2011). The remaining free parameters can be constrained by comparing the model predictions of statistics such as the galactic stellar mass function, or luminosity function in various wavelength bands, to corresponding observational measurements. Through calibrating the relevant parameters so that the model predictions are brought into agreement with the observational constraints, either 'by hand' or using sophisticated statistical methods (e.g. Henriques et al. 2013), the dependence of the model predictions on the various different parameters (and hence the relevant underlying physics) can be better understood, and predictions can be made for different statistics of the galactic population that were not used in the calibration process. In this way, comparing the model predictions to new observational results acts as a test of the predictive power of the model, and of the validity of its treatment of galaxy formation and evolution (and of the underlying cosmological framework). Where disagreements are found, this may indicate that parameter adjustments are required that the original calibration statistics were not sensitive to, or that there is some fundamental issue with the model, or that modelling of additional physical processes may be required to explain the observations. A properly calibrated semi-analytic model can be used to study the evolution history of the galactic population, potentially providing new insights into the physical processes of galaxy

formation and evolution. A large number of galaxy properties can be tracked over time within such models, including e.g. the stellar mass contributions from the galactic disc and bulge components, and the integrated luminosities (magnitudes) in different wavelength bands.

Typically, semi-analytic models exploit pre-existing DMO simulations of sufficiently high resolution and large enough volume to provide a cosmologically representative sample of well-resolved dark matter haloes, from which *merger trees* can be constructed, mapping the evolution of the halo population over time, through hierarchical clustering.[6] When a halo forms, it is assumed that gas that is drawn into the gravitational potential well is heated by shocks to form a spherically symmetric reservoir of hot halo gas. Over time this gas can cool, and condense in the halo centre, leading to star formation. By following the evolution of individual galaxies, and the merging of galaxies (when their host haloes merge), semi-analytic models can generate detailed, self-consistent, predictions for the properties of the galactic population over cosmic time.

1.2.4 Hydrodynamical simulations

Semi-analytic models are powerful tools for investigating the physical processes involved in the formation and evolution of galaxies, and for making theoretical predictions for comparison to observational results. One important advantage of semi-analytic models is their computational efficiency: given an underlying DMO simulation, the parameters of the semi-analytic model can be adjusted with relative ease, and new analytical recipes can be quickly tested, since the model calculations can be repeated many times without incurring prohibitive computational cost (without requiring the use of smaller DMO simulation volumes, or lower resolution, for calibration purposes). On the other hand, semi-analytic models describe a substantially simplified version of the true physical situation. The assumptions involved in the models must be violated in practice, at least at some level, and such models neglect the possibility that the behaviour of the dark matter may

[6] An alternative to basing semi-analytic models on the outputs of DMO simulations is to create *Monte Carlo* (i.e. random) realisations of halo merger histories (based on theoretical expectations for the distribution of halo progenitor masses), which can then be processed by semi-analytic models (e.g. Lacey & Cole 1993; Cole et al. 2000; Baugh 2006).

be influenced by baryonic physics. Similarly, while semi-analytic models can predict many integrated properties for very large numbers of galaxies, and perhaps e.g. the sizes of galaxy discs and the mass contributions from different structural components, they cannot resolve the detailed internal structure and kinematics of the galaxies.

As an alternative to the semi-analytical approach, we can simulate the important baryonic physics directly, alongside the gravitating collisionless dark matter particles, in *cosmological hydrodynamical simulations*. Taking ΛCDM initial conditions designed for DMO N-body simulations as described in §1.2.2, we can split each of the original collisionless 'dark matter' particles by mass, according to the cosmic baryon fraction ($\Omega_\mathrm{b}/\Omega_\mathrm{m} \approx 0.16$), to spawn two new particles: a collisionless dark matter particle (as before), and a gas particle (or some other fluid representation of the gas). The gravitational forces are computed in the same way as in a DMO simulation, but, in addition, we follow the hydrodynamical evolution of the thermodynamic properties of the baryonic gas over cosmic time.

The relevant equations for the time evolution of the gas are those of *ideal hydrodynamics*.[7] In fluid dynamics, we can adopt either the *Lagrangian* or *Eulerian* formalism. In the Lagrangian case, the frame of reference follows the fluid flow, and the fluid properties and equations are expressed for a given *fluid element*. In the Eulerian case, the observer is instead fixed in space, and the fluid equations relate to the fluid properties at some particular location. While these two formalisms are physically equivalent, their practical implementations are very different, in the context of numerical simulations. Lagrangian simulations follow the properties of particles (discretising mass), while Eulerian simulations track the fluid quantities within static grid cells (discretising space). Adopting the Lagrangian frame, the relevant fluid equations are: the *continuity equation* (conservation of mass),

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} + \rho \boldsymbol{\nabla} \cdot \boldsymbol{u} = 0 \,, \tag{1.19}$$

the *Euler equation* (momentum equation),

$$\frac{\mathrm{d}\boldsymbol{u}}{\mathrm{d}t} = -\frac{\boldsymbol{\nabla} p}{\rho} - \boldsymbol{\nabla}\Phi \,, \tag{1.20}$$

[7] Or even *magnetohydrodynamics* (e.g. Pakmor et al. 2011; Hopkins & Raives 2016).

the *energy equation* (conservation of energy),

$$\frac{\mathrm{d}\mathcal{E}}{\mathrm{d}t} = -\frac{p}{\rho}\boldsymbol{\nabla}\cdot\boldsymbol{u} - \frac{\mathcal{L}}{\rho}\,, \tag{1.21}$$

and the *Poisson equation* (local density and gravitational field),

$$\nabla^2\Phi = 4\pi G\rho\,, \tag{1.22}$$

where the *equation of state* is

$$p = (\gamma - 1)\rho\mathcal{E}\,, \tag{1.23}$$

where ρ is the density, $\boldsymbol{u}$ is the velocity vector, p is the pressure, $\mathcal{E}$ is the internal energy per unit mass, $\mathcal{L}$ is the net cooling rate per unit volume, γ is the adiabatic index (ratio of specific heats[8]), *of the fluid element of interest*, while Φ is the gravitational potential *at the location of the fluid element*, and t is time (e.g. Mo et al. 2010). Note that in a cosmological context there are likely to be significant contributions to the gravitational potential from density components other than the baryonic gas, and in this case the density that enters into equation (1.22) should be the total density at the position of the fluid element, including e.g. the *local* dark matter density.

The details of the hydrodynamics implementation can vary significantly between different cosmological simulation codes, primarily as a result of the differences between the Lagrangian and Eulerian approaches (e.g. Frenk et al. 1999; Fryxell et al. 2000; Kravtsov et al. 2002; Teyssier 2002; O'Shea et al. 2004; Wadsley et al. 2004; Springel 2005; Stone et al. 2008; Springel 2010; Hopkins 2015). The fluid equations can be solved for discrete gas particles of fixed mass,[9] where the local gas properties are obtained by smoothing over some number of neighbouring particles (Lagrangian *smoothed particle hydrodynamics*; SPH; Lucy 1977; Gingold & Monaghan 1977; Monaghan 1992), or for static Cartesian grid cells that discretise space (e.g. Stone & Norman 1992) and may be progressively refined in order to increase the spatial resolution

[8] That is, $\gamma = C_p/C_V$, where C_p is the heat capacity at constant pressure, and C_V is the heat capacity at constant volume. For a monotonic gas $\gamma = 5/3$. In general, for an ideal gas, $\gamma = (q+5)/(q+3)$, where q is the number of internal degrees of freedom. For reversible processes that conserve entropy $p \propto \rho^\gamma$, while $p \propto \rho$ for an isothermal processes.

[9] Ignoring possible chemical enrichment of the gas by neighbouring stars.

in regions of particular interest (Eulerian *adaptive mesh refinement*; AMR; Berger & Colella 1989). In SPH, the spatial resolution is naturally highest in high density regions, coinciding with structures of primary interest such as galaxies (since there are more particles there), and the gas particles clearly follow the fluid flow (minimising advection errors); while in AMR the spatial resolution can be varied according to some prescribed scheme. Potential disadvantages of AMR methods include advection errors, a lack of Galilean invariance, issues with angular momentum conservation, and excessive fluid mixing, while in SPH, there is usually no mixing of physical properties at all between gas particles, by construction. SPH codes usually create *artificial viscosity* in order to suppress unphysical interactions between gas particles and to generate entropy during shocks, and so SPH methods tend to have difficultly in resolving shocks and fluid instabilities (since discontinuities get smoothed out), while grid-based Eulerian techniques tend to be better suited to solving these problems. One alternative approach is to solve the hydrodynamical equations using an unstructured mesh, with mesh cells that follow the fluid flow in a Lagrangian fashion, so that fluid motions and shocks are captured at high accuracy (e.g. Springel 2010). It is even possible to construct a finite-volume scheme without using a mesh (Hopkins, 2015).

Many important physical processes are too complicated, or operate on too small a scale, to be simulated directly, within large cosmologically representative volumes. In this regime, hydrodynamical simulations implement so-called *subgrid physics models* to approximate the relevant processes that cannot be resolved (e.g. Vogelsberger et al. 2013; Crain et al. 2015). This can include, for example, prescriptions for star formation (producing massive *star particles* that represent stellar populations), metal enrichment,[10] black hole formation, and energy feedback into the gas from stars and AGN. The subgrid physics models resemble the expressions used in semi-analytic models, although in the hydrodynamical case we can of course invoke the properties of individual baryonic particles (or gas cells) within the subgrid equations. The subgrid physics equations often have free parameters that require calibration of the simulation predictions against particular observational constraints, in a similar way to semi-analytic models. The semi-analytical and hydrodynamical approaches can be seen as complementary,

[10] In astronomy, all chemical elements apart from hydrogen and helium are *metals*!

since semi-analytic models allow rapid testing of new modelling choices, and can make predictions for the properties of galaxies within cosmological volumes of sizes much larger than currently computationally feasible with full hydrodynamics (at fixed resolution); while hydrodynamical simulations offer a highly accurate treatment of the gas dynamics, can resolve the detailed structure and kinematics of galaxies and their environment at high resolution, and dispense with many of the assumptions and approximations involved in semi-analytic modelling. Variations in subgrid modelling choices tend to vastly outweigh the influence of differences due to the choice of hydrodynamics scheme, when the predictions of different cosmological hydrodynamical simulation codes are compared, for exactly the same initial conditions (Scannapieco et al., 2012).

Cutting-edge simulation codes featuring sophisticated treatments of the hydrodynamical evolution of the gas and comprehensive subgrid physics models have now been used to produce simulations of the formation of realistic galaxies within a ΛCDM Universe, starting from Gaussian initial conditions; both in homogeneously resolved volumes to create large samples of galaxies (e.g. Vogelsberger et al. 2014; Schaye et al. 2015), and also in very high resolution zoom simulations of individual systems, such as the Milky Way and the wider Local Group (e.g. Sawala et al. 2016; Grand et al. 2017).

1.3 Estimating the masses of galaxies

1.3.1 Dynamical masses

Of the fundamental forces of nature, only gravity is able to operate on large cosmological scales. The observed motions of planets, stars, and galaxies can be understood through considering the gravitational forces to which they are subjected, using the extremely simple force law of Newtonian gravity (i.e. equation 1.16), or general relativity, in the most demanding cases. The mathematical foundation of the standard cosmological framework is rooted in the solutions of the Einstein field equations of general relativity, for a homogeneous and isotropic Universe. It was through studying the gravitational dynamics of astrophysical systems that the mysterious cosmological dark matter came to be first exposed, paving the way for our modern theoretical

view of galaxy formation, driven by the cooling and ignition of baryonic gas captured in the gravitational potential wells of dark matter haloes.

Since dark matter dominates over the baryonic contribution to the cosmic energy density, the history of the highly non-linear growth of cosmic structure, from the decoupling of the CMB until the present time, can be accurately understood and explored, to first order, by following the hierarchical gravitational clustering of collisionless dark matter particles, while considering no other forces but gravity. Of course, a more sophisticated theoretical approach is required in order to model and elucidate the finer details of the hydrodynamical evolution of the baryonic gas, and the formation and evolution of stars and galaxies, including the possible influences of baryons on the dark matter density and velocity fields. Nonetheless, dark matter haloes can be generally regarded as the primary structure of galaxies, exerting their gravitational dominance over all other material. This view is particularly appropriate for dwarf galaxies, which appear to have very high dark matter to baryonic mass ratios, so the baryonic contribution to the gravitational field is fairly insignificant.

A vast wealth of information is encoded in the total dynamical masses of galaxies, and in the structure of their mass distributions. There is a great deal of interest in the halo masses and abundances of the dwarf satellite galaxies of the Milky Way, and whether or not these can be accounted for within a ΛCDM context (e.g. Boylan-Kolchin et al. 2011; Sawala et al. 2016). Answering the question of whether galactic dark matter haloes have constant density cores or NFW-like cusps in their inner density profiles also reduces to determining (an integral of) the distribution of dynamical mass (e.g. Walker & Peñarrubia 2011; Adams et al. 2012).

Various methods exist for inferring the total dynamical masses of observed galaxies, or their spherically averaged density profiles, or even their full decomposed (3D) density and velocity fields. *Schwarzschild modelling* (Schwarzschild, 1979, 1993) can be used to construct a large (ideally, exhaustive) library of random realisations of orbital models for individual stars within stellar systems, in equilibrium, for a given set of allowed gravitational potentials (e.g. Richstone 1980; Richstone & Tremaine 1988; Rix et al. 1997; van der Marel et al. 1998; Cretton & van den Bosch 1999; Gebhardt et al. 2003; Cappellari et al. 2006; Magorrian 2006). A particular (weighted) set of

orbits can then be used to compute a numerical representation of the stellar *distribution function*, $f(\boldsymbol{x}, \boldsymbol{v})$, for that model, which is the probability density distribution for finding any single star with a given position, $\boldsymbol{x}$, and velocity, $\boldsymbol{v}$, in six-dimensional *phase space*. This is a powerful and flexible technique, since the distribution function can be used to predict the projected density and velocity fields, allowing the model to interface with very high resolution observational data, in projection (typically 2D maps of the projected stellar density and velocity dispersion). Some fitting procedure can be carried out to identify the (set of) orbital model(s) that best explains the available observational data. The preferred model can then provide information on the most likely gravitational potential, and hence the mass distribution of the system. The Schwarzschild method is a highly sophisticated approach, and in general permits essentially arbitrary geometry (e.g. spherically symmetric, axisymmetric, or triaxial), does not restrict the possible configuration of the orbits, can include mass splitting between different density components (so e.g. density variations in the dark matter, gas, and black holes can be included in the set of gravitational potentials within the library), and can generate predictions for any inclination angle. However, the quality of the observational data available for galaxies does not always warrant such a thorough (and computationally involved) theoretical treatment. Simpler methods that rely on additional assumptions are usually preferred in place of full Schwarzschild modelling of the superposition of stellar orbits. As a complementary approach to Schwarzschild modelling, it is possible to construct a 'made to measure' N-body simulation that reproduces the available observational data (Syer & Tremaine, 1996; de Lorenzi et al., 2007).

At a simpler level, the measured *rotation curve*, $V_{\rm rot}(R)$, of an observed disc galaxy can be used to infer the mass distribution of the dark matter halo in which the galaxy is embedded. Here, R is the galactocentric radius in the plane of the disc, and $V_{\rm rot}(R)$ has been corrected for the inclination of the galaxy with respect to the observer. If we assume that the measured rotation curve is that of stars on perfect circular orbits in the gravitational potential of the system, $V_{\rm rot}(R)$ can be assumed equal to the *circular velocity curve*, $V_{\rm circ}(R)$, where for a spherical system

$$V_{\rm circ}^2(r) = r\frac{{\rm d}\Phi(r)}{{\rm d}r} = \frac{GM(<r)}{r}\,, \tag{1.24}$$

where $\Phi(r)$ is the gravitational potential at 3D radius r, and $M(<r)$ is the total dynamical mass enclosed within a sphere of radius r. Splitting $M(<r)$ into the contributions from different structural components, we can write

$$V_{\rm circ}^2(r) = V_{\rm halo}^2(r) + V_{\rm disc}^2(r) + V_{\rm bulge}^2(r) + \dots \,, \tag{1.25}$$

where $V_{\rm halo}^2(r) = GM_{\rm halo}(<r)/r$ is the contribution due to the dark matter halo. Defined analogously, $V_{\rm disc}^2(r)$ and $V_{\rm bulge}^2(r)$ are the contributions due to the mass in the stellar disc and bulge (if relevant), respectively. Any number of additional terms could be included, e.g. to account for a gas disc, or a central supermassive black hole. Given the measured rotation curve, $V_{\rm rot}(R) \approx V_{\rm circ}(R)$, which could be derived from the observed kinematics of a gas disc, rather than the stellar disc, equation (1.25) provides a method to constrain the mass profile of the dark matter halo, $M_{\rm halo}(<r)$, if mass models for the important baryonic components can be specified. If the mass-to-light ratio of the stars, Υ, is poorly constrained, this can lead to an annoying degeneracy between the halo and stellar mass profiles (e.g. Dutton et al. 2005). The extreme cases of $\Upsilon = 0$ (massless stars) and the maximum Υ allowed by the dynamical constrains (*maximal disc model*) can provide upper and lower limits on the possible mass contribution from the dark matter, respectively. The value of Υ (which need not be constant throughout the galaxy) can be treated as a free parameter in the mass modelling procedure. However, in order to break the *disc-halo degeneracy* it is highly desirable to have some independent constraint on Υ (i.e. a constraint on the stellar mass; e.g. Adams et al. 2012). Techniques for estimating stellar masses will be discussed in §1.3.2. If Υ is not well constrained, the range of acceptable values for Υ in the mass modelling may span a broad range of inner density profile slopes for the dark matter halo. Issues with Υ are clearly avoided if the dark matter can be assumed to dominate the gravitational potential to such an extent that equation (1.25) can be assumed to reduce to $V_{\rm circ}^2(r) = V_{\rm halo}^2(r)$. Low surface brightness galaxies and dwarf galaxies are heavily dark matter dominated, and so they are particularly attractive for mass modelling studies. Besides the disc-halo degeneracy, there are various other important issues that may need to be accounted for when trying to recover the dark matter mass profile via this technique, such as flaws in the measured rotation curve,

the presence of a central stellar bar or the shape of the halo being triaxial rather than spherical (these two effects each complicate the relationship between $V_{\rm rot}(R)$ and $V_{\rm circ}(r)$), other complex non-rotational motions in the velocity fields, or the lack of observational data at large scales needed to constrain the overall shape of the halo density profile (e.g. Dutton et al. 2005; Mo et al. 2010; Adams et al. 2012).

Mass modelling using rotation curves can be effective for inferring the dark matter distributions of the haloes of disc galaxies. Yet if significantly more detailed kinematic information is available than just the mean rotation velocity as a function of radius then this method may be throwing away highly useful data that might otherwise offer significant additional constraining power. Also, this approach is clearly not appropriate for dispersion-supported galaxies, which are characterised by their lack of significant rotation. In order to connect with higher moments of the velocity distributions, we can use *Jeans modelling* of the observed stellar kinematics (Jeans, 1922).

The *collisionless Boltzmann equation* expresses the *convective* constancy of the distribution function, $f(\boldsymbol{x},\boldsymbol{v},t)$, over time, t, for a collisionless dynamical system that has a fixed number, N, of identical tracer particles (stars) of mass m, with a smooth gravitational potential given by $\Phi(\boldsymbol{x},t)$:

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \frac{\partial f}{\partial t} + \sum_{i=1}^{3} v_i \frac{\partial f}{\partial x_i} - \sum_{i=1}^{3} \frac{\partial \Phi}{\partial x_i}\frac{\partial f}{\partial v_i} = 0\,, \tag{1.26}$$

where x_i and v_i are the i-th components of $\boldsymbol{x}$ and $\boldsymbol{v}$, respectively, and the summations are over all three spatial dimensions. See the canonical book by Binney & Tremaine (2008) for the derivation of the collisionless Boltzmann equation and of the related equations presented below. We can write equation (1.26) in a more compact form using vector notation, so that

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \frac{\partial f}{\partial t} + \boldsymbol{v}\cdot\boldsymbol{\nabla} f - \boldsymbol{\nabla}\Phi\cdot\frac{\partial f}{\partial \boldsymbol{v}} = 0\,. \tag{1.27}$$

The quantity $\mathrm{d}f/\mathrm{d}t$ is the *Lagrangian* (or *convective*) derivative of f, in the sense that it refers to the (zero) change in the local probability density according to a star as it moves through phase space. This does *not* imply that f is constant with time everywhere in phase space, but it appears to be

constant according to individual stars.[11]

By definition, the distribution function is normalised such that, integrating over all phase space,

$$\int f(\boldsymbol{x},\boldsymbol{v},t)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{v} = 1\ . \tag{1.28}$$

Note that some authors instead choose to set this integral equal to N, and in that case f would be the phase-space *number density*, rather than the *probability density* (e.g. Mo et al. 2010).

Using equation (1.26), it is possible to derive the *Jeans equation(s)*:

$$\frac{\partial\langle v_j\rangle}{\partial t} + \sum_{i=1}^{3}\langle v_i\rangle\frac{\partial\langle v_j\rangle}{\partial x_i} = -\frac{1}{n}\sum_{i=1}^{3}\frac{\partial\left(n\sigma_{ij}^2\right)}{\partial x_i} - \frac{\partial\Phi}{\partial x_j}\ , \tag{1.29}$$

where the angular brackets indicate the *expectation value* of the relevant quantity, obtained through appropriate integration of $f(\boldsymbol{x},\boldsymbol{v},t)$. For example, the spatial *mass density*, $\rho(\boldsymbol{x},t)$, and *number density*, $n(\boldsymbol{x},t)$, are given by

$$\frac{\rho(\boldsymbol{x},t)}{m} = n(\boldsymbol{x},t) = N\int f(\boldsymbol{x},\boldsymbol{v},t)\,\mathrm{d}\boldsymbol{v}\ , \tag{1.30}$$

and thus the *mean velocity* in the i-th direction is

$$\langle v_i\rangle(\boldsymbol{x},t) = \frac{N}{n(\boldsymbol{x},t)}\int v_i f(\boldsymbol{x},\boldsymbol{v},t)\,\mathrm{d}\boldsymbol{v}\ . \tag{1.31}$$

Similarly, the *velocity tensor* is defined as

$$\langle v_i v_j\rangle(\boldsymbol{x},t) = \frac{N}{n(\boldsymbol{x},t)}\int v_i v_j f(\boldsymbol{x},\boldsymbol{v},t)\,\mathrm{d}\boldsymbol{v}\ . \tag{1.32}$$

The velocity tensor can be split into the contributions from coherent stream-

[11] *Jeans theorem* states that: "Any steady-state solution of the collisionless Boltzmann equation depends on the phase-space coordinates only through integrals of motion in the given potential, and any function of the integrals yields a steady-state solution of the collisionless Boltzmann equation." (quoted from Binney & Tremaine 2008). Thus for a steady state ($\partial f/\partial t = 0$) we can write $f = f(I_1, I_2, \ldots, I_k)$ where each of the k I_n are integrals of motion (i.e. functions of $\boldsymbol{x}$ and $\boldsymbol{v}$ alone, and *not* t, that do not change along an orbit).

ing motion and random motion, such that

$$\langle v_i v_j \rangle = \langle v_i \rangle \langle v_j \rangle + \sigma_{ij}^2 \,, \tag{1.33}$$

where σ_{ij}^2 is the symmetric *velocity dispersion tensor*, and gives the contribution from random local motion. The eigenvectors of $\sigma_{ij}(\boldsymbol{x}, t)$ define the *velocity ellipsoid* at position $\boldsymbol{x}$. In terms of the distribution function,

$$\sigma_{ij}^2(\boldsymbol{x}, t) = \frac{N}{n(\boldsymbol{x}, t)} \int (v_i - \langle v_i \rangle)(v_j - \langle v_j \rangle) f(\boldsymbol{x}, \boldsymbol{v}, t) \, \mathrm{d}\boldsymbol{v} \,. \tag{1.34}$$

The quantity $n\sigma_{ij}^2$ is a *stress tensor*, and appears in equation (1.29), where it behaves like an *anisotropic pressure* term (cf. equation 1.20).

Given the potential, $\Phi(\boldsymbol{x}, t)$, and the number[12] density, $n(\boldsymbol{x}, t)$, the (three) Jeans equation(s), as given in equation (1.29), contains nine unknowns: three from the streaming motions in each direction, $\langle v_i \rangle$, and six for the *independent* components of σ_{ij}. To actually make use of equation (1.29), some additional assumptions are required, in order to reduce the number of unknowns. For example, the geometry of the dynamical system of interest could be assumed to be spherically symmetric, axisymmetric (e.g. Hunter 1977; Binney et al. 1990; Cappellari 2008), or triaxial (e.g. Mo et al. 2010). It is important to note that for a dynamical model to be physically meaningful, the distribution function must satisfy $f \geq 0$ everywhere in phase space, but this is by no means guaranteed for the distribution function that corresponds to some model that just so happens to solve the Jeans equations!

In the case of a steady-state spherically symmetric system, the set of Jeans equations can be simplified to arrive at the much more manageable *spherical Jeans equation*, in spherical polar coordinates (r, θ, ϕ):

$$\frac{1}{\rho} \frac{\mathrm{d}\left(\rho \langle v_r^2 \rangle\right)}{\mathrm{d}r} + 2\beta \frac{\langle v_r^2 \rangle}{r} = -\frac{\mathrm{d}\Phi}{\mathrm{d}r} \,, \tag{1.35}$$

[12] In equation (1.29), n only appears in the first term on the right hand side. If n is replaced by ρ here (or e.g. the luminosity density) then this does not modify the result, since ρ is just some global constant times n, and the constant (m) cancels out (recall that the tracer particles are assumed to be identical). The same will be true for equations (1.35) and (1.37), where we choose to write ρ for the tracer density.

where $\beta(r)$ is the *anisotropy parameter*, defined as

$$\beta = 1 - \frac{\langle v_\theta^2 \rangle}{\langle v_r^2 \rangle} \,, \tag{1.36}$$

and $\langle v_\theta^2 \rangle = \langle v_\phi^2 \rangle$, by symmetry. We usually refer to the root-mean-square velocity in the rest frame of the system as the *velocity dispersion* and write e.g. $\sigma_r^2 = \langle v_r^2 \rangle$, where σ_r is the velocity dispersion in the radial direction, and so on for σ_θ and σ_ϕ.[13] Equation (1.35) can be written in a slightly different, but equivalent, form:

$$V_{\rm circ}^2(r) = \frac{GM(<r)}{r} = -\sigma_r^2(r) \left[\frac{{\rm d}\log\rho(r)}{{\rm d}\log r} + \frac{{\rm d}\log\sigma_r^2(r)}{{\rm d}\log r} + 2\beta(r) \right] . \tag{1.37}$$

Given $\rho(r)$ from *deprojecting* the observed (i.e. projected) stellar density profile (assuming spherical symmetry), the total mass profile, $M(<r)$, can be recovered if *both* the profile of the velocity dispersion in the radial direction, $\sigma_r(r)$, *and* the profile of the velocity dispersion anisotropy, $\beta(r)$, are known. Since $\sigma_r(r)$ and $\beta(r)$ are intimately related, it is extremely difficult to robustly constrain $M(<r)$ using projected data alone, since $M(<r)$ is degenerate with $\beta(r)$ in this case. Some assumptions about $\beta(r)$ usually have to be made when carrying out mass modelling of observed galaxies via the spherical Jeans equation, and this may compromise the reliability of such work.

Starting with the spherical Jeans equation, and making some additional assumptions about the stellar mass distribution and velocity dispersion, Walker et al. (2009, 2010) and Wolf et al. (2010) have each derived a simple estimator for the dynamical mass enclosed within a sphere of radius $\lambda R_{\rm e}$, where $R_{\rm e}$ is the radius that encloses half the stellar luminosity, in projection (i.e. the *projected half-light radius* or *effective radius*), and λ is a constant:

$$M(<\lambda R_{\rm e}) = \frac{\mu \langle \sigma_{\rm los} \rangle^2 R_{\rm e}}{G} \,, \tag{1.38}$$

where $\langle \sigma_{\rm los} \rangle$ is the line-of-sight stellar velocity dispersion averaged over the

[13] Since $\sigma_\theta = \sigma_\phi$, we could define $\sigma_t^2 = \sigma_\theta^2 + \sigma_\phi^2 = 2\sigma_\theta^2$ and rewrite the definition of β from equation (1.36) as $\beta = 1 - \sigma_\theta^2/\sigma_r^2 = 1 - \sigma_t^2/(2\sigma_r^2)$, but this is only a matter of notation. Similarly, the total velocity dispersion, $\sigma_{\rm 3D}$, satisfies $\sigma_{\rm 3D}^2 = \sigma_r^2 + \sigma_\theta^2 + \sigma_\phi^2 = \sigma_r^2 + \sigma_t^2 = \sigma_x^2 + \sigma_y^2 + \sigma_z^2$, where σ_x, σ_y, and σ_z are the velocity dispersions in the three Cartesian directions.

whole galaxy (weighted by luminosity), and μ is a second dimensionless constant. For the estimator of Walker et al. (2009), $\lambda = 1$ and $\mu = 2.5$, while Wolf et al. (2010) have $\lambda = 4/3$ and $\mu = 4$. The different values of λ and μ result from the different assumptions applied when deriving these two estimators, where the assumptions of Wolf et al. (2010) are somewhat less restrictive than those of Walker et al. (2009). The application of these estimators is extremely simple, such that total mass within a sphere of radius close to the projected half-light radius can be inferred using only two numbers: the projected half-light radius, and the integrated line-of-sight stellar velocity dispersion. Furthermore, these two estimators are designed to be minimally sensitive to the stellar velocity dispersion anisotropy, β, since $M(< \lambda R_{\rm e})$, given the appropriate value of λ, is expected to be minimally dependent on the precise amplitude and form of $\beta(r)$, and so these simple estimators in principle offer a highly attractive opportunity to escape the dreaded *anisotropy degeneracy* between $\beta(r)$ and $M(< r)$. However, since the estimators have been analytically derived under ideal conditions with various simplifying assumptions, it is vital that we develop an understanding of their accuracy when applied to realistic galaxies in a cosmological context. We seek to address this issue in Chapter 2, making use of the latest generation of cosmological hydrodynamical simulations of Local Group analogues.

1.3.2 Stellar masses

The stellar masses of galaxies are determined by the detailed history of the accretion and cooling of gas, metal enrichment, energy feedback, galaxy mergers, and other violent events, and are closely linked to the masses of their dark matter haloes. Any realistic model of galaxy formation should be able to make reasonable predictions for the stellar masses of galaxies, since the stellar mass is such a fundamental intrinsic galactic property. However, stellar masses are certainly not observable quantities, and ignorance of the stellar mass of a galaxy can seriously hamper attempts to infer the mass distribution of its host dark matter halo (as highlighted in §1.3.1 above). If we can infer the stellar masses of observed galaxies in an accurate and robust way, then there is hope for placing reliable constraints on the galaxy formation physics models implemented in cosmological simulations.

The *spectral energy distribution* (SED) of a galaxy, f_λ, is the energy density in photons as a function of wavelength, defined such that $f_\lambda \, \mathrm{d}\lambda$ is the total energy flux in photons with wavelengths in the interval λ to $\lambda + \mathrm{d}\lambda$. Galactic SEDs thus contain all relevant observational data on the emitted light. In order to estimate the stellar mass of a galaxy, we need to be able to correctly interpret the shape and amplitude of its SED, since this is the integral of the starlight over the whole galaxy, determined by the precise state (and thus, spectrum) of each constituent star at the time of emission (and then potentially modified by absorption or scattering of the light by intervening material). While the SED can be measured in full detail using spectroscopy, observations often provide *broad-band magnitudes*, i.e. integrated magnitudes resulting from applying some particular filter to the SED. Broad-band photometry thus provides low resolution information on the amplitude of the SED in different wavelength ranges (where the precise filter shape as a function of wavelength, i.e. the *wavelength response*, is generally non-trivial). Conversely, *narrow-band* photometry uses much smaller wavelength ranges, and so provides higher resolution data on the SED, and can be designed to target particular spectral features or emission lines of interest. A *colour index* is defined as the difference between the magnitudes in two given photometric bands, e.g. if the (apparent or absolute) magnitudes in the B and V bands are m_B and m_V respectively, then the so-called $(B-V)$ *colour* is $m_B - m_V$.

Stellar population synthesis (SPS) models predict the SEDs of *simple stellar populations* (SSPs), i.e. coeval populations of stars with a given initial *metallicity*, Z (mass fraction in metals), as a function of age (e.g. Bruzual & Charlot 2003; Maraston 2005; Conroy et al. 2009). If the *star formation history* (SFH; star formation rate as a function of time) of a galaxy is known, along with its *chemical enrichment history* (metallicity of the star-forming gas as function of time), then these can be combined with an SPS model to compute the current SED of the galaxy, as the weighted sum of many SSP SEDs of appropriate age and initial metallicity. As part of this, it is necessary to specify the stellar *initial mass function* (IMF), which gives the probability distribution for creating stars of a given mass in any episode of star formation (e.g. Salpeter 1955; Kennicutt 1983; Kroupa 2001; Chabrier 2003). Uncertainty in the shape of the IMF introduces a significant source of error into the predicted galactic SEDs computed using this method. *Dust*

(small particles composed of heavy elements) within the interstellar medium of the galaxy is likely to attenuate the starlight, and so this effect should be accounted for when modelling galactic SEDs (e.g. Calzetti et al. 2000).

In the case where we have measured the SED of an observed galaxy, whether through spectroscopy or photometry with broad or narrow bands, the task of estimating the corresponding stellar mass reduces to reverse engineering the process described above. Given a large library of SSP spectra, varying in age and initial metallicity over the full plausible ranges, some fitting procedure can be used to identify the set of weighted SSP spectra that best reproduces the measured galactic SED. This is equivalent to constraining the star formation and chemical enrichment histories of the galaxy. The stellar mass follows directly from the outcome of the fitting procedure, since the SSP masses are included in the SPS model. SED fitting is further complicated by the need to account for dust attenuation, and by degeneracies between age and metallicity. In general the resulting uncertainties in the estimated stellar masses can be substantial (e.g. Conroy et al. 2009; Mitchell et al. 2013). Certain spectral features and emission lines have particularly strong constraining power. For example, Kauffmann et al. (2003) show that the amplitude of the 4000 Å break is sensitive to the mean age of the stellar population, and that the strength of the H-δ absorption line indicates whether there has been a recent burst of star formation (within the last Gyr or so). Also, Bell & de Jong (2001) show that the stellar mass-to-light ratio, Υ, is closely correlated with the optical broad-band colours, but is fairly insensitive to the SFH (see also Bell et al. 2003; Taylor et al. 2011).

Given the significant uncertainties involved in estimating the stellar masses of observed galaxies via SED fitting, it is clear that caution must be exercised when comparing observationally inferred stellar masses to the predicted stellar masses of galaxies drawn from theoretical models of galaxy formation, where the stellar masses are instead accurately computed from the known star formation histories. As a way of understanding the impact of the uncertainties involved in SED fitting, it is possible to apply this observational technique to simulated galaxies, and to compare the estimated masses from SED fitting to the true, intrinsic, stellar masses predicted by the galaxy formation model (e.g. Mitchell et al. 2013). In Chapter 3 we follow this approach to investigate the impact of the errors introduced by

SED fitting on the perceived clustering of galaxies as a function of stellar mass, making use of two recently developed semi-analytic galaxy formation models, set in a large volume, high resolution, dark matter only simulation.

1.4 Research overview

The primary aim of the work presented in this thesis is to assess the accuracy of particular techniques that are actively used to observationally infer the dynamical and stellar masses of galaxies, in the context of the use of such techniques to draw conclusions of the nature of dark matter, the physics of galaxy formation, and the validity of the standard ΛCDM cosmogony.

In Chapter 2 we test the dynamical mass estimators of Walker et al. (2009) and Wolf et al. (2010), by applying these to galaxies drawn from the high resolution APOSTLE cosmological hydrodynamical simulations of Local Group analogues (Fattahi et al., 2016; Sawala et al., 2016). These simulations are able to resolve the internal structure and kinematics of dwarf satellite galaxies orbiting within the dark matter haloes of the Milky Way and Andromeda analogues, while at the same time self-consistently reproducing the environment of the Local Group on larger scales. We investigate how the accuracy of the estimators depends on certain galaxy properties that are closely connected to the assumptions under which the estimators are derived, and that are of central importance to dynamical studies in general. Based on the levels of estimator error unveiled in our analysis, we reassess the significance with which such estimators can be used to place constraints on the inner dark matter density profile slopes of the satellite galaxies of the Milky Way. We also determine the set of estimator parameters that yields an unbiased result for the estimated mass, with minimum scatter, for our sample of dispersion-dominated simulated galaxies.

In Chapter 3 we use two versions of the GALFORM semi-analytic galaxy formation model (Gonzalez-Perez et al., 2014; Lacey et al., 2016) to investigate the impact of the errors involved in stellar mass estimation on the perceived clustering of galaxies as a function of stellar mass. These models exploit a dark matter only simulation of a large cosmologically representative volume, at sufficiently high resolution to accurately trace the evolution history of dark matter haloes on even the scale of dwarf galaxies. These models

are thus ideally suited to studies of the clustering of galaxies. The correlation function of galaxy clustering is a fundamental statistic that encodes not only the structure of the underlying dark matter, but also the evolution history of the baryonic material from which the galaxies were forged, in the centres of their host dark matter haloes. Building on the work of Mitchell et al. (2013), we estimate the masses of the simulated galaxies using SED fitting procedures designed to emulate those used in particular observational studies of the clustering of galaxies as a function of stellar mass. We then compare the model clustering predictions to the results from the observational studies, using both the true and estimated model stellar masses. We argue that only through such techniques can the clustering of galaxies as a function of stellar mass predicted by galaxy formation models be faithfully compared to observationally derived clustering signals. Following this new methodology will allow the clustering of galaxies as a function of stellar mass to be used as a powerful constraint on the predictions of future galaxy formation models. We also introduce a new scheme for modelling the mergers of satellite galaxies with the central galaxy of their parent halo. This new scheme computes merger time-scales for the satellites in a way that is more faithful to the underlying dark matter simulation, and hence more physically realistic, than the default prescription used in GALFORM. This new treatment of satellite mergers modifies the galaxy clustering on small scales, where the contribution due to satellite galaxies is particularly important.

A summary of our main findings and concluding remarks are given in Chapter 4. Important statistical definitions are given in Appendix A. The accuracy of alternative versions of the mass estimator of Wolf et al. (2010) is investigated in Appendix B. Details of the parameter differences between the GALFORM models used are described in Appendix C.

Chapter 2

Simple Estimators of Galactic Dynamical Masses

2.1 Introduction

Obtaining reliable dynamical information lies at the heart of many fundamental questions in cosmology and galactic structure. However, in nature we typically only have partial information about the kinematics of observed systems. Dynamical studies are therefore intrinsically underconstrained, and often require the use of various simplifying assumptions. In this study, we are interested in stellar dynamical systems to which the spherical Jeans equation may be applied, such as globular clusters, satellite galaxies, and massive spheroidal galaxies. The use of Jeans analysis to infer the mass distribution of a system whose gravitational potential is traced by stars is complicated by an important degeneracy between the velocity dispersion anisotropy of the stars, and the total mass profile of the system. The anisotropy is notoriously difficult to constrain with current data, and therefore studies based on the application of the spherical Jeans equation have led to many ambiguous results. For example, a topic of great interest is whether local satellite galaxies have central cusps or cores in their central dark matter

distributions. Studies based on Jeans analysis are often inconclusive, largely due to the anisotropy degeneracy.

An important advance in this subject has been the application of simple estimators to infer the dynamical mass within a sphere of radius equal to the projected (2D) stellar half-light radius, R_e (Walker et al., 2009, 2010), or the deprojected (3D) stellar half-light radius, r_{half} (Wolf et al., 2010).[1] These estimators make use of only two measured quantities: the stellar line-of-sight velocity dispersion averaged over the whole galaxy, $\langle\sigma_{los}\rangle$, and the projected half-light radius. Each estimator can be written in the form

$$M(<\lambda R_e) = \frac{\mu\langle\sigma_{los}\rangle^2 R_e}{G}\,, \tag{2.1}$$

where $M(<r)$ is the total mass enclosed within a sphere of radius r. λ and μ are dimensionless parameters, and G is the gravitational constant. These estimators are derived from the spherical Jeans equation, which is valid for spherically symmetric, dispersion-supported, collisionless, stationary systems, in dynamical equilibrium, in the absence of streaming motions. Additional assumptions are applied in the derivation of each estimator, with the assumptions used by Walker et al. (2009) being somewhat more restrictive than those of Wolf et al. (2010). The main finding of these studies is that the total mass within a sphere of radius close to the 2D or 3D stellar half-light radius is minimally sensitive to the assumed form of the stellar velocity dispersion anisotropy and the shape of the total mass profile (Wolf et al. 2010 discuss in detail how this relatively tight constraint arises).

Walker & Peñarrubia (2011) have applied these estimators to apparently distinct stellar subpopulations in the Fornax and Sculptor dwarf spheroidals (dSphs), and have argued for the presence of a central core in the dark matter density profile of these galaxies, with higher confidence for Sculptor than for Fornax. (However, this result has been disputed for Sculptor by Strigari et al. 2014.) A further example of the applicability of these mass estimators is the 'too big to fail problem' of the standard Λ cold dark matter (ΛCDM) cosmological model (Boylan-Kolchin et al., 2011), which was predicated on the basis of comparing the circular velocity curves ($V_{circ}(r) = \sqrt{GM(<r)/r}$) of

[1] R_e is the radius of the circle that encloses half of the total stellar luminosity in projection (effective radius), and r_{half} is the radius of the sphere that encloses the same luminosity fraction in 3D.

dark matter subhaloes drawn from the Aquarius (Springel et al., 2008) and Via Lactea II simulations (Diemand et al., 2008) with the dynamical mass within the 3D half-light radius of the brightest satellites of the Milky Way (MW) according to Wolf et al. (2010).

While the assumptions that underpin the spherical Jeans equation may seem relatively benign, they are not exactly satisfied by all dSphs. For example, it is clear that many galaxies, such as Fornax, are not spherically symmetric. It also seems unlikely that satellites orbiting in the potential well of a host halo are in dynamical equilibrium.

The simple mass estimators have been tested on ideal spherically symmetric systems (Walker & Peñarrubia, 2011), simulations of ideal dSphs in a static MW potential (Kowalczyk et al., 2013), stellar distributions placed in subhaloes from dark matter only simulations (Laporte et al., 2013a), and cosmological zoom simulations of isolated elliptical galaxies (Lyskova et al., 2015). More recently, the first generation of simulations with enough resolution to model the satellite galaxies of the MW reliably, in a realistic cosmological setting, have become available (Grand et al., 2016; Sawala et al., 2016; Wetzel et al., 2016). Such simulations can be used to test the accuracy of equation (2.1) in estimating the dynamical masses of model galaxies with realistic properties. In this chapter, we make use of the APOSTLE simulations of the Local Group (Fattahi et al., 2016; Sawala et al., 2016) to study the accuracy of the estimators proposed by Walker et al. (2009) and Wolf et al. (2010).

In §2.2 we describe the mass estimators and the assumptions on which they are based in greater detail. The simulations used in this work are discussed in §2.3. In §2.4 we present general properties of a sample of well-resolved dispersion-dominated galaxies drawn from the simulations, to which we apply the mass estimators in §2.5. The implications of the systematic errors on the estimated masses are discussed in §2.6. In §2.7 we find the set of estimator parameters that optimises the accuracy of the recovered mass for dispersion-dominated galaxies. Concluding remarks are given in §2.8. In Appendix B we investigate the accuracy of the estimator of Wolf et al. (2010) when written in terms of the 3D half-light radius. All simulation results presented in this chapter are at redshift zero. The assumed ΛCDM cosmological parameters are given in §2.3.1.

2.2 Mass estimators

The spherical Jeans equation relates the circular velocity curve of a spherically symmetric system to the radial distribution and velocity dispersion of a population of tracers that orbit in the gravitational potential (see §2.1 for a list of the assumptions on which this equation is based, and see §1.3.1 for an outline of its derivation):

$$\frac{GM(<r)}{r} = -\sigma_r^2(r)\left[\frac{\mathrm{d}\log\rho(r)}{\mathrm{d}\log r} + \frac{\mathrm{d}\log\sigma_r^2(r)}{\mathrm{d}\log r} + 2\beta(r)\right], \tag{2.2}$$

where $M(<r)$ is the total mass enclosed within a radius r of the centre, $\rho(r)$ is the density profile of the tracer population,[2] $\sigma_r(r)$ is the profile of its velocity dispersion, in the radial direction, and $\beta(r)$ is the tracer velocity dispersion anisotropy, which encodes the balance between the tangential and radial components of the velocity dispersion. At a given radius,

$$\beta(r) = 1 - \frac{\sigma_t^2(r)}{2\sigma_r^2(r)}, \tag{2.3}$$

where $\sigma_t(r)$ is the tangential component of the velocity dispersion, such that the total (3D) dispersion is given by

$$\sigma_{3\mathrm{D}}(r) = \sqrt{\sigma_r^2(r) + \sigma_t^2(r)}\,. \tag{2.4}$$

Thus defined, $\sigma_t(r)$ includes the contributions to the total velocity dispersion from two mutually orthogonal tangential directions, which are equivalent under the assumption of spherical symmetry.

When applied to a dSph galaxy, for example, equation (2.2) relates the total mass profile to the spatial distribution and kinematics of stars orbiting in the total gravitational potential (which is dominated by the dark matter). For an observed galaxy, we can measure the projected stellar density, $\Sigma(R)$, and velocity dispersion along the line of sight,[3] $\sigma_{\mathrm{los}}(R)$. A deprojection is required to map these observable profiles onto the 3D profiles that appear on

[2] $\rho(r)$ is the *number* density profile of the tracer population, or indeed, the mass (or luminosity) density profile, assuming that all tracers have the same mass (or luminosity).

[3] In this chapter, a lower case r is used to denote a 3D radius, and an upper case R is used to denote a 2D, projected, radius.

the right-hand-side of equation (2.2). The projected and 3D stellar density profiles, $\Sigma(R)$ and $\rho(r)$, under the assumption of spherical symmetry, are related by an Abel transform, such that

$$\rho(r) = -\frac{1}{\pi}\int_r^\infty \frac{\mathrm{d}\Sigma(R)}{\mathrm{d}R}\frac{\mathrm{d}R}{\sqrt{R^2-r^2}}\,. \tag{2.5}$$

Given only line-of-sight kinematic data for the stars, there exists an inconvenient degeneracy between $\beta(r)$ and $\sigma_r(r)$, such that

$$\Sigma(R)\sigma^2_{\rm los}(R) = 2\int_R^\infty \rho(r)\sigma_r^2(r)\left[1-\frac{R^2}{r^2}\beta(r)\right]\frac{r\mathrm{d}r}{\sqrt{r^2-R^2}}\,, \tag{2.6}$$

(Binney & Mamon, 1982). In this way, our ignorance of $\beta(r)$ influences both the deprojection of $\sigma_{\rm los}(R)$ onto $\sigma_r(r)$ in equation (2.6), and the subsequent inference of the enclosed mass profile from equation (2.2). In order to make progress, some assumption about the stellar velocity dispersion anisotropy is required, ranging from assuming isotropy ($\beta = 0$), to exploring a wide range of plausible forms for $\beta(r)$ in the fitting procedure (e.g. Wolf et al. 2010).

Through Jeans analysis of eight of the brightest dSphs of the MW, assuming constant β, Walker et al. (2009, 2010) find that the total mass within a sphere of radius equal to the projected stellar half-light radius (effective radius), $R_{\rm e}$, is relatively well constrained (compared to smaller and larger radii), and seemingly robust against the choice of β, and of the assumed shape of the total density profile (see also Strigari et al. 2007, 2008; Peñarrubia et al. 2008a,b; Wolf et al. 2010; Amorisco & Evans 2011). Assuming that the stars follow a Plummer density profile, with a constant and isotropic velocity dispersion, Walker et al. (2009) propose a simple estimator for the total mass enclosed within a sphere of radius $R_{\rm e}$, which arises immediately from equation (2.2) given these additional assumptions,

$$M(<R_{\rm e}) = \frac{5\langle\sigma_{\rm los}\rangle^2 R_{\rm e}}{2G}\,, \tag{2.7}$$

where $\langle\sigma_{\rm los}\rangle$ is the (assumed to be constant) line-of-sight stellar velocity dispersion averaged over the whole galaxy. That is, a constraint on the dynamical mass (or equivalently, a point on the circular velocity curve) can be obtained using only the stellar half-light radius and a single value for

the stellar velocity dispersion, which does not require spatially resolved kinematic data.

Wolf et al. (2010), on the other hand, argue that the 3D radius within which the sensitivity of the enclosed mass to β is minimised is in fact r_3, the radius at which the logarithmic slope of the stellar density profile, $\mathrm{d}\log\rho(r)/\mathrm{d}\log r$, equals -3. Assuming that r_3 is approximately the stellar 3D half-light radius, $r_{\rm half}$, these authors propose the estimator

$$M(< r_{\rm half}) = \frac{3\langle\sigma_{\rm los}\rangle^2 r_{\rm half}}{G}\,, \tag{2.8}$$

where they stress that $\langle\sigma_{\rm los}\rangle$ must be the luminosity-weighted mean dispersion. This estimator is based on the assumption that $\sigma_{\rm los}(R)$ remains relatively flat from the centre of the system out to beyond $R_{\rm e}$, and that $\beta(r)$ does not have an extremum within the stellar distribution. To express equation (2.8) entirely in terms of observable quantities, Wolf et al. (2010) make the further simplifying assumption that $r_{\rm half} = 4R_{\rm e}/3$, in which case

$$M\left(< \frac{4}{3}R_{\rm e}\right) = \frac{4\langle\sigma_{\rm los}\rangle^2 R_{\rm e}}{G}\,. \tag{2.9}$$

This estimator is of the same form as equation (2.7), but with $\lambda = 4/3$ and $\mu = 4$, compared to $\lambda = 1$ and $\mu = 2.5$, in the notation of equation (2.1). It is worth noting that if both equations (2.7) and (2.9) apply, then the enclosed dynamical mass increases by a factor of 1.6 from $R_{\rm e}$ to $4R_{\rm e}/3$. We test the accuracy of equations (2.7) and (2.9) using simulated galaxies in §2.5, and return to equation (2.8) in Appendix B.

2.3 Simulations

We now describe the simulations used in this chapter, followed by the definitions employed for subhaloes and galaxies, and the resolution limit we impose to ensure converged galaxy properties.

2.3.1 APOSTLE simulations

The APOSTLE simulations use the 'zoom' technique to resimulate 12 systems consistent with observational constraints on the properties of the Local Group, as described in more detail by Fattahi et al. (2016) and Sawala et al. (2016). The regions were selected from a dark matter only simulation of a cosmologically representative volume of comoving side 100 Mpc, using ΛCDM parameters consistent with WMAP7 (Komatsu et al., 2011). The density parameters at redshift zero are $\Omega_{\rm m,0} = 0.272$ (matter), $\Omega_{\rm b,0} = 0.0455$ (baryons), and $\Omega_{\Lambda,0} = 1 - \Omega_{\rm m,0} = 0.728$ (cosmological constant). The present day Hubble parameter is $H_0 = 100\,h\,{\rm km\,s^{-1}\,Mpc^{-1}}$, with $h = 0.704$. The linear power spectrum is normalised using $\sigma_8 = 0.81$ at redshift zero. The spectral index of primordial fluctuations is $n_{\rm s} = 0.967$. The pairs of MW and M31 analogues were selected on the basis of their separations, relative radial and tangential velocities, and halo masses, along with the recession velocities of outer Local Group members (see Fattahi et al. 2016 for details).

The 12 regions in the APOSTLE suite were simulated using the code from the EAGLE project (Crain et al., 2015; Schaye et al., 2015), at a series of resolution levels, which we label as LR, MR, and HR, in order of increasing resolution (low, medium, high). Table 2.1 lists the dark matter and (initial) gas particle masses in the zoom region at each resolution level, along with the gravitational force softening. All regions have been simulated at LR and MR, and two regions have also been run at HR. The EAGLE code is a version of the smoothed particle hydrodynamics (SPH) code GADGET (Springel, 2005) that includes 'subgrid' models for radiative gas cooling and heating (Wiersma et al., 2009a), reionisation, star formation (Schaye, 2004; Schaye & Dalla Vecchia, 2008), stellar mass loss and metal enrichment (Wiersma et al., 2009b), stellar feedback (Dalla Vecchia & Schaye, 2012), black hole formation and mergers (Rosas-Guevara et al., 2015), and feedback from active galactic nuclei (Booth & Schaye, 2009). See Crain et al. (2015) and Schaye et al. (2015) for full details of the subgrid models. The hydrodynamics implementation used is ANARCHY (Dalla Vecchia in preparation), which uses the conservative pressure-entropy SPH formulation derived by Hopkins (2013). See Schaller et al. (2015) for a description of ANARCHY and the impact of the changes with respect to the original GADGET SPH scheme on galaxy properties in the

Table 2.1 – Parameters for each resolution level in the APOSTLE simulations. $m_{\rm DM}$ and $m_{\rm gas}$ are the high resolution (zoom) dark matter and initial gas particle masses respectively. $\epsilon(z=0)$ is the Plummer-equivalent gravitational force softening at redshift zero. The gravitational force is Newtonian on scales larger than 2.8ϵ. There is a small amount of variation in the particle masses used in different simulations at a given resolution level; average values are quoted here (see Fattahi et al. 2016 for the individual values).

Resolution	$m_{\rm DM}$ [$\mathrm{M}_\odot$]	$m_{\rm gas}$ [$\mathrm{M}_\odot$]	$\epsilon(z=0)$ [pc]
LR	7.1×10^6	1.4×10^6	710
MR	5.8×10^5	1.2×10^5	307
HR	3.7×10^4	7.4×10^3	134

EAGLE simulations. The APOSTLE simulations use the 'reference' EAGLE model parameters as described by Schaye et al. (2015).

2.3.2 Halo finding and galaxy definition

To identify haloes in the simulations, we first make use of the friends-of-friends (FOF) algorithm, considering only dark matter particles, with a linking length of 0.2 times the mean interparticle separation (Davis et al., 1985). Baryonic particles are assigned to the same FOF group as their nearest dark matter particle. Each FOF group is then processed using SUBFIND, which identifies overdense gravitationally self-bound 'subhaloes' (Springel et al., 2001; Dolag et al., 2009). The position of a subhalo is taken to be that of the particle with the minimum value of the gravitational potential. The 'main subhalo' of a FOF group is that whose such particle has the lowest potential in the group.[4] All other subhaloes are embedded in the main subhalo (a given particle can belong to at most one subhalo). As described by Schaye et al. (2015), we combine any two subhaloes whose separation is less than the stellar half-mass radius of either subhalo, for separations of at most 3 physical kpc. This final adjustment is designed to absorb a very small

[4]Note that we still refer to this main halo from SUBFIND as a *sub*halo, even though it constitutes the main component of the FOF group. For our purposes, the distinction between the main subhalo and other subhaloes in a FOF group is not important. It is typical in the literature to label the galaxy found in the main subhalo as the 'central' galaxy of the FOF group, while other subhaloes host 'satellite' galaxies. However, this nomenclature can be confusing in some cases, e.g. in the APOSTLE simulations the MW and M31 analogues can reside in the same FOF group.

number of very low mass subhaloes that are dominated by a single baryonic particle of unusually high mass (as a result of exceptional stellar mass loss to a gas particle, or black hole growth).

We define a galaxy to be the set of subhalo star particles within a spherical aperture of radius $r_{\rm gal}$ equal to 15 percent of the virial radius, r_{200}, as measured from the subhalo centre.[5] For subhaloes where the actual value of r_{200} is not meaningful, e.g. for the subhalo of a satellite galaxy embedded in a much larger halo, we adopt the value of r_{200} obtained using a relation between r_{200} and the maximum value of the subhalo circular velocity curve, $V_{\rm max}$, calibrated using main subhaloes in the EAGLE Ref-L025N0752 simulation.[6] In practice, we make use of the directly measured value of r_{200} when computing $r_{\rm gal}$ for all main subhaloes, and for any other subhalo for which r_{200} is less than the distance to the furthest subhalo particle. All galaxy properties presented in this chapter are computed using the set of subhalo star particles within $r_{\rm gal}$, with the galaxy centre set to that of its host subhalo. We adopt the rest frame of the centre of mass of the set of star particles defined in this way.

In each simulation volume, we shall refer to the MW and M31 analogues as the 'primary' galaxies. Galaxies within 300 kpc of the centre of either primary galaxy are labelled as 'satellites', and those at larger distances are labelled as 'field' galaxies. This classification is independent of the particular FOF group in which a galaxy resides.

2.3.3 Selecting a well-resolved galaxy sample

We assume that the luminosity of a star particle (which represents a simple stellar population) is proportional to its mass. In this case, $R_{\rm e}$ is the projected radius that encloses half the stellar mass, as well as half the total stellar luminosity. Fig. 2.1 shows $R_{\rm e}$ versus stellar mass, $M_{\rm star}$, for all galaxies in the APOSTLE simulations that are composed of at least ten star particles, split by resolution level. The $R_{\rm e}$ values are averaged from projecting over 1536

[5] r_{200} is the radius of the sphere that encloses a mean density equal to 200 times the critical density of the Universe.

[6] Ref-L025N0752 has the highest resolution available in the EAGLE simulation suite, which is similar to that of the MR APOSTLE simulations. The cosmological parameters assumed in EAGLE are slightly different to those used in the APOSTLE simulations, however this is not important for our purposes.

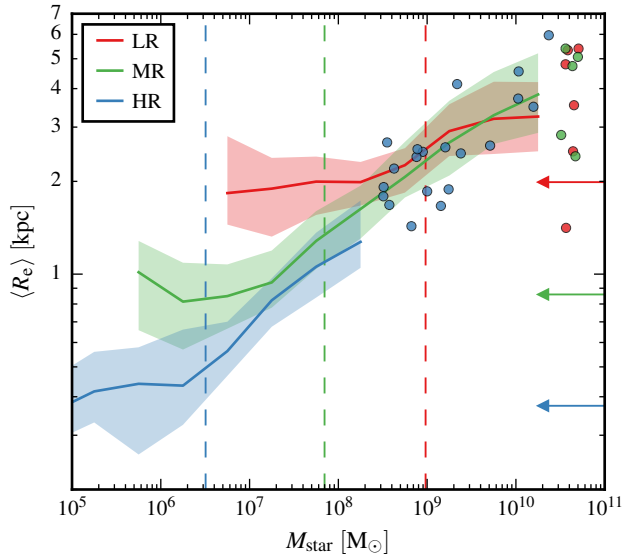

Figure 2.1 – Mean projected stellar half-mass radius, $\langle R_e \rangle$, versus stellar mass, M_{star}, for all galaxies that have at least ten star particles in the APOSTLE simulations at each resolution level (12 LR, 12 MR, and 2 HR simulations; different colours). The $\langle R_e \rangle$ values are averaged over 1536 evenly distributed projections. For each resolution, the median is shown as a solid line, and the $16^{th} - 84^{th}$ percentiles are indicated by a shaded region of the same colour. Points are shown instead for bins with fewer than ten galaxies. The vertical dashed lines show the minimum M_{star} of the set of galaxies composed of at least 1000 star particles at each resolution. The horizontal arrows indicate 2.8 times the gravitational softening at each resolution (see Table 2.1).

evenly distributed lines of sight for each galaxy, making use of the HEALPix spherical tessellation (Górski et al., 2005).[7] This set of sight-lines will be exploited throughout this chapter.

In the following, we consider all galaxies in the highest resolution realisation of each of the 12 APOSTLE regions (10 MR and 2 HR simulations) that are resolved with at least 1000 star particles. This conservative threshold has been chosen so that the stellar mass-size relation is converged at each resolution level (see also Schaye et al. 2015). The vertical lines in Fig. 2.1 show the minimum $M_{\rm star}$ of the set of galaxies with at least 1000 star particles at each resolution. Note that the masses of star particles vary according to the level of enrichment of the gas from which they formed, and the extent of their own mass loss since their birth. There is also a small amount of variation in the initial gas particle mass (and dark matter particle mass) between different simulations at each resolution (see Table 2.1).

2.4 General galaxy properties

In this section, we select a sample of simulated galaxies whose stellar kinematics are dominated by dispersion. We then investigate the basic properties of the galaxies in order to assess how realistic they are. We present radial profiles of the projected stellar density and line-of-sight velocity dispersion for galaxies that are resolved at the highest resolution level, along with profiles of their stellar velocity dispersion anisotropy. These profiles are central to the mapping between observables and the total 3D mass profile (see equations 2.2, 2.5, and 2.6), and thus represent the basic dynamical quantities that enter into the Jeans analysis. We then place the simulated galaxies in the context of the Local Group by comparing their stellar masses, projected half-mass radii, and velocity dispersions with observational data. Finally, we investigate the degree to which the simulated galaxies are spherically symmetric and in dynamical equilibrium.

[7] This corresponds to 3072 HEALPix pixels.

2.4.1 Balance between dispersion and rotational support

A key assumption underpinning the spherical Jeans equation is that the system under consideration is supported against gravitational collapse by dispersion, rather than rotational or other streaming motion. This clearly is not the case for galaxies with prominent stellar discs. To identify systems with significant rotational support, we make use of the quantity $\kappa_{\rm rot}$ as introduced by Sales et al. (2012), which they define as 'the fraction of kinetic energy invested in ordered rotation'. $\kappa_{\rm rot}$ is computed as:

$$\kappa_{\rm rot} = \frac{1}{K_{\rm star}} \sum_i \frac{m_i}{2} \left(\frac{j_{z,i}}{R_{xy,i}} \right)^2 , \tag{2.10}$$

where $K_{\rm star}$ is the total stellar kinetic energy. For each star particle, i, of mass m_i, $j_{z,i}$ is the component of its specific angular momentum in the direction of the total stellar angular momentum vector, $\boldsymbol{L}_{\rm star}$, and $R_{xy,i}$ is its distance from the axis (z) defined by $\boldsymbol{L}_{\rm star}$.

Strictly speaking, $\kappa_{\rm rot}$ is not directly sensitive to ordered rotation, because the sign of the rotation about the z axis is lost in the squared term in equation (2.10). Yet systems that exhibit strong rotation have high values of $\kappa_{\rm rot}$, which makes $\kappa_{\rm rot}$ a useful measure to discriminate between systems that are dispersion-dominated and those dominated by rotating discs. A pure disc galaxy with stars on perfect circular orbits would have $\kappa_{\rm rot} = 1$, while instead $\kappa_{\rm rot} = 1/3$ for a dispersion-supported system with isotropic orbits (since $K_{\rm star}$ includes all three orthogonal velocity components, but the summation in equation 2.10 considers only one such component). Thus, $\kappa_{\rm rot}$ decreases from unity as random motion becomes more important, with a lower limit in the region of $\kappa_{\rm rot} \sim 1/3$. However, dispersion-supported systems with radially biased orbits can have $\kappa_{\rm rot} < 1/3$, since less of the kinetic energy is invested in tangential motion than in the isotropic case. Similarly, tangentially biased orbits imply $\kappa_{\rm rot} > 1/3$.

Fig. 2.2 shows $\kappa_{\rm rot}$ versus $M_{\rm star}$ for all galaxies in our sample (as defined in §2.3.3). It can be seen that the simulations predict a broad range of stellar morphologies, according to this kinematic measure. For $M_{\rm star} \gtrsim 10^9 {\rm M}_\odot$, rotation-dominated galaxies ($\kappa_{\rm rot} > 0.5$) are slightly more prevalent than those that are dominated by dispersion ($\kappa_{\rm rot} < 0.5$). However, at progress-

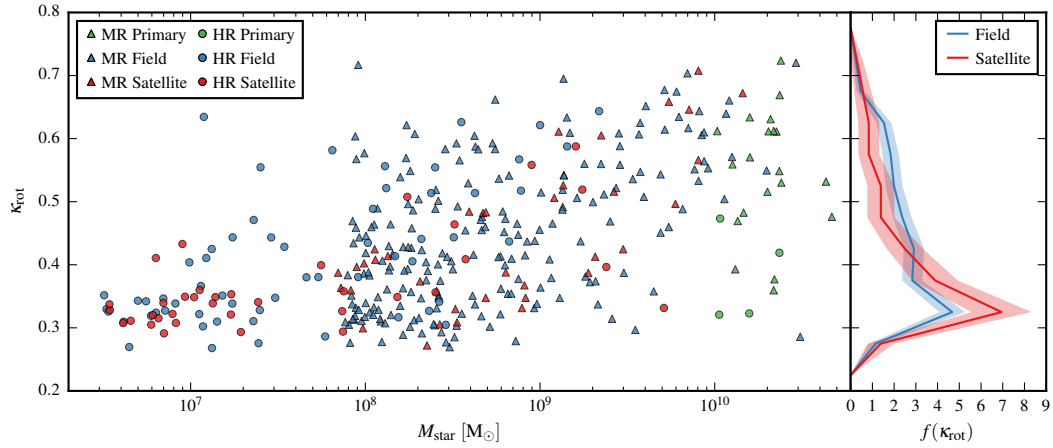

Figure 2.2 – Kinematic measure, $\kappa_{\rm rot}$, used to discriminate amongst galaxies with different levels of stellar rotational support (see equation 2.10), shown versus stellar mass, $M_{\rm star}$, for all galaxies in our simulated sample. The galaxies are classified according to whether they are one of the MW or M31 analogues (primary), within 300 kpc of either of these (satellite), or located at larger distances (field). The symbol shapes indicate the resolution level at which each galaxy has been simulated (MR or HR). Our dispersion-dominated galaxy sample consists of all galaxies that have $\kappa_{\rm rot} < 0.5$. The panel on the right shows the distributions, $f(\kappa_{\rm rot})$, of $\kappa_{\rm rot}$ for all field and satellite galaxies. Each distribution is normalised to have unit area. The shaded regions show the standard deviation assuming Poisson noise on each bin count.

ively lower stellar masses, the galaxies are progressively more likely to be dispersion-dominated. For $M_{\rm star} \lesssim 10^8\,{\rm M}_\odot$, only a small fraction of galaxies have significant levels of rotational support. The distributions of $\kappa_{\rm rot}$ for satellites and field galaxies are not significantly different. The MR and HR simulations predict similar distributions of $\kappa_{\rm rot}$ at a given $M_{\rm star}$, over the common mass range explored.

In this chapter, our main focus is on galaxies with dispersion-dominated stellar kinematics, i.e. those with $\kappa_{\rm rot} < 0.5$. The galaxy sample defined in this way includes 70 percent of our full sample.

2.4.2 Stellar density and kinematic profiles

Projected stellar density profiles, $\Sigma(R)$, for dispersion-dominated field and satellite galaxies in the HR simulations are shown in Fig. 2.3, split into bins of stellar mass. The profiles are obtained by projecting over a large number of lines of sight (as described in §2.3.3), and are stacked for clarity, showing field and satellite galaxies separately. There is no significant difference between

the profiles for field and satellite galaxies in any mass bin. There is a small variation in the shape of the profiles with increasing stellar mass, such that the stellar density is more centrally concentrated in the highest mass bin than in the other bins (see dashed line, which repeats the median profile from the lowest mass bin). Projected stellar number density profile data for bright dSphs of the MW are shown alongside the simulation predictions. The measured profiles have been rescaled assuming the best fitting of either a Plummer or Gaussian density profile model (with $M_{\rm star}$ and $R_{\rm e}$ as free parameters). We find that the data points for Leo I are best described by a Gaussian profile, while the other galaxies shown each prefer a Plummer profile. The simulation predictions closely trace the observational data, in which the degree of scatter varies for each dSph. Thus, we can see that the simulated galaxies have realistic stellar density distributions.

Line-of-sight velocity dispersion profiles, $\sigma_{\rm los}(R)$, are shown in Fig. 2.4 for the same set of simulated galaxies.[8] These profiles show each galaxy individually, using the median result from all lines of sight. The profiles tend to be quite flat with radius. The typical dispersions and radial extents of the galaxies scale closely with $M_{\rm star}$, and there are two particularly large satellites with dispersions in excess of $\sim 50\,{\rm km\,s^{-1}}$ (note that there are two HR simulations, and hence four primary galaxies, which between them host the satellites shown). Recall that an assumption used in the derivation of each of the mass estimators discussed in §2.2 is that $\sigma_{\rm los}(R)$ is relatively flat (or even constant). To see the typical profile shapes more clearly, and to identify if there are systematic differences in $\sigma_{\rm los}(R)$ for field and satellite galaxies, we show scaled versions of these profiles in Fig. 2.5, averaging over all lines of sight in our standard set. While the scatter in the profiles for both types of galaxy is large, such that there is no significant difference between the two populations, the median profiles are somewhat flatter for field than for satellite galaxies, in the two smaller mass intervals. The satellites in these panels tend to have relatively high velocity dispersions in their centres, and relatively low dispersions in their outer parts, compared with field galaxies. The median profiles intersect near the projected half-mass radius. The data

[8] When computing velocity dispersions, we weight by the particle mass. For example, the squared velocity dispersion in the x direction is given by $\sigma_x^2 = (\sum_i m_i v_{x,i}^2)/(\sum_i m_i)$, where each particle, i, has mass m_i, and velocity $v_{x,i}$ in the x direction.

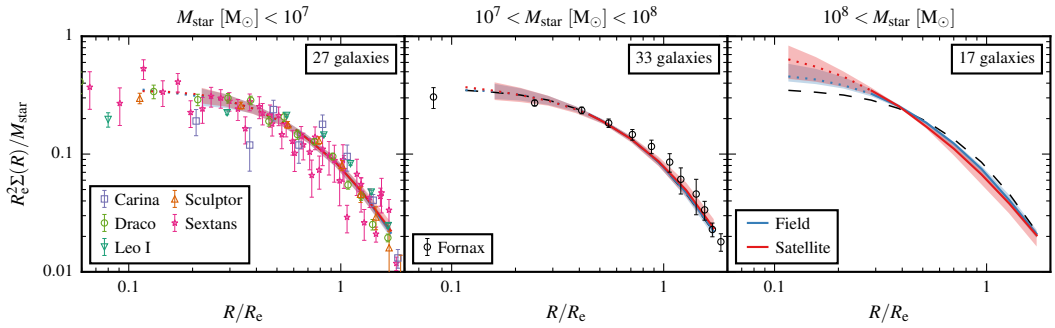

Figure 2.3 – Projected stellar density profiles, $\Sigma(R)$, for dispersion-dominated ($\kappa_{\rm rot} < 0.5$) field and satellite galaxies in the HR simulations. The solid lines are the median projected profiles obtained from 1536 evenly distributed projections for all galaxies of a given type (different colours; legend in right panel), around which the shaded regions of the same colour indicate the $16^{\rm th}-84^{\rm th}$ percentile spread. Each individual profile is scaled, prior to stacking, by the projected stellar half-mass radius, $R_{\rm e}$, for that line of sight, and the stellar mass, $M_{\rm star}$. The shaded regions are not shown below the gravitational softening (see Table 2.1; in units of the median $R_{\rm e}$ in each panel), and the median profiles are dotted below 2.8 times the softening. The panels show different stellar mass ranges, and are labelled with the number of simulated galaxies they contain. As a visual aid, the median profile for all galaxies in the lowest mass bin is repeated as a dashed black line in the higher mass bins. The symbols with error bars show data for dSph satellites of the MW, where the data for each dSph are plotted in the relevant panel for its stellar mass following McConnachie (2012). The data shown are for Carina (Muñoz et al., 2006), Draco (Odenkirchen et al., 2001), Fornax (Coleman et al., 2005), Leo I (Smolčić et al., 2007), Sculptor (Battaglia et al., 2008), and Sextans (Irwin & Hatzidimitriou, 1995). The measurements are scaled in the same way as the simulation predictions, assuming either a Gaussian (Leo I) or Plummer (all others) density profile fit.

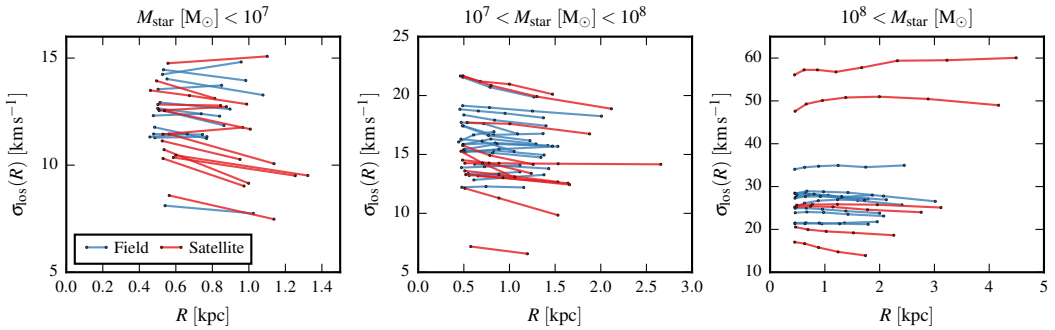

Figure 2.4 – Line-of-sight stellar velocity dispersion profiles, $\sigma_{\rm los}(R)$, for dispersion-dominated ($\kappa_{\rm rot} < 0.5$) field and satellite galaxies in the HR simulations, where each panel shows a different stellar mass range. The line shown for each galaxy is the median profile obtained from projecting over 1536 evenly distributed lines of sight, considering star particles at projected radii between 2.8 times the gravitational softening (see Table 2.1) and twice the mean projected stellar half-mass radius of the galaxy. The bin edges are evenly spaced in the logarithm of R, and the points shown are the linear means of the bin edges. The number of bins used for each galaxy is a function of its star particle count. Note that the axis limits differ between panels.

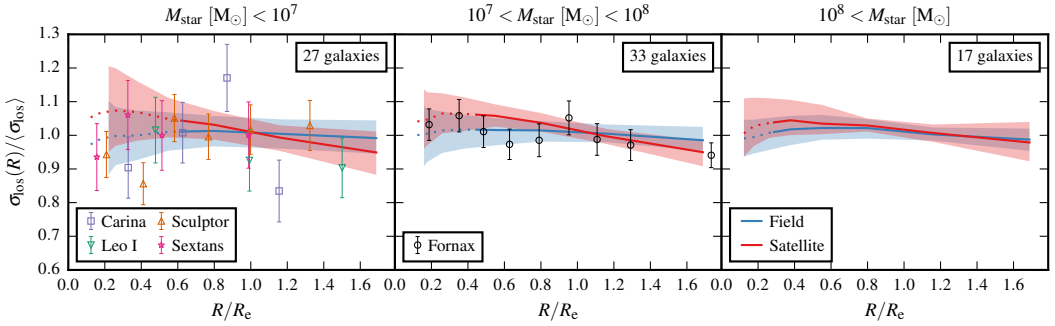

Figure 2.5 – Line-of-sight stellar velocity dispersion profiles, $\sigma_{\rm los}(R)$, for dispersion-dominated ($\kappa_{\rm rot} < 0.5$) field and satellite galaxies in the HR simulations, where each panel shows a different stellar mass range (labelled with the number of galaxies). The solid lines are the median profiles obtained from 1536 evenly distributed projections for all galaxies of a given type (different colours; legend in right panel), around which the shaded regions of the same colour indicate the $16^{\rm th}$–$84^{\rm th}$ percentile spread. Each individual profile is scaled, prior to stacking, by the projected stellar half-mass radius, $R_{\rm e}$, and the mean stellar velocity dispersion of the whole galaxy, $\langle\sigma_{\rm los}\rangle$, for that line of sight. The shaded regions are not shown below the gravitational softening (see Table 2.1; in units of the median $R_{\rm e}$ in each panel), and the median profiles are dotted below 2.8 times the softening. The symbols with error bars show data for dSph satellites of the MW, where the data for each dSph are plotted in the relevant panel for its stellar mass following McConnachie (2012). The profiles shown are for Carina, Fornax, Leo I, Sculptor, and Sextans, as computed by Strigari et al. (2010). The measurements are scaled in the same way as the simulation predictions, using the half-light radii from McConnachie (2012).

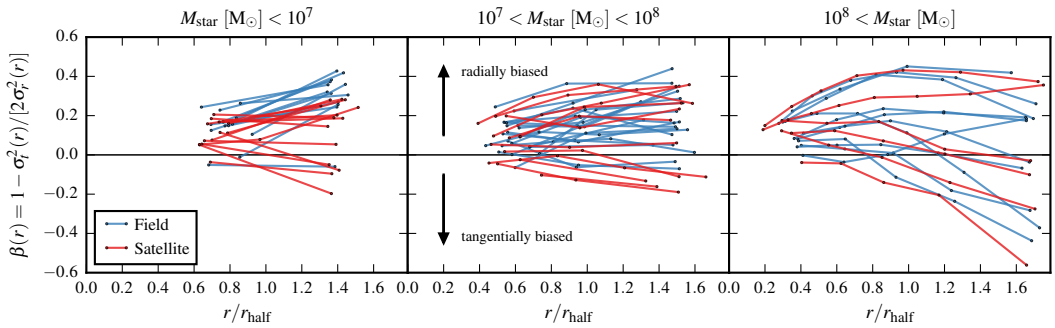

Figure 2.6 – Spherically averaged stellar velocity dispersion anisotropy profiles, $\beta(r)$, for dispersion-dominated ($\kappa_{\rm rot} < 0.5$) field and satellite galaxies in the HR simulations. $\beta(r)$ is negative for tangentially biased dispersions, zero for isotropy, and positive for radially biased dispersions (see equation 2.3). The panels show different stellar mass ranges, as labelled. The profiles include all star particles between a radius of 2.8 times the softening (see Table 2.1) and twice the 3D stellar half-mass radius, $r_{\rm half}$. The bin edges are evenly spaced percentiles of the radial distribution of the star particles, and the points show the median galactocentric radius of the star particles in each bin. The number of bins used for a given galaxy depends on how many star particles it has.

points in Fig. 2.5 show profiles for bright dSphs of the MW (Strigari et al., 2010). The simulated galaxies reproduce the flatness in $\sigma_{\rm los}(R)$ observed for the real MW dSphs.

Fig. 2.6 shows spherically averaged profiles of the velocity dispersion anisotropy, $\beta(r)$, defined in equation (2.3), again examining the highest resolution dispersion-dominated field and satellite galaxies, split by stellar mass. The galaxies predominantly have radially biased dispersions ($\beta > 0$), and only a small number of galaxies, which are in the most massive subset, have $\beta \lesssim -0.2$ at any radius. The profiles are similar in form for field galaxies and the satellites, and there are no obvious distinctions between the two populations, except that in the lowest mass bin the field galaxies tend to be slightly more radially biased than the satellites, at both small and large radii. The anisotropy introduces a key degeneracy in Jeans analysis (see §2.2), and the mass estimators of interest here assume either that $\beta(r)$ is monotonic, or indeed that β has a negligible impact on the recovered mass, and so can be assumed to equal zero.

2.4.3 Galaxy sizes and integrated dispersions

The dynamical mass estimators described in §2.2 make use of only two measurements: the projected stellar half-light radius, and the line-of-sight stellar velocity dispersion. In Fig. 2.7 we plot these two quantities against stellar mass for all galaxies in our sample. The simulations predict a relatively tight relationship between $\langle\sigma_{\rm los}\rangle$ and $M_{\rm star}$, while the relationship between $R_{\rm e}$ and $M_{\rm star}$ exhibits a somewhat higher level of scatter. There is clear convergence between the results from the two resolution levels, which for the stellar mass-size relation reflects our chosen threshold of 1000 star particles (see §2.3.3). There is no obvious difference between the velocity dispersions or sizes of field and satellite galaxies at a given stellar mass.

Observational data for galaxies within 3 Mpc of the Sun as compiled by McConnachie (2012) are shown alongside the simulation predictions in Fig. 2.7. Given our conservative limit on the number of star particles required for a galaxy to be included in our sample, we consider simulated galaxies down to stellar masses close to that of Sculptor, although the simulations do contain galaxies as faint as Draco. The simulation predictions are consistent with the observational constraints, over the large range in stellar mass shown.

It is apparent in Fig. 2.7 that the simulations predict somewhat less scatter in size at fixed stellar mass than seen in the observational data. The majority of observed galaxies, including the Magellanic Clouds and Fornax, lie on the predicted relation but a substantial fraction are smaller, for their stellar mass, than the smallest galaxies in the simulations. As shown in Fig. 2.1, at each resolution level, the stellar mass-size relation flattens out at low masses, and the minimum size scales in the same way as the gravitational softening (the horizontal arrows show the scale above which the force is Newtonian, at each resolution). It thus appears that the resolution in the simulations is not quite high enough to account for the sizes of all the observed galaxies that have stellar masses close to those of the smallest resolved galaxies in the simulations. However, for our purpose, which is to test the validity of dynamical mass estimators using the simulated galaxies, it does not matter if resolution effects have had a marginal impact on the galaxy sizes, provided that the galaxies have self-consistent internal dynamics.

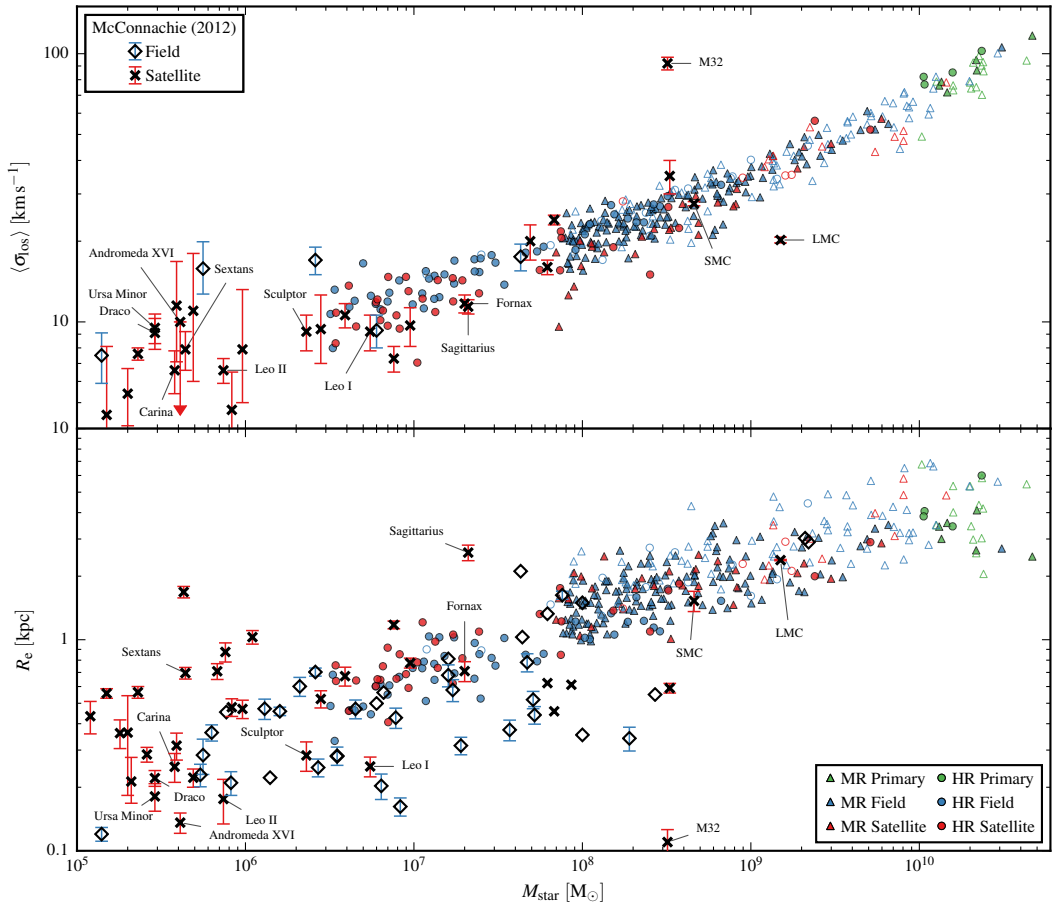

Figure 2.7 – Line-of-sight stellar velocity dispersion, $\langle\sigma_{los}\rangle$ (upper), and projected stellar half-mass radius, R_e (lower), versus stellar mass, M_{star}, for all galaxies in our simulated sample. The $\langle\sigma_{los}\rangle$ and R_e values have been computed for a single, randomly chosen, line of sight through the simulations. The galaxies are classified according to whether they are one of the MW or M31 analogues (primary), within 300 kpc of either of these (satellite), or at larger distances (field), as indicated in the lower right legend (different colours). The symbol shapes indicate the resolution level at which the galaxies have been simulated (MR or HR). Dispersion-dominated galaxies ($\kappa_{rot} < 0.5$) are plotted as filled symbols, while rotation-dominated galaxies ($\kappa_{rot} > 0.5$) are plotted as unfilled symbols of the same colour and shape. *continued...*

Figure 2.7 – *continued...* Data compiled by McConnachie (2012) for galaxies in this stellar mass range within 3 Mpc of the Sun (excluding the MW and M31) are shown alongside the simulation predictions (upper left legend). The observed galaxies are categorised as satellites if they are associated with the MW or M31, otherwise they are classed as field galaxies. The error bars indicate the published uncertainties on the radii and velocity dispersions (where available). We supplement the McConnachie (2012) data with half-light radii for the SMC and LMC from Subramanian & Subramaniam (2012) and Weinberg & Nikolaev (2001, disc model with no bar) respectively, integrating the published exponential density profile fits in each case. The labelled points are the 11 'classical' satellites of the MW, along with M32 (which has an exceptionally large dispersion and small size for its mass), and Andromeda XVI (dispersion value is upper limit; downward arrow). We assume that the luminosity of a star particle is proportional to its mass, and therefore R_e is both the half-mass and half-light radius for the simulated galaxies.

2.4.4 Equilibrium and spherical symmetry

The spherical Jeans equation applies to spherically symmetric systems that are in equilibrium. It is reasonable to expect that simple mass estimators based on this equation will fail for galaxies that deviate significantly from either of these assumptions.

In order to quantify the dynamical state of a galaxy in our sample, we consider the ratio of the total kinetic and gravitational potential energies of its star particles, defining

$$\gamma = \frac{2K_{\rm star}}{U_{\rm star}} \,, \tag{2.11}$$

where $K_{\rm star}$ is the total stellar kinetic energy, and $U_{\rm star}$ is the sum of the gravitational potential energy of each star particle due to the full mass distribution of the subhalo to which the star particles of the galaxy belong. In computing $U_{\rm star}$, we shift the zero-point of the potential to coincide with the galactic centre. Equation (2.11) is similar in appearance to the virial ratio for an isolated system (such that γ would equal 1 if the stars were an isolated self-gravitating system in equilibrium, with the potential taken to be zero at an infinite distance from the system). However, γ refers only to the stellar component, not to the whole dynamical system, and we measure the potential energy with respect to the bottom of the potential well. There is no natural γ value for equilibrium systems, and in the case of equilibrium the value of γ still depends on the form of the gravitational potential. The

upper panel of Fig. 2.8 shows γ versus $M_{\rm star}$ for all galaxies in our sample. We consider a galaxy to be out of equilibrium if it has a γ value substantially higher or lower than the median value for its stellar mass (black line). We define $\tilde{\gamma}$ to be the ratio of γ to the median line shown for each galaxy. Note that the trend in γ observed as a function of $M_{\rm star}$ in Fig. 2.8 does not imply a trend towards or away from equilibrium.

In order to quantify how close to spherically symmetric a galaxy is, we compute its stellar sphericity, s, as follows. We first compute the reduced inertia tensor, $\mathbb{I}$, of the star particles, which projects the mass distribution onto a unit sphere, so that the shape determination does not depend on the radial distribution of the particles (e.g. Bett 2012). This symmetric tensor has components I_{ij}, where $i, j \epsilon \{1, 2, 3\}$, such that

$$I_{ij} = \frac{1}{M_{\rm star}} \sum_n m_n \frac{r_{n,i} r_{n,j}}{r_n^2} \ , \tag{2.12}$$

where each star particle, n, has mass m_n and is located at a distance r_n from the galactic centre. $r_{n,i}$ and $r_{n,j}$ are the coordinates of the star particle with respect to the galactic centre, in the i and j directions, respectively. Taking the eigenvectors of $\mathbb{I}$ to be the principal axes of an ellipsoid, with axis lengths $a \geq b \geq c$ given by the square roots of the eigenvalues of $\mathbb{I}$, the sphericity is $s = c/a$ (a sphere has $s = 1$). The lower panel of Fig. 2.8 shows s versus stellar mass for all galaxies in our sample. Naturally, the rotation-dominated galaxies (unfilled symbols) tend be more flattened (and thus have lower s) than those that are dispersion-dominated (filled symbols). However, there are some galaxies which are dispersion-dominated according to our kinematic measure ($\kappa_{\rm rot}$) that have lower sphericity than the vast majority of rotation-dominated galaxies. Note that s does not distinguish between disc-like flattening and elongation, as this measure takes into account only two of the axes. The relationship between s and $\kappa_{\rm rot}$ has a large amount of scatter, although the scatter decreases (and s decreases) with increasing $\kappa_{\rm rot}$. Looking at Fig. 2.8, it appears that satellites are preferentially closer to spherical than field galaxies (see distributions in right panel, which consider only the dispersion-dominated galaxies). This effect is consistent with the expectation for satellite systems that have been tidally stripped (Barber et al., 2015). The sphericity also increases, on average, as stellar mass decreases for

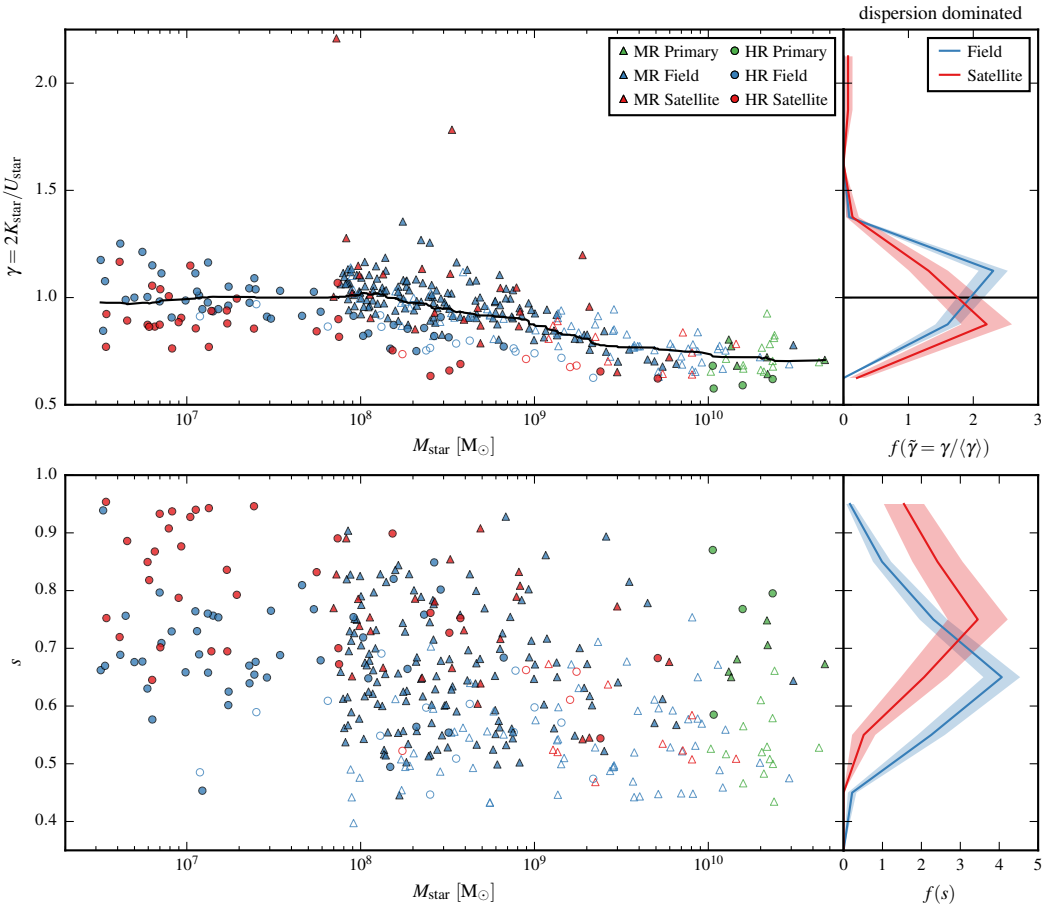

Figure 2.8 – Stellar energy ratio, γ (upper; see equation 2.11), and stellar sphericity, s (lower; derived from the reduced inertia tensor as defined in equation 2.12), versus stellar mass, $M_{\rm star}$, for all galaxies in our simulated sample. The colours show whether a galaxy is one of the MW or M31 analogues (primary), within 300 kpc of either of these (satellite), or at larger distances (field). The symbol shapes indicate the resolution level (MR or HR). Dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$) are plotted as filled symbols, while the symbols are unfilled for rotation-dominated galaxies ($\kappa_{\rm rot} > 0.5$). The black line in the upper panel shows the running median, $\langle\gamma\rangle$, of the γ values, considering at most 100 neighbouring galaxies in $M_{\rm star}$ for each galaxy. We define $\tilde{\gamma}$ as the ratio of γ to this median line for each galaxy, and assume that galaxies with $\tilde{\gamma} \approx 1$ are in equilibrium. The panels on the right show the distributions of the rescaled energy ratio, $f(\tilde{\gamma})$ (*not* the distribution of γ itself), and the sphericity, $f(s)$, on the same vertical scales as the main panels, for only the dispersion-dominated field and satellite galaxies. Each distribution is normalised to have unit area. The shaded regions show the standard deviation assuming Poisson noise on each bin count.

the data shown in Fig. 2.8.

Given the broad range of 3D shapes and the spread in energy ratios seen for the simulated galaxies in Fig. 2.8, there clearly exist galaxies within our (dispersion-dominated, or indeed, full) sample for which the assumptions of spherical symmetry and equilibrium are strongly violated. In the following, it will be particularly interesting to see how the simple mass estimators perform when applied to such galaxies.

2.5 Accuracy of mass estimators

Having demonstrated that the galaxies in our simulations have realistic projected stellar density and velocity dispersion distributions, with a range of 3D shapes, velocity dispersion anisotropies, levels of dispersion support, and departures from dynamical equilibrium, we now apply the simple mass estimators proposed by Walker et al. (2009) and Wolf et al. (2010) in order to assess their accuracy as a function of various galaxy properties.

For convenience, let us denote the ratio of the estimated mass to the true dynamical mass within some sphere as α. Specifically, for the estimator of Walker et al. (2009), from equation (2.7), we write

$$\alpha_{\rm Walker} = \frac{5\langle\sigma_{\rm los}\rangle^2 R_{\rm e}}{2GM(<R_{\rm e})}\,, \tag{2.13}$$

and, for the estimator of Wolf et al. (2010), from equation (2.9),

$$\alpha_{\rm Wolf} = \frac{4\langle\sigma_{\rm los}\rangle^2 R_{\rm e}}{GM(<4R_{\rm e}/3)}\,. \tag{2.14}$$

We emphasise that in both equations (2.13) and (2.14), the denominator is the total mass within a sphere of radius proportional to the value of $R_{\rm e}$ obtained for a given line of sight. For each galaxy in our sample, we measure $R_{\rm e}$, $\langle\sigma_{\rm los}\rangle$, $M(<R_{\rm e})$, and $M(<4R_{\rm e}/3)$ for 1536 evenly distributed lines of sight. The masses are obtained from all subhalo particles within the relevant radius, including the contributions from dark matter, gas, stars, and black holes, where such species are present.

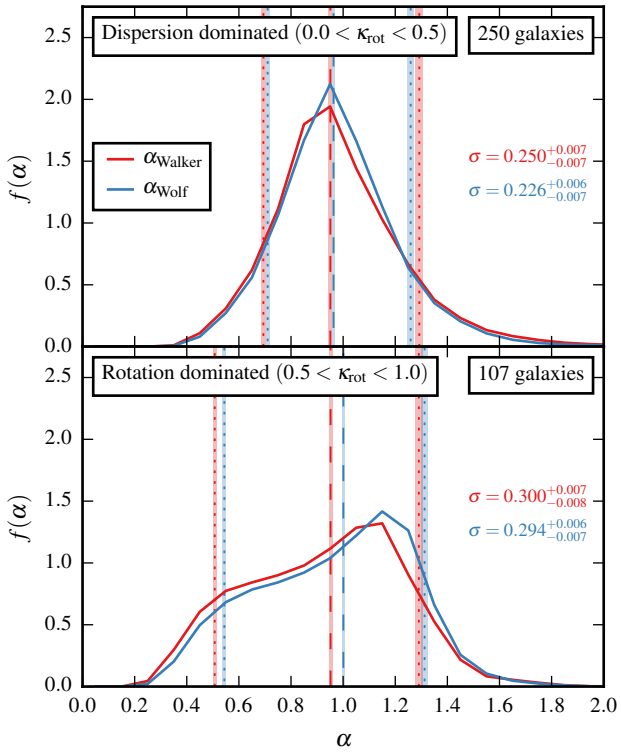

Figure 2.9 – Distributions, $f(\alpha)$, of the estimated to true mass ratio, α, for all galaxies, obtained by projecting over 1536 evenly distributed lines of sight. The galaxies are split into a dispersion-dominated sample (upper panel), and a rotation-dominated sample (lower panel), according to the value of $\kappa_{\rm rot}$ (see equation 2.10). Different colours show the distributions obtained using the Walker et al. (2009, $\alpha_{\rm Walker}$) and Wolf et al. (2010, $\alpha_{\rm Wolf}$) estimators, as labelled. The dashed lines indicate the median α values for each estimator and the dotted lines the $10^{\rm th}$ and $90^{\rm th}$ percentiles. The standard deviation, σ, of each distribution is given in the same colour as the lines. The shaded regions around the vertical lines and the quoted errors on σ are the $16^{\rm th}$ – $84^{\rm th}$ percentile confidence limits, derived from 10^4 bootstrap samples of the galaxies for each distribution. Each projection of each galaxy contributes to the relevant distribution with equal weight, and each distribution is normalised to have unit area. The panels are labelled with the number of galaxies they include. We highlight that the $f(\alpha)$ axis limits used in Figs. 2.10, 2.11, 2.12, and 2.19 vary with respect to those used here, in order to optimise the clarity of each individual figure.

2.5.1 Dispersion support

The mass estimators assume that the dynamical system is supported by dispersion. We begin by applying the estimators to all galaxies in our sample, including those with prominent stellar discs (for which we would not expect an observer to use such an estimator). Fig. 2.9 shows the distributions of α values obtained by applying both estimators to all projections of every galaxy, split according to whether the stellar motions within the galaxy are dominated by dispersion or rotation (see §2.4.1).

For both kinematic regimes, each estimator has a bias in the median of no more than 5 percent (dashed lines), accompanied by a large scatter (see dotted lines and standard deviation values). It is interesting that there is no dramatic change in the accuracy of the estimators, in the median, when switching from the dispersion-dominated to the rotation-dominated galaxies. For each sample, the two different estimators have a similar level of scatter, which is larger for the rotation-dominated galaxies. In general, the distributions are very similar in shape for the two estimators; however, the Wolf et al. (2010) estimator has a smaller median offset, 10^{th} to 90^{th} percentile spread, and standard deviation than the Walker et al. (2009) estimator, for both galaxy samples (although the difference in the scatter is much less significant for the rotation-dominated galaxies). The α distributions for the dispersion-dominated galaxies are approximately symmetric around the median, but the shape of the distributions is more complicated in the rotation-dominated case, where the peaks are above $\alpha = 1$, and there are extended tails to low α.

Overall, there is a 1σ scatter for dispersion-dominated galaxies of 25 percent when using the estimator of Walker et al. (2009), and 23 percent for the estimator of Wolf et al. (2010). In the following, we shall investigate the dependence of the α distributions for dispersion-dominated galaxies on key quantities that characterise the properties of each galaxy and are most relevant to the assumptions on which the mass estimators are based.

2.5.2 Stellar velocity dispersion anisotropy

The anisotropy, β, of the stellar velocity dispersion is an important parameter in the Jeans analysis. An important feature of the mass estimators of interest

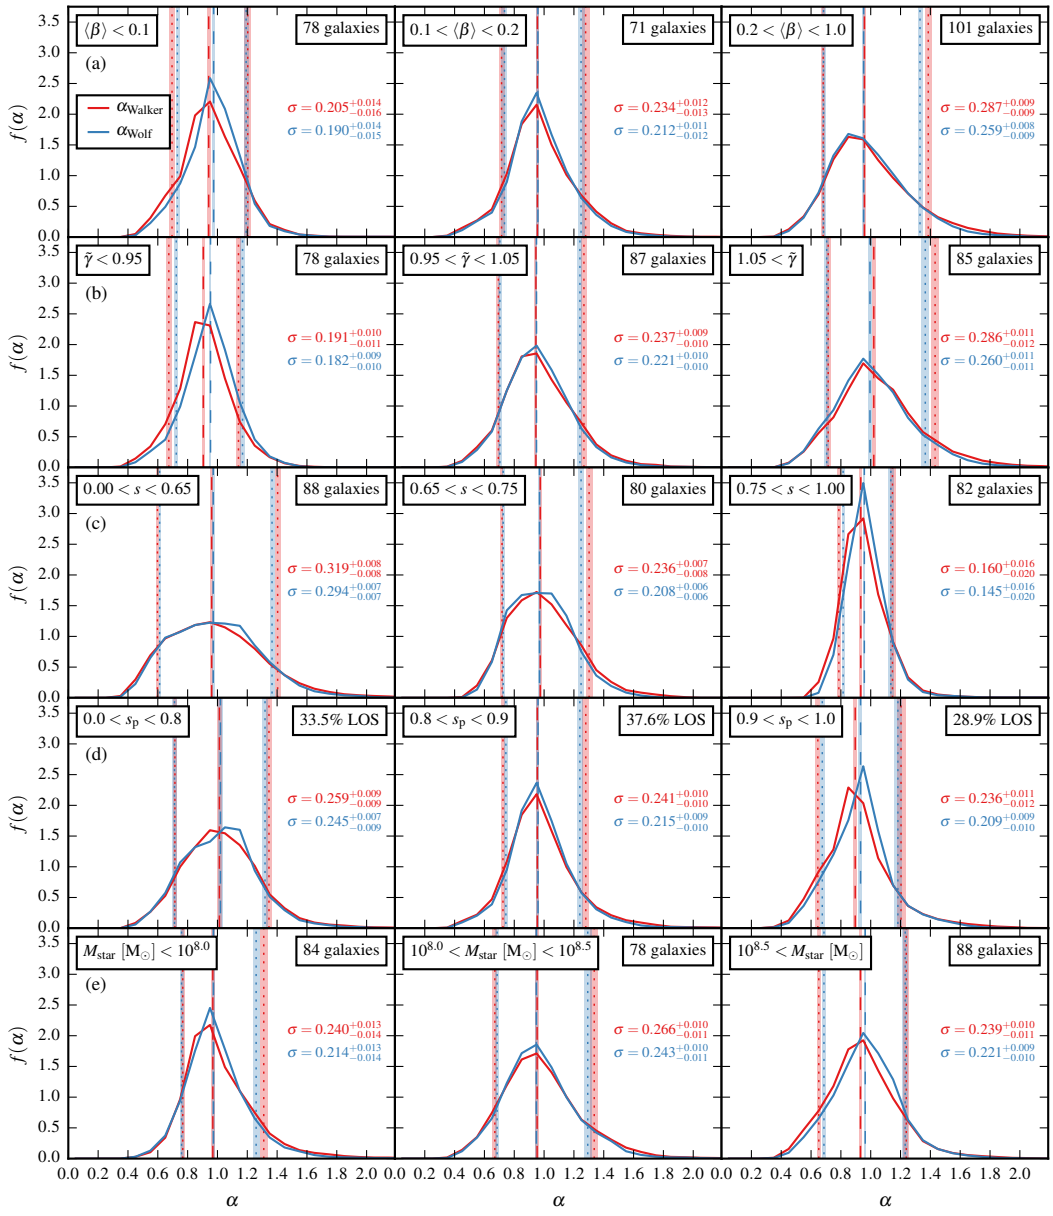

Figure 2.10 – Distribution, $f(\alpha)$, of the estimated to true mass ratio, α, for dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$), as a function of various galaxy properties. In each row, the dispersion-dominated sample is split into three bins (columns), according to: (a) the mean stellar velocity dispersion anisotropy, $\langle\beta\rangle$; (b) the stellar equilibrium measure, $\tilde{\gamma}$; (c) the stellar sphericity, s; (d) the projected stellar circularity, $s_{\rm p}$; and (e) the stellar mass, $M_{\rm star}$. Each panel is labelled with the interval it considers and the number of galaxies it contains (percentage of lines of sight in the case of circularity, where a single galaxy may contribute to more than one panel). The details of the computation of the distributions and the meaning of each line are as described in the caption of Fig. 2.9.

here is that they are designed to be minimally sensitive to β. In Fig. 2.10(a), we show the α distributions for dispersion-dominated galaxies, split into bins of $\langle\beta\rangle$, which is the stellar velocity dispersion anisotropy averaged over the whole galaxy (see equation 2.3). The stellar orbits are predominantly radially biased ($\langle\beta\rangle > 0$); there are only 46 out of 250 galaxies with $\langle\beta\rangle < 0$ (see Fig. 2.6 for the radial variation of β in the HR simulations). For each estimator, the scatter in α increases with $\langle\beta\rangle$, where the difference is greatest between the intermediate and most radially biased systems, for the $\langle\beta\rangle$ intervals shown.

2.5.3 Equilibrium

In order to assess how the accuracy of the mass estimators depends on whether the galaxy of interest is in equilibrium, as assumed in the derivation of each estimator, we make use of $\tilde{\gamma}$ as defined in §2.4.4. This quantity encodes the balance between the total stellar kinetic energy and the potential energy measured with respect to the bottom of the gravitational potential well, where galaxies with $\tilde{\gamma} \approx 1$ are assumed to be in equilibrium (the calculation of $\tilde{\gamma}$ assumes that, on average, the galaxies in our full sample are in equilibrium; see Fig. 2.8). The α distributions for dispersion-dominated galaxies divided into bins of $\tilde{\gamma}$ are shown in Fig. 2.10(b), where galaxies in the intermediate bin are taken to be close to equilibrium. The scatter in α actually increases with $\tilde{\gamma}$, such that the lowest $\tilde{\gamma}$ interval exhibits the smallest scatter for each estimator.

2.5.4 Shape

As shown in Fig. 2.8, the simulated galaxies span a broad range of 3D shapes. Fig. 2.10(c) shows the α distributions for dispersion-dominated galaxies divided into bins of stellar sphericity, s. The scatter in α diminishes sharply as the stellar mass distribution tends towards spherical symmetry ($s = 1$). This behaviour is consistent with the expectation for the applicability of the spherical Jeans equation. Clearly, the estimators perform relatively well for galaxies that are close to spherical. However, robustly identifying such systems from projected data alone is non-trivial.

We now investigate the impact of carrying out the galaxy shape determination in projection. Adapting the procedure described in §2.4.4 for computing

sphericity, we compute the reduced inertia tensor in 2D using the projected coordinates of each star particle on the sky; the resulting ellipse axis lengths, $a_p \geq b_p$, are then used to define the stellar circularity as $s_p = b_p/a_p$ (a circle has $s_p = 1$). Fig. 2.10(d) shows the α distributions in bins of circularity, where we see a much weaker trend in the scatter as a function of shape than when the full 3D information of Fig. 2.10(c) is used. As such, the scatter in α for galaxies that appear to be very close to circular on the sky ($s_p > 0.9$) is significantly larger than for galaxies that are actually close to spherical in 3D ($s > 0.75$).

2.5.5 Stellar mass

The simple mass estimators could be expected to work equally well for stellar systems on any mass scale, provided that they satisfy the spherical Jeans equation and that more detailed assumptions relating to the stellar density distribution and the stellar velocity dispersion are reasonable (see §2.2). However, many galaxy properties scale with stellar mass, which for our simulated galaxies, albeit weakly and with a large scatter, include the level of rotational support and the 3D shape (see Figs. 2.2 and 2.8). In Fig. 2.10(e) we show the α distributions for dispersion-dominated galaxies in bins of stellar mass. There is no monotonic global trend in the scatter in α as a function of stellar mass; however, for both estimators the scatter peaks in the intermediate mass interval.

2.5.6 Structural relationship with dark matter halo

As discussed above, Figs. 2.9 and 2.10 explore the dependence of the mass estimator errors on key structural and kinematical properties of the galaxy itself, i.e. properties of the set of subhalo star particles enclosed within a sphere of radius r_{gal} around the subhalo centre (including the balance between the stellar kinetic energy and the gravitational potential energy of the stars due to the mass distribution of the whole subhalo). Extending this, we now examine the dependence of the mass estimator accuracy on quantities that characterise the structural relationship between the galaxy and the dark matter component of its host subhalo.

Fig. 2.11(a) shows the α distributions for dispersion-dominated galaxies,

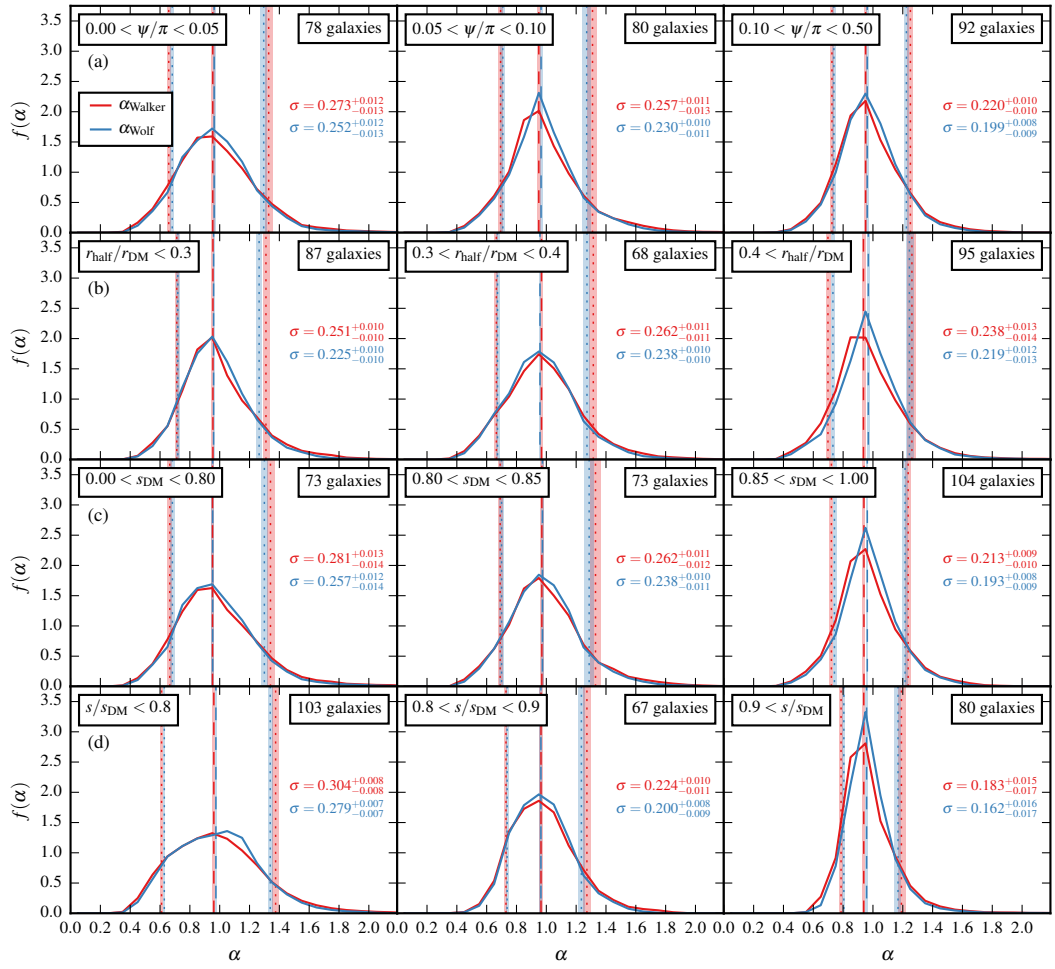

Figure 2.11 – Distribution, $f(\alpha)$, of the estimated to true mass ratio, α, for dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$), as a function of various galaxy and host subhalo dark matter properties. In each row, the dispersion-dominated galaxy sample is split into three bins (columns), according to: (a) the angle between the minor axis of the galaxy and that of the dark matter mass distribution within its host subhalo, ψ; (b) the ratio of the stellar 3D half-mass radius, $r_{\rm half}$, to the subhalo dark matter scale radius, $r_{\rm DM}$; (c) the subhalo dark matter sphericity, $s_{\rm DM}$; and (d) the ratio $s/s_{\rm DM}$, where s is the stellar sphericity. Each panel is labelled with the interval it considers and the number of galaxies it contains. The details of the computation of the distributions and the meaning of each line are as described in the caption of Fig. 2.9.

divided into bins according to the angle, ψ, between the minor axis of the galaxy and that of the dark matter component of its subhalo. For each mass component, the minor axis is taken to be given by the eigenvector of the reduced inertia tensor that has the smallest eigenvalue, and thus corresponds to the shortest axis length (i.e. c in $s = c/a$; see §2.4.4). Note that the signs of the eigenvectors are ignored, such that the angle ψ can take values from zero to $\pi/2$. The scatter in α is highest for the galaxies that are close to perfect alignment with their dark matter halo ($\psi = 0$), and the scatter decreases with increasing ψ.

In order to investigate the dependence of the estimator accuracy on how deeply the galaxy is embedded within its host subhalo, we compute a scale radius, $r_{\rm DM}$, for the subhalo dark matter mass distribution and compare this to the stellar 3D half-mass radius, $r_{\rm half}$. Given the radius at which the circular velocity curve peaks, considering the enclosed mass in subhalo dark matter particles *only*, $r_{\rm DM}$ is computed as the corresponding scale radius for an exact NFW profile (Navarro et al., 1996b).[9] Fig. 2.11(b) shows the α distributions for dispersion-dominated galaxies in bins of $r_{\rm half}/r_{\rm DM}$, which quantifies the 'concentration' of the galaxy with respect to the dark matter distribution. The scatter in α peaks in the intermediate bin shown, reducing somewhat in the most concentrated bin (lowest $r_{\rm half}/r_{\rm DM}$), followed by the least concentrated systems, but the differences are small.

Finally, it is interesting to examine the dependence of the estimator error on the 3D shape of the host dark matter halo, since the mass estimators assume that the dynamical system is spherically symmetric as a whole, and the scatter in α is found to be strongly dependent on the 3D shape of the stellar distribution (see Fig. 2.10c). We compute the sphericity of the subhalo dark matter distribution, $s_{\rm DM}$, in the same way as described in §2.4.4 for the stellar sphericity, s. Fig. 2.11(c) shows the α distributions for dispersion-dominated galaxies in bins of $s_{\rm DM}$. The scatter in α decreases with increasing $s_{\rm DM}$; however the strength of this dependence is smaller than that for the stellar sphericity, comparing to Fig. 2.10(c). This indicates that the shape of the stellar mass distribution has the largest impact on the estimator error,

[9] Given the enclosed dark matter mass profile, $M_{\rm DM}(<r)$, if the radius at which the *dark matter only* circular velocity curve, $\sqrt{GM_{\rm DM}(<r)/r}$, peaks is $r_{\rm max(DM)}$, then the scale radius for an exact NFW halo with the same value of $r_{\rm max(DM)}$ is $r_{\rm DM} \simeq r_{\rm max(DM)}/2.163$.

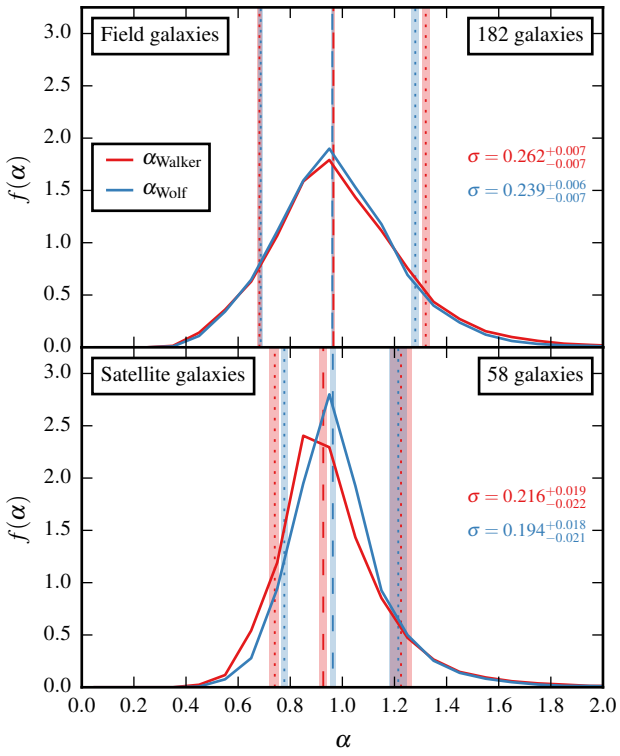

Figure 2.12 – Distributions, $f(\alpha)$, of the estimated to true mass ratio, α, for field (upper panel) and satellite (within 300 kpc of MW or M31; lower panel) galaxies in the dispersion-dominated sample ($\kappa_{\rm rot} < 0.5$). The details of the computation of the distributions and the meaning of each line are as described in the caption of Fig. 2.9.

with a weaker influence from the shape of the mass distribution of the dark matter. We note that although s does exhibit a weak increase on average as a function of $s_{\rm DM}$, there is a very large scatter in s at a given $s_{\rm DM}$, i.e. the stellar and dark matter 3D shapes are only loosely connected. The dark matter distribution is usually closer than the galaxy to being spherically symmetric (only 24 out of 250 galaxies have $s > s_{\rm DM}$). Fig. 2.11(d) shows the α distributions for dispersion-dominated galaxies in bins of $s/s_{\rm DM}$. The decrease in the scatter in α is more significant as a function of this ratio than as a function of $s_{\rm DM}$ (cf. Fig. 2.11c), but less pronounced than as a function of s alone (cf. Fig. 2.10c). Therefore, the difference between the 3D shapes of the galaxy and its host dark matter halo (quantified here by $s/s_{\rm DM}$) plays a role in the accuracy of simple mass estimators.

2.5.7 Field versus satellite galaxies

Fig. 2.12 shows the α distributions separately for dispersion-dominated field and satellite galaxies (satellites are within 300 kpc of a MW or M31 analogue). Although there are fewer satellites, and hence poorer statistics, the scatter in α is smaller for the satellite galaxies than for the field galaxies, for both estimators. This may be closely related to the fact that, on average, the satellite galaxies are closer to spherical symmetry than those in the field (see Fig. 2.8). There are hints that the satellites also tend to be more strongly supported by dispersion (see Fig. 2.2) and tend to have stellar velocity dispersions that are closer to isotropic (see Fig. 2.6 for HR), compared with field galaxies. We also find that the satellite galaxies tend to be more extended relative to the dark matter distribution of their subhalo (higher $r_{\rm half}/r_{\rm DM}$) than the field galaxies.

2.5.8 Angular dependence

The analysis presented so far has considered the estimator error distributions resulting from summing over all galaxies (or lines of sight) satisfying certain criteria based on integrated stellar and dark matter properties (or per-projection observable properties). We now investigate how the estimator accuracy varies with viewing angle relative to the orientation of the galaxy. Given Cartesian coordinates (x, y, z), we adopt the standard convention for the spherical polar coordinates (r, θ, ϕ), such that $x = r\sin(\theta)\cos(\phi)$, $y = r\sin(\theta)\sin(\phi)$, and $z = r\cos(\theta)$. For each galaxy, we rotate the simulation coordinate system to align with the eigenvectors of the stellar reduced inertia tensor (see §2.4.4). The z axis ($\theta = 0$) is aligned with the shortest principal axis (minor axis), the x axis ($\theta = \pi/2, \phi = 0$) with the longest principal axis (major axis), and the remaining principal axis with the y direction ($\theta = \pi/2, \phi = \pi/2$).

The upper panels of Fig. 2.13 show the variation in the α values for each estimator, relative to the mean α value from averaging over all lines of sight for a given galaxy, as a function of θ and ϕ, defined using the *stellar* reduced inertia tensor. The angular variations of the projected half-mass radius, $R_{\rm e}$, and the line-of-sight integrated stellar velocity dispersion, $\langle\sigma_{\rm los}\rangle$, are also shown, again relative to the corresponding galactic means. The lower

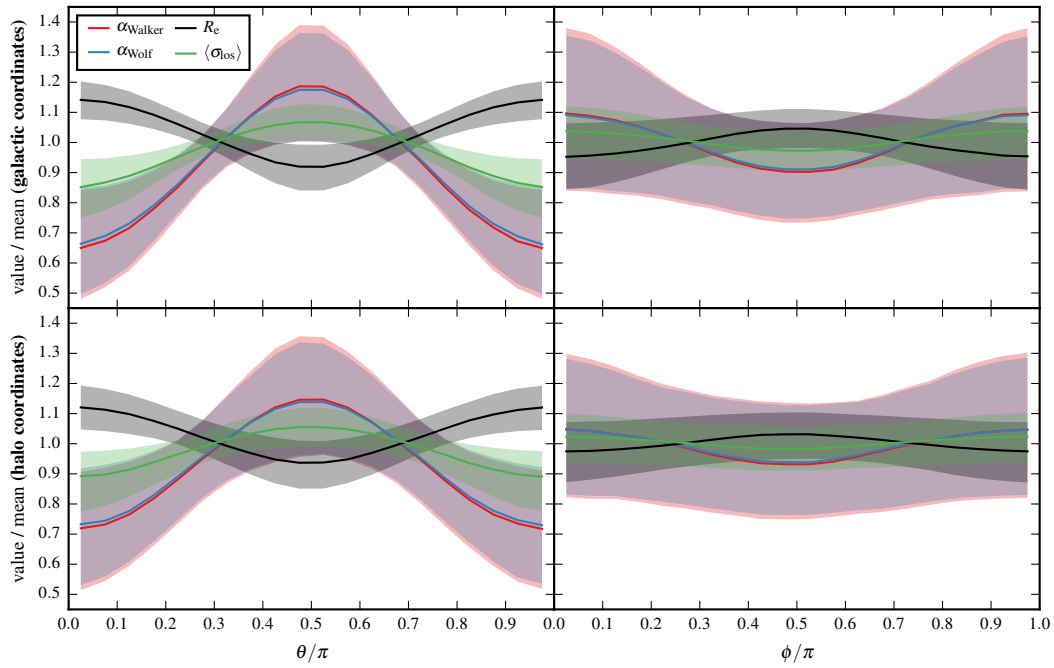

Figure 2.13 – Angular variation of the estimated to true mass ratio for the Walker et al. (2009, $\alpha_{\rm Walker}$) and Wolf et al. (2010, $\alpha_{\rm Wolf}$) estimators, the projected stellar half-mass radius, $R_{\rm e}$, and the line-of-sight stellar velocity dispersion, $\langle\sigma_{\rm los}\rangle$, for dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$), computed by projecting over 1536 evenly distributed lines of sight. The curves are shown as functions of the spherical polar angles, θ and ϕ, where the coordinate system is aligned with the eigenbasis of the reduced inertia tensor of either the galactic star particles (upper panels) or the dark matter particles of the host subhalo (lower panels). The z axis ($\theta = 0$) is aligned with the shortest principal axis, and the x axis ($\theta = \pi/2, \phi = 0$) is aligned with the longest principal axis. The spherical polar angular ranges are defined as $0 \leq \theta \leq \pi$ and $0 \leq \phi < 2\pi$. Since projections in opposite directions are equivalent in this work, the 1536 unique lines of sight are identified by their angular coordinates in the half-sphere defined by $0 \leq \theta < \pi$ and $0 \leq \phi < \pi$ within this figure. For each galaxy, the values of $\alpha_{\rm Walker}$, $\alpha_{\rm Wolf}$, $R_{\rm e}$, and $\langle\sigma_{\rm los}\rangle$ for each projection have been divided by the corresponding mean value over all projections of that galaxy, before computing the curves shown. The solid lines show the median values, and the shaded regions of the same colour indicate the $16^{\rm th}$ – $84^{\rm th}$ percentile spread (the θ and ϕ axes use the same fixed bin width, in units of π). Each panel makes use of the full set of 1536 lines of sight, so at fixed θ (left) or ϕ (right), the median and percentile range values shown result from the data for all lines of sight (with that particular θ or ϕ) for all galaxies.

panels in Fig. 2.13 repeat the same analysis, but using instead the θ and ϕ coordinates with respect to the eigenbasis of the reduced inertia tensor of the *dark matter* particles within the subhalo. Thus, for a given galaxy, the (θ, ϕ) coordinates contributing to the upper and lower panels of Fig. 2.13 differ unless the stellar and dark matter reduced inertia tensor eigenvectors are exactly aligned, with the same eigenvalue ordering (since this sets the ordering of the principal axis lengths, and hence the definitions of θ and ϕ).

Looking at the upper panels of Fig. 2.13, it is clear that α is strongly sensitive to θ, i.e. the angle with respect to the stellar minor axis (z), and has a weaker dependence on ϕ, i.e. the angle with respect to the major axis (x; in the $x-y$ plane). The results for the two different estimators are very similar, although the estimator of Wolf et al. (2010) exhibits a slightly smaller scatter at all angles. As a function of θ (upper left panel), the mass estimates are maximally biased low (in the median for all galaxies), relative to the galactic mean values, for lines of sight coincident with the stellar minor axis ($\theta \sim 0$ and $\theta \sim \pi$). The median estimates then smoothly increase towards the $x-y$ plane, peaking for $\theta/\pi \sim 0.5$. The scatter in α, relative to the galactic mean, is relatively large for lines of sight along the minor axis or within the $x-y$ plane, but is minimised for $\theta/\pi \sim 0.3$ and $\theta/\pi \sim 0.7$, being approximately symmetrical around $\theta/\pi = 0.5$. The scatter in both $R_{\rm e}$ and $\langle\sigma_{\rm los}\rangle$ exhibits the same behaviour as a function of θ as the scatter in α, with extremes at approximately the same values of θ. The velocity dispersion varies with θ in the same way as α does, peaking for $\theta/\pi \sim 0.5$, with minima coinciding with the minor axis. However, the dependence of $R_{\rm e}$ on θ is reversed, compared to that of α and $\langle\sigma_{\rm los}\rangle$, such that the $R_{\rm e}$ curve is approximately the reflection of the $\langle\sigma_{\rm los}\rangle$ curve around the galactic mean. The coupled variation of $R_{\rm e}$ and $\langle\sigma_{\rm los}\rangle$ suppresses the variation in the product $R_{\rm e}\langle\sigma_{\rm los}\rangle^2$, and hence the variation in the estimated mass (see equations 2.13 and 2.14). The α, $R_{\rm e}$, and $\langle\sigma_{\rm los}\rangle$ curves intersect (with values ~ 1) for θ close to where the scatter in each quantity is minimised ($\theta/\pi \sim 0.3$ and $\theta/\pi \sim 0.7$).

The results as a function of ϕ in galactic coordinates (upper right panel in Fig. 2.13) are similar to those just described for the θ dependence, but with smaller median offsets from unity and larger scatters, exhibiting extremes such that α and $\langle\sigma_{\rm los}\rangle$ are maximised for $\phi \sim 0$ and $\phi \sim \pi$ (i.e. for lines of sight in the $x-z$ plane), and minimised for $\phi/\pi \sim 0.5$ (i.e. the $y-z$ plane); while

the ϕ values of the turning points are shifted by an angle of $\pi/2$ relative to this for R_e. Note that since projections in opposite directions are equivalent for our purposes, the angular coordinates of the 1536 unique lines of sight are evaluated within the half-sphere defined by $0 \leq \theta < \pi$ and $0 \leq \phi < \pi$ in Fig. 2.13.

Using instead the coordinate system based on the subhalo dark matter, as shown in the lower panels of Fig. 2.13, results in dependencies of α, R_e, and $\langle \sigma_{los} \rangle$ on the spherical polar angles that are qualitatively very similar to the data shown in the upper panels, where the coordinates based on the galactic star particles are used. However, the strengths of the trends seen in galactic coordinates are diluted when switching to the halo-based coordinates, such that the median variations are reduced for each quantity shown (i.e. the solid lines are slightly closer to unity). The scatter around the median also increases for each quantity at all angles when switching to the halo coordinates, except for $\phi/\pi \lesssim 0.2$ and $\phi/\pi \gtrsim 0.8$, where for each curve shown, the $16^{th} - 84^{th}$ percentile width actually decreases. The differences between the upper and lower panels in Fig. 2.13 indicate that the accuracy of the mass estimators, for a given line of sight, is more sensitive to the shape and relative alignment of the galaxy, rather than the corresponding properties of the dark matter (sub)halo in which the galaxy is embedded.

Fig. 2.14 shows the mean α value for the Walker et al. (2009) estimator, as a function of $\cos(\theta)$ and ϕ, for all dispersion-dominated galaxies, using the coordinate system based on the reduced inertia tensor of the stars within each galaxy, as in the upper panels of Fig. 2.13 (but note that Fig. 2.13 considers the α values relative to the galactic mean, for each galaxy, while in Fig. 2.14 the actual individual α values are used directly to compute the mean for each pixel). Since the lines of sight are drawn from an evenly spaced spherical tessellation, they must evenly sample the $\cos(\theta)$ versus ϕ plane, and so the pixels in Fig. 2.14 each consider approximately the same number of data points (galaxy projections). The results for the Wolf et al. (2010) estimator are very similar to those shown in Fig. 2.14. As can be seen in Fig. 2.14, the most accurate mean mass estimates form an hourglass shape in the $\cos(\theta) - \phi$ plane (white). The masses are most heavily underestimated for $\cos(\theta) \sim \pm 1$ (blue; corresponding to lines of sight along the stellar *minor* axis), and most heavily overestimated for $\cos(\theta) \sim 0$ with $\phi/\pi \sim 0$ or $\phi/\pi \sim 1$

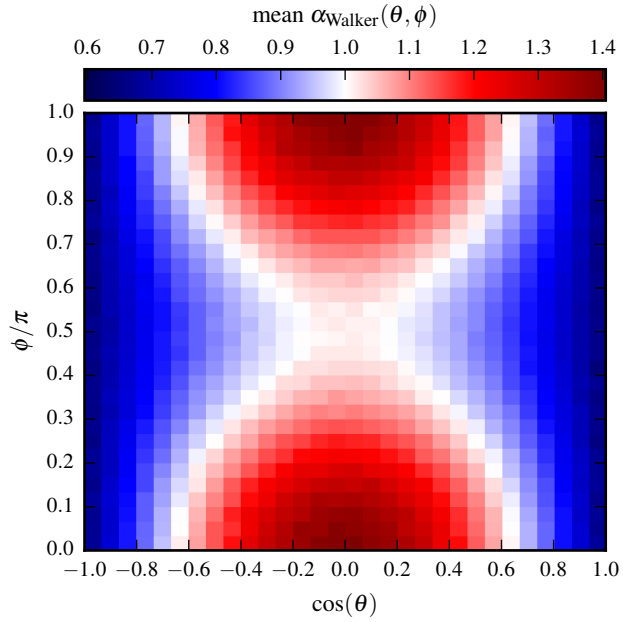

Figure 2.14 – Estimated to true mass ratio, $\alpha_{\rm Walker}(\theta,\phi)$, resulting from applying the estimator of Walker et al. (2009) to all galaxies in our dispersion-dominated sample ($\kappa_{\rm rot} < 0.5$), by projecting over 1536 evenly distributed lines of sight for each galaxy. θ and ϕ are the spherical polar angles in the galactic coordinate system defined by the eigenbasis of the stellar reduced inertia tensor, where the z axis ($\theta = 0$) is aligned with the shortest stellar principal axis, and the x axis ($\theta = \pi/2, \phi = 0$) is aligned with the longest principal axis. The angular ranges are defined as $0 \leq \theta \leq \pi$ and $0 \leq \phi < 2\pi$; however, since projections in opposite directions are equivalent in the context of our analysis, we consider the coordinates of the set of 1536 unique lines of sight within the half-sphere defined by $0 \leq \theta < \pi$ and $0 \leq \phi < \pi$ in the grid shown here (extending the grid to $\phi = 2\pi$ would result in repetition of the data). Since the lines of sight result from an evenly distributed spherical tessellation, plotting ϕ versus $\cos(\theta) = z/r$ ensures that the number of projections per grid pixel is approximately constant. The $\alpha_{\rm Walker}(\theta,\phi)$ value shown in each pixel is the mean over all projections of each galaxy within the relevant angular range. The colour scale is centred on a mass ratio of unity (i.e. white for accurate mean estimates, blue for mean underestimates, and red for mean overestimates). The results for the Wolf et al. (2010) estimator are very similar to those shown here. The distributions of the mass ratios, α, for both estimators are shown in the upper panel of Fig. 2.9 (summing over all projections), and the α variations relative to the galactic mean are shown in the upper panels of Fig. 2.13 (integrating over θ and ϕ separately).

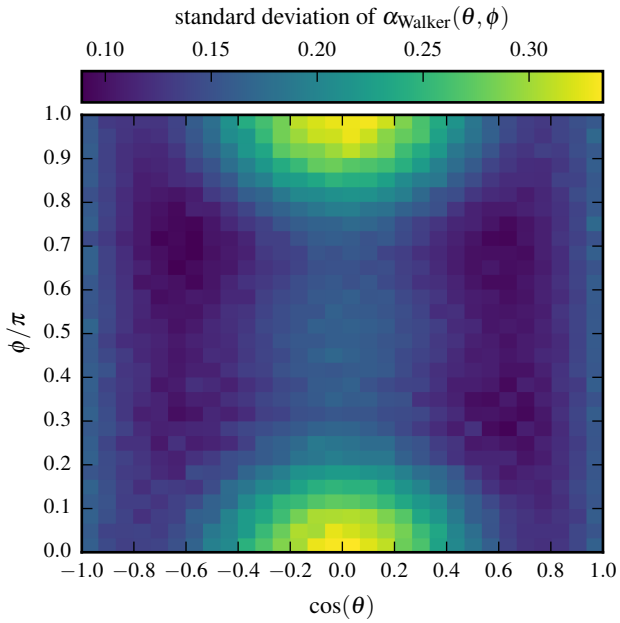

Figure 2.15 – The same as Fig. 2.14, but showing instead the standard deviation of the $\alpha_{\rm Walker}(\theta,\phi)$ values in each pixel, rather than the mean value. The details of the analysis are otherwise exactly as described in Fig. 2.14.

(red; corresponding to lines of sight along the stellar *major* axis). Lines of sight coincident with the stellar principal axis of *intermediate* length, where $\cos(\theta) = 0$ and $\phi/\pi = 0.5$ (or $\phi/\pi = 1.5$; not shown), correspond to the centre of the grid shown in Fig. 2.14, where the mean mass estimates are relatively close to perfect accuracy.

Fig. 2.15 is the same as Fig. 2.14, except that it shows the standard deviation of the α values, as a function of $\cos(\theta)$ and ϕ; otherwise the calculations involved are identical for the two figures. It can be seen in Fig. 2.15 that the scatter in α is highest for lines of sight along the stellar major axis, and lowest for the two regions with $\cos(\theta) \sim \pm 0.6$ (i.e. $\theta/\pi \sim 0.3$ and $\theta/\pi \sim 0.7$) and $0.2 \lesssim \phi/\pi \lesssim 0.8$ (in keeping with the upper panels of Fig. 2.13). The scatter for lines of sight coincident with the stellar minor axis and the intermediate stellar principal axis is fairly similar, in stark contrast to the major axis. Comparing Figs. 2.14 and 2.15, we see that the locations of the largest scatter in α are also the locations of the highest mean α values. For the three stellar principal axes, the dynamical masses are either overestimated with a large scatter (major axis), underestimated with a modest scatter (minor axis), or in fact close to accurate with a modest scatter

(intermediate axis).

2.5.9 Summary and discussion

To summarise the preceding results, Fig. 2.16 shows α for each estimator as a function of various galaxy properties, in bins of approximately equal numbers of galaxies (or sight-lines). The first panel shows the dependence on $\kappa_{\rm rot}$ for the full galaxy sample, and the subsequent panels consider only the dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$). The distinction between the two estimators is small, although where they deviate, the Wolf et al. (2010) estimator has the smallest bias and scatter in almost all cases.

Considering the scatter, the estimators are most precise for dispersion supported galaxies that are close to spherical, have at most mildly radially biased stellar velocity dispersions ($\langle\beta\rangle \lesssim 0.2$), and have relatively low stellar kinetic energy for their gravitational potential energy. The strongest dependence of the scatter is on the 3D shape of the stellar mass distribution, of which almost all useful information is lost in projection. Considering only the dispersion-dominated galaxies, as shown in Fig. 2.16, the scatter also reduces for galaxies that are misaligned with respect to their host dark matter halo, and for galaxies that are highly extended, in units of their halo scale radius. It is interesting to note that the (dispersion-dominated) satellite galaxies exhibit a lower α scatter overall compared to their counterparts in the field (see Fig. 2.12). Satellite galaxies also tend to be closer to spherical and tend to be more extended relative to their host dark matter halo, compared to the field galaxies (the highest $r_{\rm half}/r_{\rm DM}$ bin in Fig. 2.16 contains 21 satellite galaxies and only four field galaxies). We note that the main physical galaxy properties of interest are related at some level: galaxies that are close to spherical tend to have velocity dispersions that are close to isotropic and have relatively little rotational support. The scatter in α is much more sensitive to the shape of the galaxy than to that of the mass distribution of the dark matter within its host subhalo, such that there remain strong trends in the scatter in α as a function of the ratio of stellar and dark matter sphericities, $s/s_{\rm DM}$.

Any trends in the median α values shown in Fig. 2.16 are weak. However, considering the dependence on the projected stellar circularity, $s_{\rm p}$, the bias

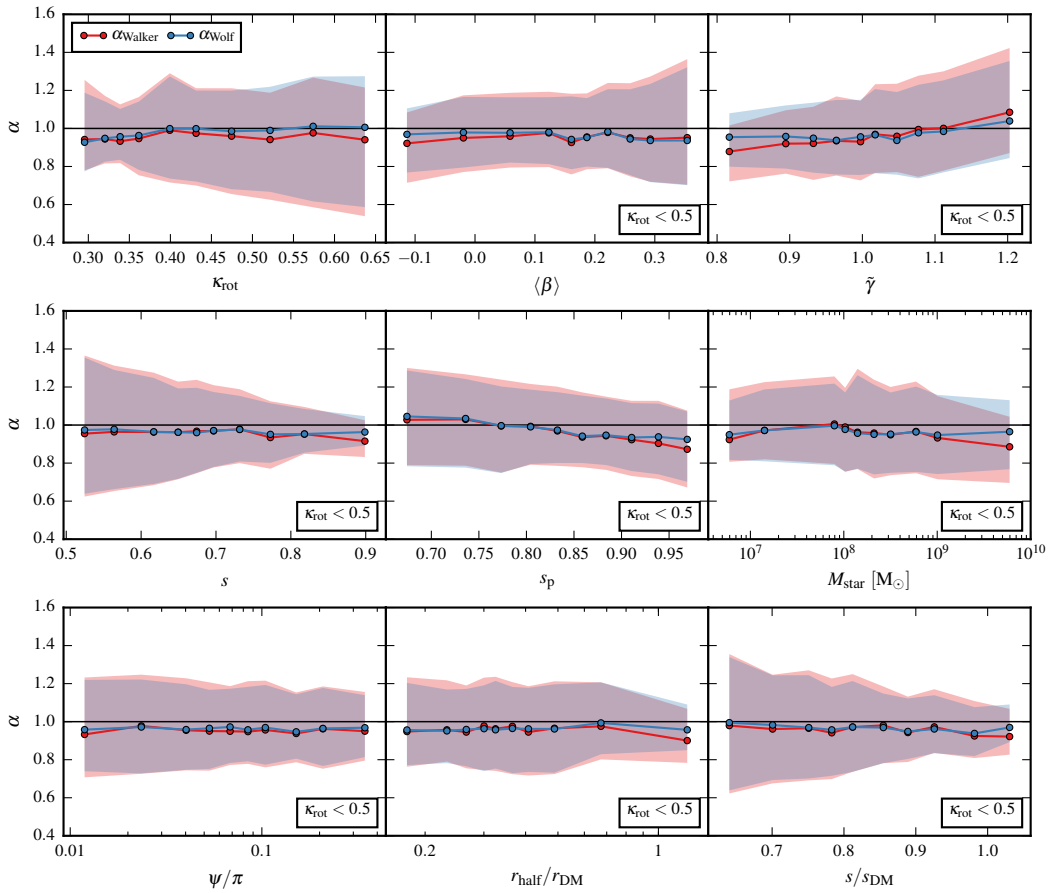

Figure 2.16 – Ratio, α, of the estimated to true mass obtained using the estimators of Walker et al. (2009, $\alpha_{\rm Walker}$) and Wolf et al. (2010, $\alpha_{\rm Wolf}$), by projecting over 1536 evenly distributed lines of sight, as a function of various galaxy and subhalo dark matter properties. The properties shown are: the stellar kinematic measure used to discriminate between galaxies with different levels of rotational support ($\kappa_{\rm rot}$), the mean stellar velocity dispersion anisotropy ($\langle\beta\rangle$), the stellar equilibrium measure ($\tilde{\gamma}$), the stellar sphericity (s), the projected stellar circularity ($s_{\rm p}$), the stellar mass ($M_{\rm star}$), the angle between the stellar and subhalo dark matter minor axes (ψ), the ratio of the stellar and dark matter scale radii ($r_{\rm half}/r_{\rm DM}$), and the ratio of the stellar and dark matter sphericities ($s/s_{\rm DM}$). The first panel shows α versus $\kappa_{\rm rot}$ for all galaxies; all other panels consider only those galaxies whose stellar motions are dominated by dispersion ($\kappa_{\rm rot} < 0.5$). For each estimator (colours), the solid lines show the median mass ratios and the shaded regions the $16^{\rm th}$ – $84^{\rm th}$ percentile spread. The bin edges are evenly spaced percentiles of the quantity on the horizontal axis, and the bin centres (points) are the median values in each bin.

towards overestimation of the mass for galaxies that are very elongated on the sky gradually decreases with increasing circularity (we note that this trend is exaggerated if we instead consider the full galaxy sample, including the rotation-dominated galaxies). The median α also exhibits an increase with increasing $\tilde{\gamma}$ (i.e. for increasing stellar kinetic energy relative to the stellar gravitational potential energy, at fixed stellar mass).

As shown in Fig. 2.13 (upper panels), the value of α is strongly sensitive to the viewing angle, θ, relative to the stellar minor axis, z, such that the estimated masses are lowest for lines of sight close to the minor axis, and highest for lines of sight in the $x-y$ plane. The dependence of α on the other spherical polar viewing angle, ϕ, relative to the stellar major axis, x (in the $x-y$ plane), is much weaker than the dependence on θ, but shows similar behaviour. As shown in Figs. 2.14 and 2.15, the dynamical masses are overestimated with a large scatter for lines of sight coincident with the major axis, underestimated with a modest scatter for the minor axis, and close to accurate with a modest scatter for the intermediate principal axis.

We note that Wolf et al. (2010) suggest an error of around 0.05 dex (12 percent; cf. Fig. 2.16) for their estimator (in addition to the uncertainty due to the errors on the measured size and velocity dispersion), in the special case where the unknown $\beta(r)$ has an extremum within the stellar distribution, or $\sigma_{\rm los}(R)$ has not been measured over the full extent of the galaxy (i.e. the dispersion much beyond $R \sim R_{\rm e}$ is not included when $\langle\sigma_{\rm los}\rangle$ is computed). Otherwise, Wolf et al. (2010) advocate considering only the measurement errors on the velocity dispersion and half-light radius when evaluating the uncertainties in masses obtained using their estimator.

Kowalczyk et al. (2013) apply a modified version of the estimator of Wolf et al. (2010) to idealised simulations of dSphs in a static MW potential, and find that the error decreases as the shape of the stellar mass distribution tends towards spherical symmetry, as also found here. However, they also find that for the simulated dSphs there is no clear dependence of the error on the level of stellar rotation or velocity dispersion anisotropy, while we do detect such dependencies here (see Fig. 2.16).

Laporte et al. (2013a) investigate the accuracy of the Walker et al. (2009) estimator for idealised stellar distributions placed in dark matter subhaloes from the Aquarius simulations, finding fluctuations of between 10 and 20

percent with respect to the true mass. An interesting result of this study, which naturally takes into account the triaxiality of dark matter haloes formed in cosmological simulations, is that it finds that the enclosed dynamical mass is more strongly overestimated for stellar populations that are more deeply embedded (more highly concentrated) within their dark matter halo, concurring with the spherically symmetric tests of Walker & Peñarrubia (2011). However, we do not detect significant evidence of such a trend in the APOSTLE simulations (see the $r_{\rm half}/r_{\rm DM}$ panel of Fig. 2.16). Additionally, Laporte et al. (2013a) demonstrate variations of the estimated mass, $R_{\rm e}$, and $\langle\sigma_{\rm los}\rangle$, as a function of viewing angle, that are in good qualitative agreement with the results presented in §2.5.8.

Lyskova et al. (2015) apply the Wolf et al. (2010) estimator to cosmological zoom simulations of 40 isolated elliptical galaxies, finding a 1σ scatter of around 8 percent for $V_{\rm c}(4R_{\rm e}/3)$ (note that the fractional error on $V_{\rm c}(r)$ is half that on $M(<r)$, assuming zero error on the radius), with a bias of 3 percent above the true circular velocity (cf. Fig. 2.16).

Our results complement these earlier studies, but we have used a large sample of realistic galaxies with a broad range of stellar masses, formed in a self-consistent ΛCDM cosmological context, in the environment of the Local Group, using sophisticated treatments of the baryonic physics important for galaxy formation.

2.6 Effects of uncertainties

We now explore the impact of the systematic uncertainties associated with the simple mass estimators on key results in the literature that are based on applying the estimators to observational data. Despite the high resolution of the APOSTLE simulations, our sample of galaxies does not include systems with stellar masses smaller than that of Sculptor (due to our chosen resolution threshold; see §2.3.3 and Fig. 2.7). However, the estimator error has no clear dependence on stellar mass (see Fig. 2.10e). Therefore, we assume that the uncertainties explored in §2.5 remain relevant for galaxies that are fainter than those in our sample.

2.6.1 Dynamical masses of MW dSphs

Wolf et al. (2010) find that simply propagating the observational errors on the half-light radius and the line-of-sight velocity dispersion through their estimator equation yields a similar uncertainty for the estimated dynamical mass to that obtained from their full Jeans analysis (which, like the estimator, assumes spherical symmetry). The authors argue that this is consistent with their claim that the uncertainty on the recovered mass is dominated by observational errors, rather than underlying systematic effects. However, as we have demonstrated, the estimator's precision is sensitive to various properties of the target galaxy.

To illustrate the significance of the systematic error on the estimated mass, Table 2.2 lists the mass within the 3D half-light radius for dSphs of the MW, which we compute using equation (2.9), given the projected half-light radius and line-of-sight velocity dispersion values tabulated by Wolf et al. (2010).[10] For each galaxy, we list the 1σ fractional error, σ_M, on the estimated mass, obtained by propagating the observational errors through equation (2.9), as advocated by Wolf et al. (2010). The final column gives the fractional error on the mass if, in addition, we add in quadrature a *representative* 1σ systematic error of 20 percent in the estimated to true mass ratio. This fiducial value for the scatter is appropriate on average for the dispersion-dominated satellite galaxies in the simulations (see Fig. 2.12), or all dispersion-dominated galaxies that appear near-circular on the sky (see Fig. 2.10d). For many of the bright satellites, the observational errors on the measured quantities are sufficiently small that the increase in the overall mass error due to including the representative systematic uncertainty is similar in size to the original mass error. However, in many other cases the original error is so large that the contribution from the systematic uncertainty is relatively insignificant.

2.6.2 Density profiles of Sculptor and Fornax

Walker & Peñarrubia (2011) employ a Markov Chain Monte Carlo (MCMC) technique to infer projected half-light radii and velocity dispersions for

[10] Note that Wolf et al. (2010) tabulate the mass within the 3D half-light radius from their full Jeans analysis, rather than from their estimator.

Table 2.2 – Dynamical mass, $M_{\rm Wolf}$, within the 3D half-light radius ($r_{\rm half} \approx 4R_{\rm e}/3$) of dSph satellites of the MW, according to the estimator of Wolf et al. (2010), as given in equation (2.9). The estimates have been computed using the projected half-light radii, $R_{\rm e}$ (as reproduced here), and line-of-sight velocity dispersions tabulated by Wolf et al. (2010). σ_M is the 1σ fractional error on $M_{\rm Wolf}$, obtained by propagating the observational uncertainties on the size and dispersion measurements through equation (2.9), approximating the errors as symmetric. To illustrate the impact of the systematic errors explored in this chapter, the final column gives the fractional error, $\tilde{\sigma}_M = (\sigma_M^2 + \sigma_{\rm sys}^2)^{1/2}$, on the estimated mass if a *representative* systematic error of $\sigma_{\rm sys} = 0.2$ is assumed for the estimated to true mass ratio (ignoring any bias; see Fig. 2.12). The galaxies are listed in order of decreasing $M_{\rm Wolf}$.

Galaxy	$R_{\rm e}$ [pc]	$\log_{10}(M_{\rm Wolf}\ [{\rm M}_\odot])$	σ_M	$\tilde{\sigma}_M$
Fornax	714^{+40}_{-40}	7.88	0.07	0.21
Ursa Minor	445^{+44}_{-44}	7.74	0.14	0.25
Sextans	768^{+47}_{-47}	7.56	0.10	0.23
Canes Venatici I	564^{+36}_{-36}	7.48	0.15	0.25
Leo I	295^{+49}_{-49}	7.35	0.19	0.27
Sculptor	282^{+41}_{-41}	7.33	0.15	0.25
Draco	220^{+11}_{-11}	7.32	0.11	0.23
Bootes I	242^{+22}_{-20}	7.26	0.50	0.54
Ursa Major I	318^{+50}_{-39}	7.23	0.30	0.36
Carina	254^{+28}_{-28}	6.99	0.13	0.24
Leo II	177^{+13}_{-13}	6.86	0.17	0.26
Leo T	115^{+17}_{-17}	6.81	0.44	0.48
Ursa Major II	140^{+25}_{-25}	6.77	0.45	0.50
Hercules	229^{+19}_{-19}	6.74	0.36	0.41
Coma Berenices	77^{+10}_{-10}	6.18	0.37	0.42
Canes Venatici II	74^{+14}_{-10}	6.16	0.46	0.51
Leo IV	116^{+26}_{-34}	6.07	1.06	1.08
Segue 1	29^{+8}_{-5}	5.70	0.56	0.59
Willman 1	25^{+5}_{-6}	5.57	0.50	0.54

chemo-dynamically distinct stellar subpopulations in the Sculptor and Fornax dSphs. In each galaxy there is a centrally concentrated, kinematically cold, metal-rich stellar component (population 1), and a more extended, metal-poor component, with higher velocity dispersion (population 2). The procedure used by Walker & Peñarrubia (2011) treats each galaxy as a superposition of two spherically symmetric stellar populations (each described by a Plummer profile), with Gaussian velocity and metallicity distributions. Assuming that the two populations independently trace the gravitational potential, the resulting size and dispersion values define a mass slope, $\Gamma = \Delta \log M(< r)/\Delta \log r$, using equation (2.1), where the actual values of λ and μ are not relevant for computing the slope:

$$\Gamma = \frac{\log[M(< r_2)/M(< r_1)]}{\log[r_2/r_1]} \approx 1 + \frac{\log[\sigma_2^2/\sigma_1^2]}{\log[r_2/r_1]} \,. \tag{2.15}$$

Here r_1 and r_2 are 3D radii equal to (or some multiple of) the projected half-light radii, R_e, of populations 1 and 2, and σ_1 and σ_2 are the mean velocity dispersions, $\langle \sigma_{los} \rangle$, of each population, respectively. Building a posterior probability distribution function, $P(\Gamma)$, using the set of r_1, r_2, σ_1, and σ_2 values at each point in the MCMC chains, Walker & Peñarrubia (2011) find median slopes of $\Gamma = 2.61^{+0.43}_{-0.37}$ for Fornax and $\Gamma = 2.95^{+0.51}_{-0.39}$ for Sculptor, where the ranges indicate the $16^{th} - 84^{th}$ percentile confidence intervals.

A measurement of Γ using the masses enclosed at two non-zero radii places an upper limit on the inner logarithmic slope, γ_{DM}, of the dark matter density profile, such that $\gamma_{DM} < 3 - \Gamma$. Given that $\gamma_{DM} = 1$ for an NFW density profile (Navarro, Frenk & White, 1996b, 1997), while e.g. $\gamma_{DM} = 0$ for a constant density core, Walker & Peñarrubia (2011) exclude NFW (or steeper) profiles with confidence of (at least) 95.9 and 99.8 percent for Fornax and Sculptor respectively. These significance values, $s(\gamma_{DM})$, are

$$s(\gamma_{DM}) = \frac{\int_{3-\gamma_{DM}}^{\infty} P(\Gamma)\, d\Gamma}{\int_{-\infty}^{\infty} P(\Gamma)\, d\Gamma} \,. \tag{2.16}$$

If the estimates of the enclosed dynamical masses for the two stellar populations each happen to be biased by exactly the same (arbitrary) factor, with respect to the corresponding true masses, then clearly such a bias

cancels out in the calculation of Γ, i.e. the recovered mass profile slope is insensitive to any coherent bias in the estimates of the enclosed masses at the two radii (see equation 2.15). Therefore, only *differences* in the enclosed mass estimation biases for the two populations are relevant in the context of the methodology of Walker & Peñarrubia (2011).

Laporte et al. (2013a) generate stellar distribution functions in dark matter subhaloes drawn from the Aquarius simulations, in order to assess the sensitivity of the accuracy of the dual population method of Walker & Peñarrubia (2011) to the triaxiality of the gravitational potential (i.e. the lack of spherical symmetry in the dark matter distribution). The lack of symmetry tends to introduce an anticorrelation between $R_{\rm e}$ and $\langle\sigma_{\rm los}\rangle$ that suppresses the error on the recovered mass (see also Fig. 2.13 and §B.2). Laporte et al. (2013a) argue that the level of spherical symmetry is not important to the estimator accuracy, in the content of the analysis of Walker & Peñarrubia (2011), despite the fact that they obtain fluctuations of between 10 and 20 percent for the estimated masses (cf. Fig. 2.16). These authors find, as in the spherically symmetric tests of Walker & Peñarrubia (2011), that Γ tends to be systematically *underestimated*, and thus the exclusion confidences for NFW profiles in Sculptor and Fornax are deemed to be *conservative* (i.e. the results for Γ can be considered as reliable lower limits, corresponding to reliable upper limits on the density slope, since $\gamma_{\rm DM} < 3 - \Gamma$). This is because the enclosed mass tends to be more strongly overestimated for tracer populations that are more deeply embedded. However, as noted in §2.5.9, we do not find any significant evidence of such a trend in the bias in the estimated to true mass ratio (α) as a function of the concentration of the stellar distribution ($r_{\rm half}/r_{\rm DM}$) for the dispersion-dominated galaxies within our simulations (see Fig. 2.16).

An important aspect of the tests of Laporte et al. (2013a), which use idealised independent stellar distributions placed in realistic dark matter haloes, is that the shapes (equidensity surfaces) of the stellar populations exactly trace the gravitational potential, and so the shapes of any two stellar populations generated in this way are inherently closely correlated, by construction (Laporte et al., 2013b). This close coupling assumed between the two populations may produce artificially small errors on the recovered mass slope, as it does not take into account the possibility that the stellar mass

may be distributed in a way that does not exactly follow the contours in the potential, and that the two populations may have quite different 3D shapes and relative alignments.

Kowalczyk et al. (2013) carry out a similar study to that of Laporte et al. (2013a), but making use of idealised simulations of dSph galaxies orbiting in a static MW potential. The galaxies are initialised as perfect stellar discs, embedded in spherical dark matter subhaloes, with the two dSph stellar populations assumed to originate from the inner and outer disc. These authors make use of the velocity dispersion measured within R_e when computing Γ, in contrast to the method of Walker & Peñarrubia (2011), which uses the dispersion averaged over the whole galaxy.[11] They find that Γ can be biased both low *and* high with respect to the true slope, depending on the line of sight, and that the results are most accurate for galaxies that are close to spherical. If Γ was overestimated for Sculptor or Fornax by Walker & Peñarrubia (2011) then the confidence with which they exclude an inner NFW slope would be artificially high. However, as pointed out by Laporte et al. (2013a), since Kowalczyk et al. (2013) use the dispersion measured within R_e for each population, this may yield significantly different results to the method used by Walker & Peñarrubia (2011). Indeed, Laporte et al. (2013a) demonstrate that if they measure the dispersion within R_e then their analysis using the Aquarius simulations can also overestimate Γ. Note that in this case Γ is *more accurate* (i.e. *less biased*), albeit typically with a larger scatter. However, since the objective of applying the methodology of Walker & Peñarrubia (2011) is to place *conservative* upper limits on the inner density profile slopes, in the context of ruling out NFW profiles, more accurate values of Γ are deemed to be undesirable, since in this case the slope is much more likely to be *overestimated*, hence leading to unreliable (i.e. non-conservative) NFW profile exclusion confidence values.

If it can be assumed that the biases in the estimated masses for the two stellar populations are each determined *entirely* by the 'choice' of sight-line orientation relative to the whole dynamical system (galaxy) for a single observer, and no other factors are relevant (including e.g. the 3D shapes and

[11] Kowalczyk et al. (2013) refer to the estimator they use as that of Wolf et al. (2010). In fact, Wolf et al. (2010) advocate averaging the dispersion over all stars belonging to the galaxy.

relative orientations of the distinct stellar populations), then the mass bias will of course be identical for the two stellar populations,[12] and the line of sight choice will not affect the recovered Γ value. However, it seems somewhat unlikely that this assumption would hold in general for real galaxies, since the two metallicity populations are likely to have formed by different processes and at different times (Benítez-Llambay et al., 2016). The two populations might have different levels of dispersion support, velocity dispersion anisotropy, 3D shapes, alignments relative to their host halo and to each other, and may be in different dynamical states relative to equilibrium; each factor potentially introducing scatter into the two mass estimates. Differences in these important properties between the two populations could result in the introduction of additional uncorrelated biases into the two recovered mass values, and hence lead to propagation of unanticipated errors into the measurement of Γ.

Clearly, the validity of assuming that the two populations yield mass estimates with identical (and hence irrelevant) random biases depends inversely on the extent to which it is valid to assume that the populations are independent. Nonetheless, if we are to exclude NFW profiles within dSphs with multiple stellar populations, then surely the conservative approach is to assume that the two populations are indeed independent, with uncorrelated random mass errors. In fact, the APOSTLE galaxies studied here exhibit a range of alignment angles and 3D shapes relative to their host haloes (see Fig. 2.16). Therefore it seems entirely plausible that the dual stellar populations found in certain dSphs could have different orientations and shapes with respect to each other (and their host halo), depending on their physical origins.

In the special (and perhaps unlikely) limiting case where the two stellar populations are in perfect morphological alignment with each other, Fig. 2.14 corresponds to the prediction from our simulations for the mean value of the estimated to true mass ratio for the Walker et al. (2009) estimator, as a function of viewing angle relative to the (here assumed perfectly aligned) principal axes of the stellar populations. Even though this mean bias would cancel out when computing Γ, there is also a substantial scatter around

[12] Ignoring from now on the claimed dependence on the stellar population concentration within the halo ($r_{\rm half}/r_{\rm DM}$), since this acts to make the NFW exclusion limits more conservative.

the mean expectation, for any randomly chosen line of sight, as shown in Fig. 2.15. Thus even for perfect alignment of the populations with each other (and a single fixed line of sight), the fractional errors on the estimated masses for the two populations may be significantly different, thus perturbing the inferred value of Γ, as a result of e.g. differences in the 3D shapes, levels of dispersion support, or velocity dispersion anisotropies of the two populations. Considering the distribution of the standard deviation values, computed directly from the pixels shown in Fig. 2.15, the median value of the standard deviation (evenly weighted over the spherical surface) is $\sigma[\alpha_{\rm Walker}(\theta,\phi)] = 0.140^{+0.051}_{-0.028}$, where the quoted range corresponds to the $16^{\rm th}$ and $84^{\rm th}$ percentiles. Thus under the assumption of exact alignment of the populations, this value of $\sigma[\alpha_{\rm Walker}(\theta,\phi)]$ could be taken as the scatter in the estimated masses for the individual populations, for a single randomly chosen line of sight. Yet clearly the expected scatter increases as the populations move away from the state of perfect alignment assumed in this special case, and the mean biases for different sight-line orientations relative to the individual populations become relevant (see Fig. 2.14).

To represent the various uncertainties involved in estimating the enclosed masses for the two stellar populations, we assume in the following that, for each population, the estimated mass has an associated 1σ fractional error of $\sigma_{\rm sys} = 0.2$. If the important properties of the two stellar populations are assumed to be independent, then the choice of $\sigma_{\rm sys} = 0.2$ seems reasonable in general given our analysis of the uncertainties involved in estimating the dynamical masses of realistic simulated galaxies in §2.5. This representative scatter of 20 percent is relatively high compared to the standard deviations of the α values for most line of sight orientations relative to the stellar principal axes for the galaxies in our sample, although much larger scatters are observed for many viewing angles. Within the context of obtaining *conservative* exclusion confidences for NFW-like inner density profile slopes, our choice of $\sigma_{\rm sys}$ might in fact be considered somewhat low; in which case the results presented below constitute an underestimate of the impact of the mass estimator errors on the reliability of the results of Walker & Peñarrubia (2011). However, we are ignorant of any underlying correlations between the properties of the dual populations in Sculptor or Fornax that may conspire to bias the estimated enclosed masses in a similar way for the two radii, thus

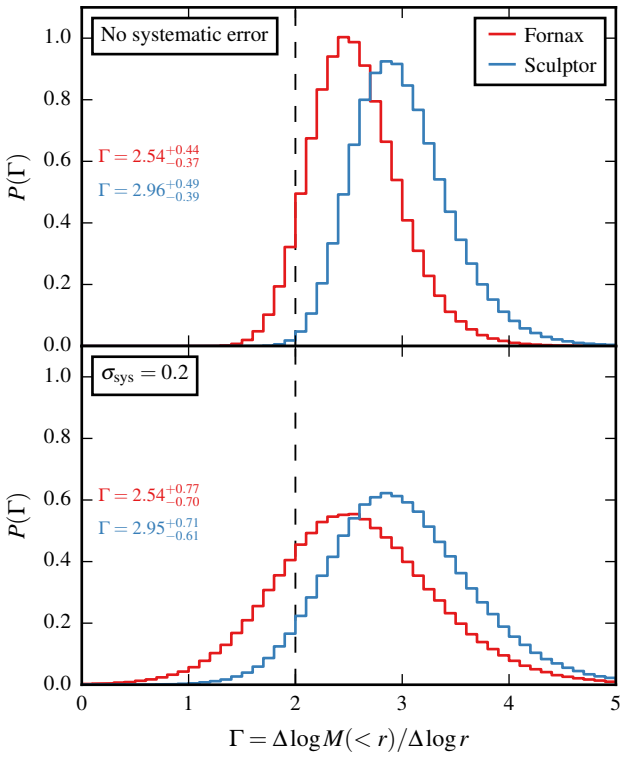

Figure 2.17 – Distributions, $P(\Gamma)$, of the logarithmic mass profile slope, Γ, derived using chemo-dynamically distinct stellar subpopulations in Sculptor and Fornax (see equation 2.15). The upper panel shows the distributions from sampling the posterior probability distribution functions for the half-light radii and velocity dispersions from Walker & Peñarrubia (2011). The lower panel shows the results of repeating this procedure, but introducing a 1σ scatter of $\sigma_{\rm sys} = 0.2$ in the mass estimates for both populations. The median Γ value and $16^{\rm th} - 84^{\rm th}$ percentile confidence interval are given in each panel for each distribution, in the same colours as the lines. Each distribution is normalised such that $\int_{-\infty}^{\infty} P(\Gamma)\,{\rm d}\Gamma = 1$. The vertical dashed line shows the central (maximum) mass slope for an NFW profile.

reducing the error on the inferred mass slope.

In order to assess the impact of the systematic errors inherent in simple mass estimators of the form assumed in equation (2.15) on the results of Walker & Peñarrubia (2011), we begin by creating a simple model of their analysis procedure. Approximating the published posterior probability distribution functions for $\log_{10}(r_2)$, r_1/r_2, $\log_{10}(\sigma_1^2)$, and $\log_{10}(\sigma_2^2)$ as Gaussian (using the tabulated medians and $16^{\rm th} - 84^{\rm th}$ percentile ranges), we generate 10^6 values for Γ by independently drawing random values consistent with the four distributions, and applying equation (2.15). This approach ignores the correlations between the free parameters of the system, which are ac-

counted for in the MCMC chains. However, the resulting $P(\Gamma)$ distributions for the two dSphs, shown in the upper panel of Fig. 2.17, are very similar to those obtained by Walker & Peñarrubia (2011, cf. their fig. 10). Indeed, our median values for the slopes of $\Gamma = 2.54^{+0.44}_{-0.37}$ for Fornax and $\Gamma = 2.96^{+0.49}_{-0.39}$ for Sculptor are in excellent agreement with their results.

The lower panel of Fig. 2.17 shows the result of including a representative 1σ fractional systematic error of $\sigma_{\rm sys} = 0.2$ in the mass estimates for both populations. In detail, we repeat the process described above, but for each set of radii and dispersions, we multiply $M(< r_1)$ and $M(< r_2)$ (or equivalently, the squared dispersions) by random values, E_1 and E_2, respectively, drawn separately from a Gaussian distribution centred at one with standard deviation $\sigma_{\rm sys}$, and then proceed to apply equation (2.15) as before.[13] Including the systematic scatter spreads out the $P(\Gamma)$ distributions, with a negligible impact on the median values, giving $\Gamma = 2.54^{+0.77}_{-0.70}$ for Fornax and $\Gamma = 2.95^{+0.71}_{-0.61}$ for Sculptor. In terms of the significance as defined in equation (2.16), the $P(\Gamma)$ distributions including the systematic error disfavour NFW slopes with confidence of only 77.7 and 94.5 percent for Fornax and Sculptor respectively. Clearly, the systematic errors associated with mass estimators of the form given in equation (2.1) have the potential to reduce the confidence with which it can be claimed that the results of Walker & Peñarrubia (2011) rule out the presence of dark matter cusps as steep as that of the NFW profile in Fornax and Sculptor, depending on the extent to which the individual stellar populations are independent from each other.

A deeper understanding of the uncertainties associated with the methodology of Walker & Peñarrubia (2011) will come from high resolution simulations where the chemo-dynamically distinct stellar populations in dSphs are treated in a realistic way, ideally as the result of cosmological initial conditions, such as in the APOSTLE simulation suite. Dwarf galaxies with multiple stellar populations may reside in our simulations, facilitating direct tests of methods that use estimators of enclosed dynamical masses to infer density slopes. Such analysis is beyond the scope of the present work; we leave the identification and classification of such galaxies, and analysis of their physical origins, to future studies using APOSTLE and similar high resolution cosmological hydrodynamical simulations.

[13] We impose a lower limit of 0.01 for each of E_1 and E_2.

2.7 An optimal estimator

The mass estimators discussed in the preceding sections have a common form, as parameterised in equation (2.1), where λ is the 3D radius within which the mass is estimated, in units of R_e, and μ is the estimator prefactor. Extending this general idea, it may be that the radius within which the velocity dispersion is measured influences the accuracy of the recovered mass. Including an additional dimensionless parameter, ν, we can write a more flexible generalisation,

$$M(<\lambda R_e) = \frac{\mu \langle \sigma_{los}(<\nu R_e)\rangle^2 R_e}{G}\,, \tag{2.17}$$

where $\langle \sigma_{los}(<R)\rangle$ is the line-of-sight stellar velocity dispersion measured within a projected radius, R, of the galactic centre. Up to now, we have averaged the velocity dispersion over the whole galaxy ($\nu \to \infty$), in keeping with the estimators proposed by Walker et al. (2009) and Wolf et al. (2010).

Setting aside the theoretical motivation for certain choices of the parameters in equation (2.17) which, at least for the estimators of Walker et al. (2009) and Wolf et al. (2010), are based on the assumptions that underpin the spherical Jeans equation combined with additional simplifications, we now ask: empirically, for the population of dispersion-dominated galaxies in our simulations, exactly which set of estimator parameters (λ, μ, ν) yields an unbiased estimate of the true dynamical mass within some radius, with minimum scatter?

Considering our dispersion-dominated galaxy sample, the upper panel of Fig. 2.18 shows the standard deviation of the estimated to true mass ratio, α, for the unbiased estimator obtained by setting μ such that the mean α is one for a given λ and ν. This grid has been computed by projecting over 1536 evenly distributed lines of sight, with a grid spacing of 0.01 in λ and ν. The scatter in the estimates varies strongly with λ, and depends more weakly on ν. However, the impact of ν on the scatter becomes more significant near the location of minimum scatter in this plane (diamond marker). The middle panel in Fig. 2.18 shows the scatter as a function of λ for the case where ν is chosen to minimise the scatter for each value of λ (solid line), and also where the velocity dispersion averaged over the whole galaxy is used instead

Table 2.3 – Parameters (λ, μ, ν) that yield an estimated to true mass ratio, α, of one on average, with minimum scatter, calibrated using all galaxies in our dispersion-dominated sample (see equation 2.17). The three sets of parameters shown result from different constraints on ν, which sets the projected radius within which the line-of-sight stellar velocity dispersion is measured. The final row gives the scatter for each parameter set (standard deviation of α). The quoted uncertainties are the $16^{\text{th}} - 84^{\text{th}}$ percentile confidence limits, from 10^4 bootstrap samples of the galaxies, for each constraint on ν.

Parameter	$0.1 < \nu < 4.0$ (optimum set)	$\nu = 1$	$\nu \to \infty$
λ	$1.77^{+0.08}_{-0.06}$	$1.76^{+0.09}_{-0.05}$	$1.91^{+0.09}_{-0.03}$
μ	$5.99^{+0.37}_{-0.33}$	$5.94^{+0.42}_{-0.28}$	$6.95^{+0.48}_{-0.18}$
ν	$1.04^{+0.10}_{-0.17}$	1	∞
Scatter	$0.204^{+0.005}_{-0.006}$	$0.204^{+0.005}_{-0.006}$	$0.220^{+0.005}_{-0.006}$

(dashed line). In the case where ν is allowed to vary, the scatter is minimised for $\nu \approx 1$, while including the whole galaxy in the dispersion measurement leads to a preference for higher λ.

The sets of parameters that minimise the scatter for these two treatments of ν are given in Table 2.3, along with the associated scatter in the α distributions. We shall refer to the parameters found in the case where ν is allowed to vary as the 'optimum' set. The parameters in the case where the dispersion is measured within exactly R_e is also given in Table 2.3 ($\nu = 1$) and yield an equivalent level of scatter as in the optimum case. The uncertainties on the parameters and the scatter values given in Table 2.3 are derived from bootstrap resampling of the galaxies, where for each constraint on ν, we draw 10^4 samples of the same number of galaxies as in the true sample, and compute the 16^{th} and 84^{th} percentiles of the distributions of the parameters and scatter in α resulting from applying the minimisation procedure to each sample. Fig. 2.19 shows the full α distribution in the optimum case. The reduction in the minimum scatter due to switching from $\nu \to \infty$ to the optimum parameter set is 7.0 percent.

The mean α values and associated scatter obtained from applying the estimators of Walker et al. (2009) and Wolf et al. (2010) to the same galaxy sample as used in the calibration are shown in the lower panels of Fig. 2.18. Both of these estimators have a slightly smaller scatter than in the unbiased

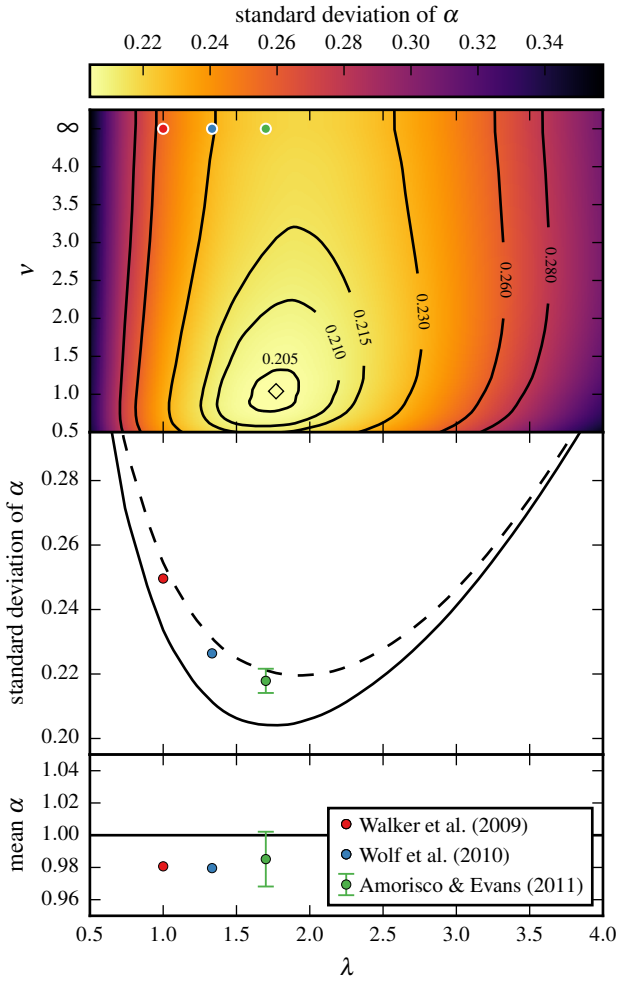

Figure 2.18 – Calibration of mass estimator parameters (λ, μ, ν) using all dispersion-dominated galaxies in our simulated sample (see equation 2.17). The upper panel shows the standard deviation of the estimated to true mass ratio α as a function of λ and ν, where the proportionality constant μ is chosen such that the estimator is *unbiased* (mean α equals one) for each combination of λ and ν. The scatter shown is linearly interpolated between $\nu = 4$ and the case where the dispersion is averaged over the whole galaxy ($\nu \to \infty$). The lines show contours of constant standard deviation, as labelled. The diamond symbol indicates the location of minimum scatter (see Table 2.3). The middle panel shows the scatter as a function of λ, for the choice of ν which minimises the scatter at each value of λ (solid line), and using instead the velocity dispersion averaged over the whole galaxy (dashed line). The points show the corresponding results from applying the estimators of Walker et al. (2009), Wolf et al. (2010), and Amorisco & Evans (2011), which each use the dispersion averaged over the whole galaxy. Note that the α distribution for a given estimator can have a lower scatter than the dashed line, provided that it is biased in the mean. The lower panel shows the mean α for each of these three estimators.

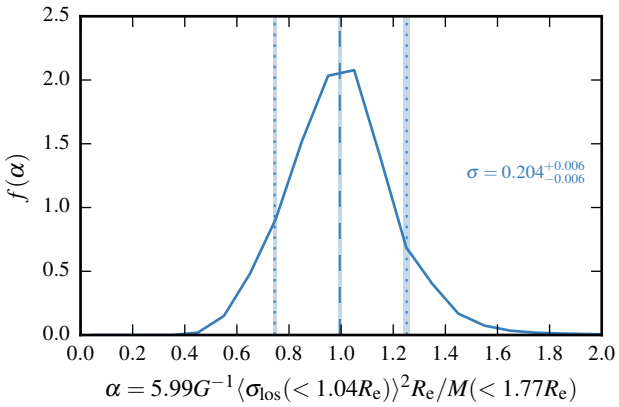

Figure 2.19 – Distribution, $f(\alpha)$, of the estimated to true mass ratio, α, for the unbiased estimator with minimum scatter of the form given in equation (2.17). The optimum parameters are given in Table 2.3 (as in the horizontal axis label). The distribution considers all dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$), using projections from 1536 evenly distributed lines of sight. The dashed line shows the median, and the dotted lines show the $10^{\rm th}$ and $90^{\rm th}$ percentiles. The standard deviation, σ, is also shown. The shaded regions around the vertical lines and the quoted errors on σ are the $16^{\rm th}$ – $84^{\rm th}$ percentile confidence limits, derived from applying the estimator to 10^4 bootstrap samples of the galaxies. Note that the similar σ errors given in Table 2.3 for this estimator are instead derived from the distribution of optimum σ values obtained from repeatedly calibrating equation (2.17) on bootstrap samples of the same galaxies. Each projection of each galaxy contributes to the distribution with equal weight, and the distribution is normalised to have unit area.

case that uses the dispersion averaged over the whole galaxy (as these two estimators do), for their λ. Using the optimum parameter set leads to reductions in the scatter of 18.2 and 9.9 percent relative to the Walker et al. (2009) and Wolf et al. (2010) versions, respectively, while correcting for the bias exhibited by each estimator, as shown in Fig. 2.18.

We note that Amorisco & Evans (2011) obtain an empirical estimator with similar parameters to our optimum result, from phase-space modelling of the dSphs of the MW (they find $\lambda = 1.7$ and $\mu = 5.8 \pm 1.0$; cf. Table 2.3). The choice of $1.7R_{\rm e}$ for the enclosing radius minimises the dependence in their analysis of the recovered mass on the assumed halo density profile and its associated scale-length, assuming that the dSphs have isotropic velocity dispersions in their centres, with relatively flat $\sigma_{\rm los}(R)$ profiles. The mean α and scatter resulting from applying this estimator to our dispersion-dominated sample are shown in Fig. 2.18 alongside the Walker et al. (2009) and Wolf et al. (2010) results, where the error bars indicate the spread due to the quoted uncertainty on μ.[14]

2.8 Conclusions

In this chapter we have presented the key intrinsic and observable properties of galaxies drawn from the APOSTLE simulations of the Local Group that are relevant to Jeans analysis in general, and in particular to the use of simple mass estimators of the form given in equation (2.1), as advocated by Walker et al. (2009) and Wolf et al. (2010). The simulated galaxies have realistic stellar density distributions and line-of-sight stellar velocity dispersions, combined with a range of 3D shapes, levels of dispersion support, and stellar velocity dispersion anisotropy.

Applying the mass estimators proposed by Walker et al. (2009) and Wolf et al. (2010) to many projections of each galaxy in our sample, we have found that the estimators each have a small bias in the estimated to true mass ratio,

[14] Amorisco & Evans (2011) state that their analysis makes use of the velocity dispersion at the galactic centre, however it is not clear to what level the dispersions used can be considered to be 'central'. The dataset used is that compiled by Walker et al. (2009, 2010), who note that the dispersions are global averages. Therefore we use the velocity dispersion averaged over the whole galaxy when applying the estimator of Amorisco & Evans (2011), in order to be consistent with the data used in their analysis.

α, combined with a relatively large scatter. This scatter is 23 or 25 percent at the 1σ level overall for dispersion-dominated galaxies, respectively (see Fig. 2.9).

The dependence of α on various galaxy properties is summarised in Fig. 2.16. The scatter in α depends strongly on the shape of the galaxy, such that the α distributions are sharply peaked for stellar mass distributions that are close to spherical symmetry. However, the dependence of the scatter on the observable 2D shape on the sky is much weaker, i.e. it is difficult to identify ideal spherically symmetric systems in projection. The scatter also increases for progressively larger radial bias in the stellar velocity dispersion anisotropy, and progressively higher levels of rotational support, and for galaxies that exhibit closer alignment with their host dark matter haloes. We also find that the scatter is lowest for galaxies that have relatively low stellar kinetic energy for their gravitational potential energy, and for galaxies that are highly extended, in units of their halo scale radius. The mass estimator accuracy is more sensitive to the 3D shape (and alignment with respect to the observer) of the galaxy than to that of the host dark matter halo. The dispersion-dominated satellites of the MW and M31 analogues found in the simulations yield a smaller scatter in α (typically 20–22 percent) than the dispersion-dominated field galaxies (typically 24–26 percent), as shown in Fig. 2.12.

Considering the dependence of α on the orientation of the line of sight with respect to the galaxy, we find that the dynamical masses are overestimated with a relatively large scatter for lines of sight coincident with the stellar morphological major axis, underestimated with a modest scatter along the minor axis, and close to accurate with a modest scatter for the observations along the stellar principal axis of intermediate length (see Figs. 2.14 and 2.15).

Adopting a representative value of 20 percent for the 1σ systematic scatter in α for satellite galaxies, we have shown that the systematic uncertainties inherent in the simple estimators significantly increase the errors on the estimated dynamical masses of dSph satellites of the MW from Wolf et al. (2010) in cases where the observational errors on the half-light radii and velocity dispersions are relatively small. Including the same representative scatter within a simple model of the analysis of Walker & Peñarrubia (2011)

demonstrates that systematic uncertainties can also greatly reduce the significance with which the mass profile slopes derived for two independent stellar populations in Sculptor and Fornax are inconsistent with an NFW dark matter density profile. This result depends on the level to which the errors on the estimated masses for the two populations can be assumed to be independent.

Finally, we have investigated the scatter in the estimated to true mass ratio for the set of unbiased estimators of the form given in equation (2.17), considering all the dispersion-dominated galaxies in our simulated sample. We find that using the optimum set of parameters given in Table 2.3 results in a reduction in the scatter of between 10 and 20 percent with respect to similar estimators in the literature (Walker et al., 2009; Wolf et al., 2010; Amorisco & Evans, 2011), combined with the removal of bias in the mean. The scatter is optimised when the line-of-sight velocity dispersion is measured within a radius close to the projected stellar half-mass radius, $R_{\rm e}$ (rather than over the whole galaxy as is typical):

$$M(<1.77R_{\rm e}) = \frac{5.99\langle\sigma_{\rm los}(<1.04R_{\rm e})\rangle^2 R_{\rm e}}{G} \,. \tag{2.18}$$

This equation gives our optimum result in the case where the radius within which the velocity dispersion is measured is treated as a free parameter. If instead we fix the dispersion measurement to consider stars within exactly $R_{\rm e}$, the resulting scatter is the same as that obtained using the formally optimal parameter set (see Table 2.3).

Chapter 3

Galaxy Clustering as a Function of Stellar Mass

3.1 Introduction

Galaxy formation involves the interplay between a variety of physical processes, such as the suppression of star formation in massive haloes due to the shutting down of gas cooling by active galactic nuclei, and the heating of the interstellar medium by supernovae. These processes vary in importance with redshift and the mass of the dark matter halo that hosts the galaxy. Semi-analytic models use simplified expressions and approximations to describe these complex physical processes, allowing predictions to be generated for how dark matter haloes are populated with galaxies (Benson et al., 2000a; Cole et al., 2000; Hatton et al., 2003; Bower et al., 2006; Croton et al., 2006; Monaco et al., 2007; Somerville et al., 2008).

The model processes are constrained through comparison to observed statistics of the galaxy population, such as the luminosity function (e.g. Gonzalez-Perez et al. 2014). We can also constrain the galaxy formation physics by measuring galaxy clustering as a function of intrinsic properties, such as stellar mass. Appealing to the clustering as a function of galaxy prop-

erties provides a constraint on how dark matter haloes are populated with galaxies (Kauffmann et al., 1999; Benson et al., 2000b). Galaxy formation models have been used to make predictions for comparison to observational measurements of the clustering as a function of different properties, such as luminosity, star formation rate, and stellar mass (e.g. Norberg et al. 2001; Kim et al. 2009; Meneux et al. 2009; Guo et al. 2011; Li et al. 2012; Guo et al. 2013; Henriques et al. 2013; Marulli et al. 2013; Bielby et al. 2014). Some models explicitly consider galaxy clustering when setting their parameters, e.g. Guo et al. (2011). For models that have been constructed without using galaxy clustering as a constraint, comparison to clustering measurements is a test of their predictive power. The models considered in this chapter have not been calibrated to reproduce any clustering measurements, or stellar mass function data.

Despite the perceived complexity of semi-analytic models, Contreras et al. (2013) show that the clustering predictions made by different groups are robust, at least for samples defined by stellar mass. They considered the Durham GALFORM (Bower et al., 2006; Font et al., 2008) and Munich LGALAXIES (Bertone et al., 2007; de Lucia & Blaizot, 2007; Guo et al., 2011) model families. These two sets of models differ in their implementation of various physical processes of galaxy formation. Contreras et al. find differences in the clustering predictions on small scales, which they trace to the handling of galaxy mergers. We explore the impact of the choice of galaxy merger scheme further in this chapter.

The stellar masses of real galaxies have to be derived from observable properties such as broad-band photometry, which requires the assumption of a star formation history (SFH), stellar initial mass function (IMF), stellar population synthesis (SPS) model, and dust extinction model (e.g. Pforr et al. 2012; Mitchell et al. 2013). The choice of the particular set of assumptions used has an impact on the recovered masses, and so there are important systematic uncertainties inherent in the derived stellar masses of observed galaxies. These uncertainties may influence comparisons of galaxy clustering as a function of stellar mass, between theoretical models and observational measurements. Thus if the clustering as a function of stellar mass is to be used to constrain theoretical models, we must take care to treat appropriately the difficulties in estimating the stellar masses of observed galaxies.

Contreras et al. (2013) compared theoretical predictions by examining galaxy samples of fixed abundance. We extend this study, motivated by the work of Mitchell et al. (2013, see also Marchesini et al. 2009 and Pforr et al. 2012), and demonstrate a new methodology that we find is essential for comparing model predictions to observational clustering data. Our technique consists of carrying out broad-band spectral energy distribution (SED) fitting to compute stellar mass estimates for model galaxies, in the manner that is typically applied to observed galaxies. Such a treatment is particularly important for understanding the influence of dust extinction, and the impact of Eddington bias in higher mass bins where the stellar mass function is steep.

We present the clustering predictions of two new GALFORM variants: the models of Gonzalez-Perez et al. (2014) and Lacey et al. (2016) hereafter Gon14 and Lac14.[1] These models have been calibrated to reproduce observations in the WMAP7 cosmology (Komatsu et al., 2011). The models make use of the MS-W7 N-body simulation (Guo et al., 2013), which is a new version of the Millennium Simulation (Springel et al., 2005) updated to use the WMAP7 cosmological parameters. These models take advantage of extensions to the galaxy formation physics implemented in GALFORM, making use of an empirical law to determine star formation rates (Lagos et al., 2011). For discussions of the implications for galaxy formation models of the WMAP7 cosmology, see Guo et al. (2013) and Gon14.

We compare the model predictions to derived results from the Sloan Digital Sky Survey (SDSS; York et al. 2000; Abazajian et al. 2009), the Galaxy and Mass Assembly Survey (GAMA; Driver et al. 2011), and the VIMOS Public Extragalactic Redshift Survey (VIPERS; Guzzo et al. 2014). Together, these surveys have measured the clustering of galaxies as a function of stellar mass up to redshift $z \sim 1$ (Li et al. 2006; Farrow et al. 2015; Marulli et al. 2013; respectively). In particular, we compare to these observational results at redshifts 0.1, 0.2, and 0.6 respectively. The specific star formation rate falls dramatically over this interval (e.g. Weinmann et al. 2011), so the interplay between the different galaxy formation processes could change over

[1] The model of Lacey et al. (2016) was known as Lac14 at the time the research presented in this chapter was completed, and published as Campbell et al. (2015). The 'Lac14' model used here is identical to that described by Lacey et al. (2016).

the cosmic times considered. Through comparing to these measurements, we assess both the importance of carrying out SED fitting to the model photometry, and the level of agreement between the measurements and the model predictions.

The layout of this chapter is as follows. Details of the galaxy formation models and observational data are given in §3.2 and §3.3. Our methodology is described in §3.4. Our results are presented and discussed in §3.5. Concluding remarks are given in §3.6. All magnitudes are on the AB system (Oke, 1974). The assumed ΛCDM cosmological parameters are listed in §3.2.2. Comoving length units are used throughout this chapter.

3.2 Theoretical modelling

This section describes the semi-analytic galaxy formation models used in our study of galaxy clustering as a function of stellar mass. We first give an overview of GALFORM (§3.2.1), followed by a description of the N-body dark matter only simulation used (§3.2.2), then contrast the two GALFORM models compared (§3.2.3), and describe a new scheme for the treatment of satellite galaxy orbits and merger time-scales (§3.2.4).

3.2.1 Overview of GALFORM

Galaxy formation within dark matter haloes, as followed in recent GALFORM variants, can be broken down into several key processes: (i) formation and hierarchical growth of dark matter haloes, (ii) shock-heating of baryonic material falling into haloes, followed by radiative cooling and disc formation, (iii) quiescent star formation in discs (and bursts due to instabilities), (iv) suppression of gas cooling (and hence of star formation) through feedback from supernovae, active galactic nuclei (AGN), and photoionisation of the intergalactic medium, (v) chemical enrichment of the stars and gas, and (vi) mergers of satellite galaxies with the central galaxy of their halo, due to dynamical friction, which can cause bursts of star formation. For an overview of the development of the GALFORM model, see Benson & Bower (2010). For details of the galaxy formation physics, see Baugh (2006) and Benson (2010).

In order to connect the predictions of galaxy formation models to the

properties of observed galaxies, a stellar population synthesis (SPS) model must be assumed, along with a model to describe dust attenuation. SPS models, such as those of Bruzual & Charlot (2003), Maraston (2005), and Conroy et al. (2009), compute the spectral energy distribution (SED) of a coeval stellar population with a given initial metallicity, as a function of age. Convolving this stellar SED with the star formation history of a galaxy (i.e. its star formation rate as a function of time; SFH), taking into account its chemical enrichment history (the metallicity of stars forming at a given time), yields the SED of the galaxy itself (Cole et al., 2000). SPS models require a stellar initial mass function (IMF) to be specified, which gives the distribution of masses of stars formed in a given episode of star formation.

Attenuation of starlight by dust in GALFORM is modelled in a physically motivated way, in which the stars and dust are mixed together, based on radiative transfer and the geometry of a disc and bulge (Silva et al., 1998; Ferrara et al., 1999; Cole et al., 2000; Lacey et al., 2011; Gonzalez-Perez et al., 2013). For a given photometric band, the galactic SED is weighted by the wavelength response of the filter and integrated to yield the flux.

3.2.2 The MS-W7 simulation

The cosmological parameters from WMAP7 (Komatsu et al., 2011) have been used in an N-body simulation similar to the Millennium Simulation of Springel et al. (2005). This MS-W7 simulation (Guo et al., 2013) has present day density parameters of $\Omega_{\mathrm{m},0} = 0.272$, $\Omega_{\mathrm{b},0} = 0.0455$, and $\Omega_{\Lambda,0} = 0.728$, for matter, baryons, and dark energy, respectively. The present day Hubble parameter is $H_0 = 100\,h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, where $h = 0.704$. The spectral index of primordial fluctuations is $n_{\mathrm{s}} = 0.967$, and the linear perturbation amplitude is $\sigma_8 = 0.810$. The simulation follows 2160^3 particles from redshift 127, in a volume of comoving side $L_{\mathrm{box}} = 500\,h^{-1}\,\mathrm{Mpc}$.

3.2.3 The Gon14 and Lac14 models

The Gon14 and Lac14 GALFORM models are based on the halo merger trees of the MS-W7 simulation. They are separate developments of the Lagos et al. (2012) model, which used an empirical star formation rate law introduced by Lagos et al. (2011). The Lagos et al. (2012) model in turn is based on that

of Bower et al. (2006), which introduced AGN feedback into GALFORM. The new models will be made publicly available in the Millennium Database.[2]

The IMF, $\zeta(M_\star)$, is defined such that the number of newly formed stars (per solar mass, e.g.) with stellar mass, $M_\star$, between $\log_{10}(M_\star/\mathrm{M}_\odot)$ and $\log_{10}(M_\star/\mathrm{M}_\odot) + \mathrm{d}\log_{10}(M_\star/\mathrm{M}_\odot)$ is given by $\zeta(M_\star)\,\mathrm{d}\log_{10}(M_\star/\mathrm{M}_\odot)$. A power law is often used to model (parts of) $\zeta(M_\star)$, such that

$$\zeta(M_\star)\,\mathrm{d}\log_{10}\left(\frac{M_\star}{\mathrm{M}_\odot}\right) = \zeta_0 \left(\frac{M_\star}{\mathrm{M}_\odot}\right)^{-x} \mathrm{d}\log_{10}\left(\frac{M_\star}{\mathrm{M}_\odot}\right), \tag{3.1}$$

for some normalisation, ζ_0, and slope, x. The Gon14 model uses a Kennicutt (1983) IMF, which is a broken power law. This has a slope of $x = 0.4$ for $M_\star < 1\,\mathrm{M}_\odot$, and $x = 1.5$ for $M_\star > 1\,\mathrm{M}_\odot$. In the Lac14 model, distinct IMFs are employed in quiescent star formation and bursts. In the former case, a Kennicutt IMF is used as in Gon14. However, for bursts, the IMF is taken to be a single power law with $x = 1$. Such non-universality of the IMF is argued to be necessary to match the observed number counts and redshift distribution of sub-millimetre galaxies (Baugh et al., 2005).

The SPS model used in the Gon14 model is a private release of the Bruzual & Charlot series from 1999, which is intermediate between Bruzual & Charlot (1993) and Bruzual & Charlot (2003). The Lac14 model uses the Maraston (2005) SPS model. The Maraston model attributes much more luminosity to stars in the thermally pulsating asymptotic giant branch (TP-AGB) phase than is done in the Bruzual & Charlot model. Such stars emit strongly in the near infra-red (NIR; e.g. MacArthur et al. 2010). The influence of TP-AGB stars in the Maraston model has been the source of some debate in the literature (e.g. Marigo & Girardi 2007; Kriek et al. 2010; MacArthur et al. 2010; Zibetti et al. 2013). Gon14 study the influence of using various alternative SPS models, including that of Maraston. They find that the choice of SPS model does not affect the evolution of the rest-frame optical and ultra-violet (UV) luminosity functions, but models incorporating strong TP-AGB emission yield significantly different evolution of the rest-frame NIR luminosity function (see also Tonini et al. 2009 and Henriques et al. 2011). This choice gives an improved match to the bright end of the rest-frame K-band luminosity function at high redshifts in the Lac14 model.

[2] `http://virgodb.cosma.dur.ac.uk:8080/Millennium`

Additional parameter differences between the Gon14 and Lac14 models are described in Appendix C.

3.2.4 Subhalo dynamical friction for satellites

In the standard models considered here, time-scales for the merging of satellites with the central galaxy in their host halo due to dynamical friction are computed in the models as described by Cole et al. (2000). This method assumes that when a new halo forms, each satellite galaxy enters the halo on a random orbit. The merger time-scale is then computed using an analytical formula that assumes the halo to be isothermal. While the Gon14 model makes use of the equations presented by Lacey & Cole (1993), a modified expression is used in the Lac14 model. This expression has been empirically fitted to numerical simulations to account for the tidal stripping of subhaloes (Jiang et al., 2008, 2010), but otherwise the treatment is the same; i.e. an analytic time-scale is computed as soon as a galaxy enters a larger halo. The satellite is considered to have merged with its central galaxy once the merger time-scale has elapsed, provided that this transpires before the halo merges to form a larger system, in which case a new merger time-scale is computed. Note that this scheme does not take into account that the satellite may still reside in a resolvable dark matter subhalo at the time the merger takes place.

We have implemented an alternative treatment of mergers, which makes use of the subhalo information from the simulation. The motivation for using this new scheme is that it is more faithful to the underlying N-body simulation, minimising the reliance on analytically determined orbits. Instead of assuming random initial orbits for satellites, they track the positions of their associated subhaloes. For cases where the subhalo containing a satellite can no longer be resolved following mass stripping, the position and velocity of the subhalo when it was last identified are used to compute a merger time-scale. This time-scale is then used in the same way as in the default scheme described above. The merger time-scale calculation assumes an NFW (Navarro, Frenk & White, 1996b, 1997) halo mass distribution to compute the orbital parameters of the satellite, combined with the analytical time-scale result of Lacey & Cole (1993). If a halo formation event occurs at a time after the subhalo is lost, a new merger time-scale for the satellite

is calculated in the same way, using instead the position and velocity of the particle that was the most bound particle of the subhalo when it was last identified in the simulation.

In this new GALFORM merger scheme, a satellite galaxy is not allowed to merge while it remains associated with a resolved subhalo. This treatment is similar to the scheme employed in the LGALAXIES model. The choice of merger scheme has an impact on the small-scale clustering, and contributed to the differences between the predictions of LGALAXIES and GALFORM reported by Contreras et al. (2013). The differences between the clustering predictions using the two types of merger scheme can be explained in terms of the concentration of the radial distribution of satellites, and also changes in the number of satellites. Requiring satellites to track their resolved subhaloes, rather than computing an analytical merger time-scale as soon as a galaxy becomes a satellite, results in a more radially extended distribution of satellites (as demonstrated by Contreras et al.).

No model parameters have been recalibrated when using the new merger scheme. This would be likely to involve changing the strength of AGN feedback, and the time-scale for gas return from supernovae. We leave such calibration for future work. When comparing the model predictions using the new merger scheme to observational estimates of galaxy clustering as a function of stellar mass (in §3.5.2.3), we require that the model stellar mass functions reproduce those inferred from observations, through abundance matching (see §3.4.1.2). It is possible that making use of the subhalo mass at the time it was last identified results in shorter merger time-scales than would be computed if any subsequent mass stripping of the subhalo could be taken into account.

3.3 SED fitting and observational data

In theoretical models such as GALFORM, the galactic stellar masses are predicted from the star formation histories. However, for observed galaxies, the stellar masses are not known directly but must be derived from observables. SED fitting is a popular technique for computing stellar masses. This section first describes the SED fitting procedure and then gives details of the observational data considered in this chapter. We consider different surveys in order

Table 3.1 – Observed galaxy samples. The galaxy numbers refer to galaxies contributing to the correlation function, $w_{\rm p}(\sigma)$, in Figs. 3.7, 3.8, and 3.9. r and i are the SDSS r-band and CFHTLS i-band apparent magnitudes, respectively. $M_{0.1_r}$ is the SDSS r-band absolute magnitude (at $z = 0.1$). The median redshifts for GAMA and VIPERS are given in order of increasing stellar mass interval (see Figs. 3.8 and 3.9). The GALFORM comparison redshifts are restricted to the set of output snapshots of the MS-W7 simulation. $\pi_{\rm max}$ is the $w_{\rm p}(\sigma)$ integration limit (see equation 3.8). The lower part of the table lists the Schechter (1976) function parameters, α, $\tilde{M}_\star$, and $\tilde{\Phi}$ (see equation 3.2), for the stellar mass functions of the samples. The SDSS mass function fit is as measured by Wang et al. (2006). The GAMA stellar mass function has been measured by Baldry et al. (2012) for $z < 0.06$ (see §3.4.1.2), and is represented by the sum of two Schechter functions (left, right), with a common characteristic mass. The VIPERS fit is for $0.5 < z < 0.6$ (Davidzon et al., 2013).

Sample property	SDSS	GAMA	VIPERS
Source	Li et al. (2006)	Farrow et al. (2015)	Marulli et al. (2013)
Galaxies used	157884	50770	17100
Apparent magnitude limits	$14.5 < r < 17.77$	$r < 19.8$	$i < 22.5$
Absolute magnitude limits	$-23 < M_{0.1_r} - 5\log_{10}(h) < -16$		
Redshift range	$0.01 < z < 0.3$	$0.14 < z < 0.24$	$0.5 < z < 0.7$
Median redshift	~ 0.1	0.19, 0.20	0.62, 0.62, 0.62
GALFORM redshift for comparison	0.089	0.17	0.62
$\pi_{\rm max}$ $[h^{-1}\,{\rm Mpc}]$	40	47	30
α	-1.073 ± 0.003	$-0.35 \pm 0.18, -1.47 \pm 0.05$	$-0.95^{+0.03}_{-0.02}$
$\log_{10}(\tilde{M}_\star\ [h^{-2}\,{\rm M}_\odot])$	$10 + \log_{10}(4.11 \pm 0.02)$	$(10.66 \pm 0.05) + \log_{10}(0.7^2)$	$10.87^{+0.02}_{-0.02} + \log_{10}(0.7^2)$
$\tilde{\Phi}$ $[10^{-3}\,h^3\,{\rm Mpc}^{-3}\,{\rm dex}^{-1}]$	$(20.4 \pm 0.1)/\ln(10)$	$(3.96 \pm 0.34, 0.79 \pm 0.23)/0.7^3$	$1.42^{+0.06}_{-0.07}/0.7^3$

to probe a range of lookback times (see Table 3.1). A detailed discussion of SED fitting can be found in Mitchell et al. (2013).

3.3.1 SED fitting

Broad-band SED fitting is essentially the reverse of the process described in §3.2.1 for computing galaxy SEDs and broad-band photometry in GALFORM. A grid of template SEDs is generated, given an assumed SPS model, IMF, SFH (including assumptions about chemical enrichment), and dust extinction. The observed SED (i.e. broad-band photometry) is then used to identify the maximum-likelihood template SED (e.g. Mitchell et al. 2013). The stellar mass associated with the best-fitting SED is then assigned to the observed galaxy. The SFH is usually taken to be of a simple exponentially declining form, in contrast to the complicated form predicted in theoretical galaxy formation models (see e.g. Baugh 2006). Mitchell et al. show that on average, the detailed form of the SFH is not important (see also Simha et al. 2014). In SED fitting, dust attenuation is usually dealt with by assuming the so-called

Calzetti law (Calzetti et al., 2000), which is equivalent to assuming that the dust forms a uniform sheet between the galaxy and the observer. GALFORM assumes a physically motivated distribution of dust in model galaxies (see §3.2.1), and applies dust extinction in massive galaxies that is very different from the Calzetti law, resulting in systematic errors of up to an order of magnitude in $M_\star$ (Mitchell et al., 2013). Conroy et al. (2009) report that the uncertainties in stellar masses derived from broad-band SED fitting are in the region of 0.3 dex at redshift zero, considering the uncertainties in the details of different stages of stellar evolution, while at $z = 2$ the uncertainty for bright red galaxies rises to ~ 0.6 dex.

3.3.2 Sloan Digital Sky Survey

The Sloan Digital Sky Survey (SDSS) uses photometry in the u, g, r, i, and z filter bands to identify candidates for spectroscopic follow-up (York et al., 2000). Li et al. (2006) estimated the clustering of galaxies as a function of stellar mass using the New York University value-added galaxy catalogue (Blanton et al., 2005), which is based on the second SDSS data release (Abazajian et al., 2004). This catalogue has photometry covering 3514 square degrees, and spectroscopy covering 2627 square degrees (about 85 percent complete), for redshifts $z \lesssim 0.3$. Li et al. define their sample of galaxies with the magnitude and redshift limits given in Table 3.1, yielding a total of 196238 galaxies. Subsets of this sample, defined in stellar mass, are used to study galaxy clustering.

Li et al. use the approach of Kauffmann et al. (2003) to estimate stellar masses. The amplitude of the 4000 Å break, D_{4000}, (Balogh et al., 1999) and the strength of the Balmer H-δ absorption line (Worthey & Ottaviani, 1997) are measured from a spectrum obtained with a 3 arcsec diameter fibre. These measurements encode information about the age of the galaxy's stellar population (D_{4000}), and can be used as indicators of whether recent star formation has been predominantly quiescent, or due to bursts (H-δ). The stellar mass-to-light ratio in the z-band is estimated for each galaxy, by fitting D_{4000} and H-δ to a Monte Carlo library of stellar populations, based on the Bruzual & Charlot (2003) SPS model. A Kroupa (2001) IMF is assumed in building the library, i.e. $x = -0.3$ for $M_\star < 0.08\,\mathrm{M}_\odot$, $x = 0.3$

for $0.08\,\mathrm{M}_\odot < M_\star < 0.5\,\mathrm{M}_\odot$, and $x = 1.3$ for $M_\star > 0.5\,\mathrm{M}_\odot$ (see equation 3.1). Exponentially declining SFHs are used, with additional random bursts. The dust extinction applied in the best-fitting model is assumed for each galaxy, making use of a power law attenuation curve, corresponding to a foreground dust screen (Charlot & Fall, 2000). The stellar mass is then found by using the derived mass-to-light ratio in the z-band, combined with full z-band photometry (i.e. not limited by the fibre diameter). In this way, the mass-to-light ratio and dust attenuation derived within 1.5 arcsec of the galactic centre are extrapolated over the full galaxy. The masses estimated following Kauffmann et al. using spectral features have been shown to have a scatter of about 0.1 dex with respect to those obtained using SED fitting to broad-band photometry (Blanton & Roweis, 2007; Li & White, 2009).

The clustering results presented by Li et al. use six bins spaced logarithmically in stellar mass,[3] covering $10^{8.69} < M_\star\,[h^{-2}\,\mathrm{M}_\odot] < 10^{11.69}$. All but the highest mass bin correspond to samples that are volume limited in stellar mass, where each volume limited stellar mass bin covers a different redshift interval. However, the highest mass bin is flux limited, and thus incomplete in stellar mass. A correction is made by Li et al. to the correlation function computed for this bin, by weighting the contribution from each galaxy pair by the maximum volume over which they could be detected in the survey volume. They find that applying the same approach to flux limited samples for the lower stellar masses produces good agreement with the clustering results for the volume limited samples. Thus we assume that the results do not suffer from incompleteness in stellar mass due to magnitude limiting.

Li et al. find increasing clustering amplitude as a function of stellar mass, with a sharp increase above the characteristic stellar mass ($\sim 10^{10.6}\,h^{-2}\,\mathrm{M}_\odot$), which is similar to the clustering trend they find when selecting galaxies by their r-band luminosity.

3.3.3 Galaxy and Mass Assembly Survey

The Galaxy and Mass Assembly Survey (GAMA) is a multi-wavelength survey (far UV to far IR), with complete spectroscopy to $r = 19.8$ (Driver et al. 2011; Liske et al. 2015; see Table 3.1). There are two clustering studies in

[3] The stellar mass bin edges of Li et al. (2006) have been converted from units of $\mathrm{M}_\odot$ to $h^{-2}\,\mathrm{M}_\odot$, using their $h = 0.7$ (Kauffmann et al., 2003).

GAMA focusing on stellar mass (Palamara et al. in preparation; Farrow et al. 2015). For the present analysis, we compare to the clustering of galaxies in differential stellar mass bins as computed by Farrow et al. using the GAMA-II data. This is selected from the seventh SDSS data release (Abazajian et al., 2009), with sky coverage of 180 square degrees, out to $z \lesssim 0.5$, and spectroscopic completeness of > 98 percent.

To estimate stellar masses, Farrow et al. use an empirical relation based on the observed $g-i$ colour and absolute magnitude in the i-band, as fitted by Taylor et al. (2011) to galaxies from the second GAMA data release, with $r < 19.4$ (Liske et al., 2015).[4] To derive this relation, Taylor et al. implement broad-band SED fitting using the Bruzual & Charlot (2003) SPS model, with a Chabrier (2003) IMF,[5] and exponentially declining SFHs. Extinction by dust is modelled using the Calzetti et al. (2000) law. The photometric bands used in the fitting are the GAMA u, g, r, i, and z. Taylor et al. use likelihood-weighting of all template SEDs, which can suppress discreteness effects due to the lack of interpolation between the small number of metallicities typically available in SPS models, improving on the common practice of taking the mode of the likelihood distribution (e.g. Mitchell et al. 2013).

Farrow et al. use samples selected in redshift and stellar mass to compute the correlation function. We compare to their intermediate redshift range (see Table 3.1), considering the stellar mass range $10^{9.5} < M_\star\ [h^{-2}\,\mathrm{M}_\odot] < 10^{11.5}$. They find that clustering amplitude increases with stellar mass. A decrease in clustering strength with redshift is noted for masses below $10^{10.5}\,h^{-2}\,\mathrm{M}_\odot$, with no significant redshift evolution above this.[6]

[4] This release is based on GAMA-I (Baldry et al., 2010).

[5] The Chabrier (2003) IMF has $x = 1.3$ for $M_\star > 1\,\mathrm{M}_\odot$, and adopts a smooth transition below $1\,\mathrm{M}_\odot$ to a slope similar to the Kroupa (2001) IMF (see equation 3.1).

[6] The analysis presented here was completed (and published as Campbell et al. 2015) before the final publication of the paper by Farrow et al. (2015). There are some minor differences between the GAMA data used here, and the data published by Farrow et al. (2015). In particular, we note that the stellar mass bins used here are different to those of Farrow et al. (2015), as is the line-of-sight integration limit for the projected correlation function. These small changes are insignificant from the point of view of our analysis. To avoid potential confusion with the final GAMA clustering data as published by Farrow et al. (2015), we label their projected correlation function measurements as *Farrow et al. (in prep.)* within Fig. 3.8, in keeping with Campbell et al. (2015).

3.3.4 VIMOS Public Extragalactic Redshift Survey

The VIMOS Public Extragalactic Redshift Survey (VIPERS) consists of spectroscopic observations of galaxies selected using Canada-France-Hawaii Telescope Legacy Survey (CFHTLS) photometry (Guzzo et al., 2014). Marulli et al. (2013) present galaxy clustering as a function of stellar mass in the first VIPERS data release (Garilli et al., 2014). This dataset has sky coverage of about 11 square degrees, with spectroscopic completeness of roughly 40 percent. The redshift range sampled is $0.5 < z < 1.2$, and the magnitude limit is as given in Table 3.1. A selection in colour is used to exclude galaxies with $z < 0.5$, which is not completely efficient at selecting galaxies towards the lower limit of the surveyed redshift range, i.e. the selection does not correspond exactly to a step function at $z = 0.5$ (Garilli et al., 2014). However, in practice this has little impact on the number of galaxies recovered across the full redshift range, and is not important for our comparisons, according to tests with the model galaxies (see §3.4).

Marulli et al. use broad-band SED fitting to estimate the stellar masses of their sample of VIPERS galaxies, as described by Davidzon et al. (2013). The SPS model of Bruzual & Charlot (2003) is used, with a Chabrier (2003) IMF; i.e. the same choices as made by Taylor et al. (2011) for GAMA (see §3.3.3 above). Both the Calzetti et al. (2000) and Prevot-Bouchet (Prevot et al., 1984; Bouchet et al., 1985) dust attenuation laws are used,[7] with the best-fitting option being chosen for each galaxy. Both exponentially declining and constant SFHs are used in the fitting. The photometric bands used are: the CFHTLS u, g, r, i, and z; the Wide-Field Infra-Red Camera (WIRCAM; Puget et al. 2004) K; the Galaxy Evolution Explorer (GALEX; Gil de Paz et al. 2007) far-UV and near-UV; the UKIRT Infra-Red Telescope Deep Sky Survey (UKIDSS; Lawrence et al. 2007) Y, J, H, and K; and the Spitzer Wide-Area Infra-Red Extragalactic Survey (SWIRE; Lonsdale et al. 2003) 3.6 and 4.5 μm.

Samples selected in redshift and stellar mass are used by Marulli et al. to compute the correlation function. The clustering strength increases with stellar mass in each redshift range. We compare the model clustering predictions to the lower redshift interval considered by Marulli et al. (see Table 3.1), for stellar masses $M_\star > 10^{9.5}\, h^{-2}\, \mathrm{M}_\odot$.

[7] The Prevot-Bouchet law results from modelling the dust attenuation of the Small Magellanic Cloud, while the Calzetti law was calibrated using a sample of starburst galaxies.

3.4 Analysis

Our analysis consists of two distinct components: (i) computation of the real-space clustering as a function of stellar mass in the Gon14 and Lac14 GALFORM models at redshifts of 0.089, 0.32, and 0.62, and (ii) comparison of the projected model clustering as a function of stellar mass to observational data at redshifts of 0.089, 0.17, and 0.62 (i.e. the MS-W7 snapshots closest to the median redshifts of the observed galaxy samples; see Table 3.1). In (i), the lower and upper redshifts are the same as those in (ii), for comparing to SDSS and VIPERS data respectively; however the intermediate redshift (0.32) is chosen to be roughly evenly spaced in lookback time between these redshifts, covering 1.2 to 5.9 Gyr, rather than using the redshift of the comparison to GAMA (0.17). In (ii), we carry out SED fitting to obtain stellar mass estimates for the model galaxies, to allow a more robust comparison to the observations. The predictions obtained using the new subhalo dynamical friction merger scheme are considered in both parts of our analysis.

For each model galaxy, GALFORM outputs the true stellar mass, real-space coordinates, peculiar velocity, and photometry including dust attenuation (see §3.2.1). Redshift-space coordinates are computed by taking the line of sight as the third Cartesian axis and taking into account the peculiar motions along this axis (this is the distant observer approximation for projected clustering). We convert the stellar masses predicted by GALFORM from units of $h^{-1}\,\mathrm{M}_{\odot}$ to $h^{-2}\,\mathrm{M}_{\odot}$ using $h = 0.704$ (see §3.2.2), to be consistent with the mass units of the observational data.

Apparent magnitude limits were imposed to match those of each survey, at certain stages in our analysis that will be indicated in §3.4.1.2 and §3.4.2.2. These made use of SDSS r-band filter wavelength response data to match GAMA, and CFHTLS i-band filter response data to match VIPERS (see Table 3.1). In matching SDSS, we did not impose (apparent or absolute; cf. Table 3.1) magnitude limits, as the clustering results of Li et al. (2006) correspond to galaxy samples that are volume limited (see §3.3.2).

VIPERS uses a selection in colour to exclude galaxies with $z < 0.5$, as noted in §3.3.4, which leads to incomplete sampling of galaxies close to $z = 0.5$. Guzzo et al. (2014) note that the erroneous exclusion of galaxies ceases for $z \gtrsim 0.6$. We have verified that this selection is indeed unimportant

for the models by $z = 0.62$.

This section describes our adjustments to the model galaxy stellar masses (§3.4.1), and the computation of the correlation functions used to describe their clustering (§3.4.2).

3.4.1 Stellar masses

When computing clustering predictions to compare to observations, we considered three sets of model stellar masses: (i) the true stellar masses as predicted by GALFORM, (ii) estimates of the masses from SED fitting to the model photometry, and (iii) masses resulting from abundance matching to the stellar mass functions reported for the observed galaxies (as an adjustment following the SED fitting). In particular, we present the clustering results of the models with the standard merger scheme using (i), (ii), and (iii), and the results of the new subhalo dynamical friction merger scheme using (iii). The determination of (ii) and (iii) will now be described.

3.4.1.1 SED fitting

Mitchell et al. (2013) implemented SED fitting (see §3.3.1) to estimate stellar masses for model galaxies, using broad-band photometry. These stellar mass estimates, when compared to the true values calculated in the model, can be used to investigate the influences of the various assumptions that are required in SED fitting (e.g. differences in the choice of IMF, the SFH, recycling of stellar mass back into the interstellar medium, the metallicities available in the SPS models, and dust attenuation), on the derived properties of the galactic population.

Following Mitchell et al., we carried out SED fitting for both GALFORM models, at each redshift, in order to obtain estimates of the stellar masses of the model galaxies. As noted in §3.3, the GAMA and VIPERS stellar masses are themselves derived through broad-band SED fitting (the GAMA stellar masses use an empirical formula based on this; Taylor et al. 2011), so our intention here is to carry out equivalent fitting procedures for the model galaxies. The SDSS stellar masses were derived by fitting to particular spectral features, rather than to broad-band photometry (see §3.3.2); we nonetheless carried out broad-band SED fitting for this comparison.

In all cases the Bruzual & Charlot (2003) SPS model was assumed, with a Chabrier (2003) IMF, exponentially declining SFHs, and the Calzetti et al. (2000) dust extinction law. The metallicities used to compute the template SEDs were matched to those used in the GAMA and VIPERS SED fitting by Taylor et al. (2011) and Davidzon et al. (2013) respectively, at the relevant redshifts. Taylor et al. use the full native metallicity grid of the Bruzual & Charlot SPS model, i.e. $Z \in \{0.0001, 0.0004, 0.004, 0.008, 0.02, 0.05\}$, while Davidzon et al. use $Z \in \{0.004, 0.02\}$. To compare to SDSS, the full set of metallicity values was employed in the fitting, as in the comparison to GAMA. Each template SED had a non-evolving metallicity, and we did not interpolate between the available SPS model metallicities, in keeping with the typical observational SED fitting procedures. All the photometric bands listed in §3.3.3 and §3.3.4 were used in the SED fitting to match GAMA and VIPERS respectively, with the exception of the UKIDSS K-band, which was only included by Davidzon et al. in the absence of WIRCAM K-band magnitudes for VIPERS. The standard SDSS u, g, r, i, and z filter set was used in the fitting to match SDSS. Taylor et al. compute likelihood-weighted stellar masses (see §3.3.3). We conformed to this for the fitting to match GAMA, but used the more standard approach of selecting the SED at the mode of the likelihood distribution for the fitting to match SDSS and VIPERS.

3.4.1.2 Schechter function matching

The stellar mass function of galaxies, $\Phi(M_\star)$, is defined such that the number of galaxies per unit volume with stellar mass in the range $\log_{10}(M_\star/\mathrm{M}_\odot)$ to $\log_{10}(M_\star/\mathrm{M}_\odot) + \mathrm{d}\log_{10}(M_\star/\mathrm{M}_\odot)$ is $\Phi(M_\star)\,\mathrm{d}\log_{10}(M_\star/\mathrm{M}_\odot)$. The stellar mass function is conventionally modelled using a Schechter (1976) function, defined such that

$$\Phi(M_\star)\,\mathrm{d}\log_{10}\left(\frac{M_\star}{\mathrm{M}_\odot}\right) = \ln(10)\tilde{\Phi}\left(\frac{M_\star}{\tilde{M}_\star}\right)^{\alpha+1}\exp\left(-\frac{M_\star}{\tilde{M}_\star}\right)\mathrm{d}\log_{10}\left(\frac{M_\star}{\mathrm{M}_\odot}\right), \quad (3.2)$$

where $\tilde{M}_\star$ is a characteristic mass, α is a power law slope, and $\tilde{\Phi}$ sets the normalisation of the mass function.[8]

Differences between the numbers of model and observed galaxies in a

[8]The factor of ln(10) in equation (3.2) arises since we could instead have defined the stellar mass function as $\phi(M_\star)$ such that $\phi(M_\star)\mathrm{d}M_\star = \Phi(M_\star)\,\mathrm{d}\log_{10}(M_\star/\mathrm{M}_\odot)$. In that case,

given mass interval may give rise to discrepancies in the clustering results for samples selected by stellar mass, even if the underlying clustering signal is identical. Such mass function differences are dependent on the details of the model physics, combined with the procedure for estimating stellar masses through SED fitting. In order to eliminate any differences between the model and observationally inferred stellar mass functions, the Schechter functions representing the stellar mass functions of the galaxies in each observational sample (see Table 3.1) were imposed on the model galaxies at the corresponding redshifts. This process is equivalent to rescaling the estimated stellar mass of each GALFORM galaxy, in order to make the model mass functions match the observational results. The procedure used was to match the shape and normalisation of the target Schechter function, while maintaining the ordering of the model galaxies in estimated stellar mass from SED fitting, as follows. Given the lowest stellar mass of interest (for the clustering samples, see §3.3), and the Schechter function fitted to the measurement for the observed galaxies, we generated the expected number of galaxies in the simulation volume by randomly sampling stellar masses consistent with this mass function. The generated masses were assigned in order to the GALFORM galaxies, such that the highest generated mass was ascribed to the GALFORM galaxy with the highest mass estimate from SED fitting, and so on. It is important to note that the galaxy formation models considered here have not been calibrated to reproduce observationally inferred stellar mass function data. They have however, been calibrated to match the local K-band luminosity function.

Measuring $\Phi(M_\star)$ requires one to know the number of galaxies within some stellar mass range of interest, in some known volume. Difficulties in achieving this arise from the fact that galaxy surveys are defined by apparent magnitude limits; that is, a sample of survey galaxies defined by stellar mass limits will inherently also be restricted to some range in apparent magnitude. It is, however, still possible to identify samples of galaxies that

$\Phi(M_\star) = \ln(10) M_\star \phi(M_\star)$. Therefore,

$$\phi(M_\star)\,\mathrm{d}M_\star = \tilde{\Phi}\left(\frac{M_\star}{\tilde{M}_\star}\right)^{\alpha} \exp\left(-\frac{M_\star}{\tilde{M}_\star}\right)\frac{\mathrm{d}M_\star}{\tilde{M}_\star}\,, \tag{3.3}$$

and thus we can see that $\tilde{\Phi}$ normalises the Schechter function, without any other scaling prefactor, when written in terms of $\phi(M_\star)$.

are complete, i.e. volume limited, in stellar mass (e.g. the clustering samples of Li et al. 2006). The level of incompleteness in stellar mass, i.e. the fraction of missing galaxies, in a sample defined by stellar mass, depends on the apparent magnitude limits of the survey and the redshift range of interest.

In the typical case, where we are concerned with a faint apparent magnitude limit, it is useful to be able to estimate a lower mass threshold above which the measured mass function can be considered to be 'reliable' (i.e. the same as what would have been measured with no magnitude limit, for a volume limited sample). Pozzetti et al. (2010) describe the method used by Davidzon et al. (2013) to estimate such a threshold mass for VIPERS. At our comparison redshift (see Table 3.1), this VIPERS threshold value is approximately $10^{9.6}\,h^{-2}\,\mathrm{M}_{\odot}$, which corresponds to the mass below which the faint limit causes more than about 20 percent of galaxies to be missed, in a given mass interval. The VIPERS Schechter function specified in Table 3.1 has been fitted by Davidzon et al. only for masses above the reliability threshold defined in this way. Clearly the i-band faint limit has an important influence on the completeness of the stellar mass function at this redshift, for the lowest masses of interest (note that the minimum mass we consider for the clustering calculations using these galaxies is $10^{9.5}\,h^{-2}\,\mathrm{M}_{\odot}$). In light of this, we imposed the i-band faint limit on the model galaxies *before* matching the VIPERS mass function at this redshift, in order to reproduce the 'underestimated' mass function measurement at lower masses, i.e. the incomplete mass function of the galaxies actually used in the clustering analysis of Marulli et al. (2013).

The double Schechter function representing the GAMA stellar mass function has been fitted for $z < 0.06$ (Baldry et al., 2012), whereas the clustering results of Farrow et al. (2015) considered here use galaxies with $0.14 < z < 0.24$ (see Table 3.1). Baldry et al. use the SED fitting procedure of Taylor et al. (2011) to estimate stellar masses, supplemented by the corresponding empirical relation of Taylor et al. for a small number of galaxies with missing photometry (see §3.3.3). Baldry et al. compare their stellar mass function to that obtained by Pozzetti et al. (2010) using the Cosmic Evolution Survey (zCOSMOS; Lilly et al. 2007) for $0.1 < z < 0.35$, using similar SED fitting for the stellar mass estimation, and find good agreement between the two measurements. Furthermore, we have verified

that the stellar mass function does not evolve significantly when using the sample of galaxies considered by Farrow et al. ($0.14 < z < 0.24$), with respect to the Baldry et al. measurement. The Baldry et al. mass function fit can be regarded as complete for the masses of interest here, thanks to being constrained at low redshift (the fitting considered the mass function as measured for $M_\star > 10^{7.7}\,h^{-2}\,\mathrm{M}_\odot$, while the clustering data we compare to are for $M_\star > 10^{9.5}\,h^{-2}\,\mathrm{M}_\odot$). Thus, we matched the mass functions of the GALFORM models to the double Schechter fit, while keeping all the model galaxies, and only *afterwards* imposed the r-band faint limit on the models (see Table 3.1). In this way, the complete model stellar mass functions were made to reproduce the complete observationally inferred mass function, before introducing the relative incompleteness due to the particular selection relevant for comparison to the clustering data of Farrow et al.

The SDSS mass function Schechter fit was derived directly from the second SDSS data release (Abazajian et al., 2004; Wang et al., 2006), using r-band apparent magnitude and redshift limits the same as for the clustering data of Li et al. (2006), but without the additional absolute magnitude limits imposed for the clustering samples (see Table 3.1). We assumed the Wang et al. fit to be approximately complete in stellar mass, and thus imposed it upon the model galaxies without applying any magnitude limiting, as required for comparison to the volume limited clustering samples of Li et al. (see §3.3.2). It should be noted, however, that the measured mass function is likely to be significantly suppressed by the SDSS faint r-band limit for $M_\star \lesssim 10^{8.2}\,h^{-2}\,\mathrm{M}_\odot$ (Baldry et al., 2008; Li & White, 2009). This incompleteness has minor implications for our comparison to the Wang et al. mass function data and Schechter fit (for $M_\star > 10^{8.69}\,h^{-2}\,\mathrm{M}_\odot$) in §3.5.2, but is not important for our comparison to the volume limited clustering data; except where the mass function is imposed, in which case any errors introduced by the assumption of completeness in the Schechter fit should only have an impact on the lowest masses considered.

In §3.5.2.2 and §3.5.2.3 we make use of the mass function matching to compare the results of the two GALFORM merger schemes to the observational clustering data (as outlined at the beginning of §3.4.1). In this way, we force the models using either scheme to reproduce the same mass function, and thus the differences in the clustering results using the two schemes, when

compared in this way, are not due to any differences in the predicted stellar mass function that are introduced by changing to the new merger scheme.

3.4.2 Clustering

We now define the correlation function, $\xi(r)$, and the projected correlation function, $w_{\rm p}(\sigma)$, and give details of their computation.

3.4.2.1 Correlation function

For a cosmologically representative volume, V, the probability, $\mathrm{d}P$, of finding a galaxy in some volume element, $\mathrm{d}V_1$, at a comoving distance of r_{12} from another galaxy in another volume element, $\mathrm{d}V_2$, defines the spatial two-point autocorrelation function, $\xi(r)$, such that

$$\mathrm{d}P = n_V^2[1 + \xi(r_{12})]\,\mathrm{d}V_1\,\mathrm{d}V_2\,, \tag{3.4}$$

where n_V is the mean number density of galaxies within V (Peebles, 1980). Thus, following Rivolo (1986),

$$1 + \xi(r) = \frac{1}{n_V^2 V}\frac{N(r)}{V_{\rm s}(r)}\,, \tag{3.5}$$

where $N(r)$ is the number of pairs with separation r, computed by summing over the number of pairs involving each galaxy in the volume individually (so this is twice the number of independent pairs), considering a spherical shell of radius r and volume $V_{\rm s}(r)$. A simple appreciation of the uncertainty in $\xi(r)$ can be gained from considering Poisson statistics (Iovino & Shaver, 1988; Martinez et al., 1993). The error on $\xi(r)$ is then $\Delta\xi(r)$ such that

$$\Delta\xi(r) = \left(\frac{2}{n_V^2 V}\frac{1 + \xi(r)}{V_{\rm s}(r)}\right)^{1/2}\,. \tag{3.6}$$

It is common to fit a power law to $\xi(r)$, parameterised by a correlation length, r_0, and slope, γ, (Peebles, 1980), i.e.

$$\xi(r) = \left(\frac{r}{r_0}\right)^{-\gamma}\,. \tag{3.7}$$

Here r_0 characterises the clustering 'strength', where $\xi(r_0) = 1$. This parameterisation is not suitable for describing the precise clustering measurements that are possible today over a wide range of scales, but we can make use of it to describe the galaxy clustering over small ranges of pair separations.

$\xi(r)$ was computed for the models at redshifts $z \in \{0.089, 0.32, 0.62\}$ (using the true stellar masses, with no magnitude limits), with and without the new merger scheme, in three equally spaced bins in stellar mass. Pair separations in the range $0.1 < r\ [h^{-1}\,\mathrm{Mpc}] < 30$ (the choice of the large-scale limit will be discussed in §3.4.2.3) were divided into 30 bins of equal logarithmic width. $\xi(r)$ was computed using the pair counts binned in r, following equation (3.5), where $V = L_{\mathrm{box}}^3$ (see §3.2.2).

The power law given in equation (3.7) was fitted to each $\xi(r)$, using the $\Delta\xi(r)$ from equation (3.6) to weight the fit. We preferred to fit over a relatively small range in r, where $\xi(r)$ is close to being an exact power law, in the neighbourhood of $\xi(r) = 1$, such that r_0 relates closely to the 'true' correlation length. In this way, bins within $3 < r\ [h^{-1}\,\mathrm{Mpc}] < 10$ were used to fit the power law for each $\xi(r)$. This range in r samples the two-halo term in the correlation function, i.e. the separations considered relate to galaxies in different haloes. As such, the fitted power laws are insensitive to clustering on small scales (one-halo term), and thus we consider them using only the standard merger scheme.

Our results for $\xi(r)$ using the two GALFORM models, comparing the two merger schemes, are presented in §3.5.1.

3.4.2.2 Projected correlation function

The distance to a galaxy moving with the cosmic expansion can be inferred through measuring its recession velocity (i.e. redshift) and using Hubble's law. If the galaxy has some peculiar velocity relative to the Hubble flow, this will perturb the determined distance. As such, it is necessary to distinguish between real-space and redshift-space coordinates, i.e. true spatial positions, and those determined through measuring redshifts, respectively (Kaiser, 1987).

For observed galaxies, the redshift-space correlation function may be computed. This encodes information about the peculiar motions of galaxies on different separation scales. Through considering the components of the

pair separation orthogonal and parallel to the line of sight (σ and π respectively), we can define the two-dimensional correlation function $\xi(\sigma,\pi)$,[9] by analogy with equation (3.5). We may integrate $\xi(\sigma,\pi)$ along the line of sight, defining the projected correlation function, $w_{\rm p}(\sigma)$, (Davis & Peebles, 1983):

$$w_{\rm p}(\sigma) = \int_{-\pi_{\rm max}}^{\pi_{\rm max}} \xi(\sigma,\pi)\,{\rm d}\pi\ . \tag{3.8}$$

This statistic is traditionally used to describe the real-space clustering of galaxies drawn from observational surveys, as redshift-space distortions do not influence $w_{\rm p}(\sigma)$. This is true provided that $\pi_{\rm max}$ is sufficiently large, such that the integral includes all correlated galaxy pairs, encompassing their redshift-distorted coordinates (Davis & Peebles, 1983; Norberg et al., 2009). Norberg et al. show using simulations that $w_{\rm p}(\sigma)$ is sensitive to the choice of $\pi_{\rm max}$, with differences between the projected clustering recovered on large scales for $\pi_{\rm max}$ $[h^{-1}\,{\rm Mpc}] \in \{30, 64\}$. They note that using $\pi_{\rm max} = 64\,h^{-1}\,{\rm Mpc}$ (cf. Table 3.1) results in a difference of ten percent between real and redshift space for $\sigma \sim 10\,h^{-1}\,{\rm Mpc}$, rising to 50 percent by $\sigma \sim 30\,h^{-1}\,{\rm Mpc}$.

Assuming equation (3.7) for $\xi(r)$, and infinite $\pi_{\rm max}$, the corresponding power law for $w_{\rm p}(\sigma)$ is given by (Davis & Peebles, 1983)

$$w_{\rm p}(\sigma) = \sigma^{1-\gamma} r_0^{\gamma} \frac{\Gamma(1/2)\Gamma([\gamma-1]/2)}{\Gamma(\gamma/2)}\ , \tag{3.9}$$

where Γ is the Gamma function. Equation (3.9) is used in this work for illustrative purposes only.

$w_{\rm p}(\sigma)$ was computed in real and redshift space for the two GALFORM models at each redshift $z \in \{0.089, 0.17, 0.62\}$, for each of the four variants identified in §3.4.1, i.e. the default models with their true and estimated masses, and the models with masses resulting from matching the observational mass functions (using both the new and default merger schemes). At each redshift the appropriate magnitude limits were imposed (see the beginning of §3.4). The stellar mass binning was matched to that of each survey. Separations

[9] Throughout this chapter, we reserve $\xi(r)$ to denote the spherically averaged two-point correlation function computed using real-space pair separations, r. We use (σ, π) to refer to the pair separation components in *either* real or redshift space, and state whether $w_{\rm p}(\sigma)$ has been computed from real or redshift-space coordinates (i.e. real or redshift-space values of the line-of-sight separation, π).

orthogonal to the line of sight were computed with the same binning as the r values for $\xi(r)$ (see §3.4.2.1). Line-of-sight separations were measured in bins of fixed width, $\Delta\pi$, using the distant observer approximation (as noted at the beginning of §3.4). $\xi(\sigma, \pi)$ was evaluated for each bin of the σ versus π grid, using equation (3.5), where the volume (V_s) considered was instead that of a cylindrical shell of inner and outer radii set by the σ bin edges and length $\Delta\pi$. Discretising equation (3.8), $w_p(\sigma) = \sum_{\pi_i} \xi(\sigma, \pi_i)\Delta\pi$ was computed for bin centres π_i, such that $|\pi_i| < \pi_{max}$ ($30\,h^{-1}$ Mpc; see §3.4.2.3 below).

The results of our comparison of the model clustering predictions to the observationally inferred $w_p(\sigma)$ data are presented in §3.5.2.

3.4.2.3 The limit on separation scales

The finite size of the simulation volume sets an upper limit on the separation scales where the galaxy clustering can be considered to be reliable, with respect to that which would be computed using an arbitrarily large cosmological volume.

To establish the pair separations where the box size becomes important, we computed the projected correlation function along three mutually orthogonal lines of sight. This was carried out for each redshift and stellar mass range of interest in this chapter, in both real and redshift space.

We found the scatter in $w_p(\sigma)$ for $\sigma < 30\,h^{-1}$ Mpc, around the mean for the three projections, to be about 0.01 dex. For $\sigma \gtrsim 30\,h^{-1}$ Mpc, the scatter rose sharply with increasing σ, reaching roughly 0.1 dex by $\sigma \sim 50\,h^{-1}$ Mpc. These results for the $w_p(\sigma)$ variation across mutually orthogonal lines of sight are largely insensitive to the choice of line-of-sight integration limit, over the range $30 < \pi_{max}\ [h^{-1}\,\mathrm{Mpc}] < 50$.

Based on this test, we used an upper limit of $30\,h^{-1}$ Mpc for σ when computing $w_p(\sigma)$, and for r when computing $\xi(r)$. $\pi_{max} = 30\,h^{-1}$ Mpc was used when comparing to each survey. We have checked that using the particular π_{max} values given in Table 3.1 for each survey instead (which are larger for the SDSS and GAMA measurements) does not have a significant impact on our comparisons. Our upper limit to pair separations is the same as that noted by Orsi et al. (2008) as the largest separation where the galactic and dark matter $\xi(r)$ are related by a scale-independent bias within the original Millennium Simulation volume. Gonzalez-Perez et al. (2011)

also found this scale to be the upper limit for describing the redshift-space correlation function boost with the Kaiser (1987) formalism, due to the finite size of the simulation volume.

3.5 Results and discussion

Our results are now presented for the real-space $\xi(r)$ computed using the GALFORM models as a function of true stellar mass and redshift (§3.5.1), and the comparisons to the $w_{\rm p}(\sigma)$ observational clustering estimates (§3.5.2).

3.5.1 Predicted real-space clustering

We first consider the predictions of the default models (§3.5.1.1), and then examine the influence of the choice of satellite merger scheme on the small-scale clustering of galaxies (§3.5.1.2).

3.5.1.1 Standard models

The galactic stellar mass function, $\Phi(M_\star)$, is shown as a function of redshift in Fig. 3.1, as predicted using the Gon14 and Lac14 GALFORM models. Fig. 3.2 shows the real-space $\xi(r)$ computed as a function of stellar mass and redshift. Fig. 3.3 shows the power law parameters, r_0 and γ, fitted to each $\xi(r)$ bin.

As shown in Fig. 3.1, the evolution of the shape of $\Phi(M_\star)$ as a function of redshift is fairly similar between the models. The number density above the knee of the mass function decreases with increasing redshift in each model. There are more high-mass galaxies in the Gon14 model than in the Lac14 model (this is similar to the differences between the predicted K-band luminosity functions). This regime roughly corresponds to the highest mass bin for the correlation functions shown in Fig. 3.2 (as indicated in Fig. 3.1).

For clarity, we divide the predicted correlation functions by a reference power law (see Fig. 3.2 caption). The model clustering predictions shown in Fig. 3.2 are close to a power law over only a limited range of scales. There is a clear transition between the one-halo term on small scales and the two-halo term on large scales (at $r \sim 2h^{-1}\,\mathrm{Mpc}$). The two-halo term has the same shape as the large-scale dark matter $\xi(r)$ from linear theory, where the galaxy clustering bias with respect to the dark matter increases with both redshift

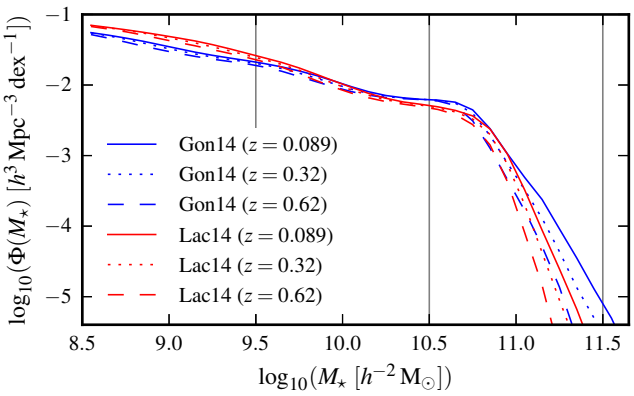

Figure 3.1 – Galactic stellar mass function, $\Phi(M_\star)$, predicted by the Gon14 and Lac14 GALFORM models, as a function of redshift. The vertical lines indicate the stellar mass samples used to measure $\xi(r)$ in Fig. 3.2. The lower $\Phi(M_\star)$ axis limit corresponds to 50 model galaxies per bin in the MS-W7 simulation volume.

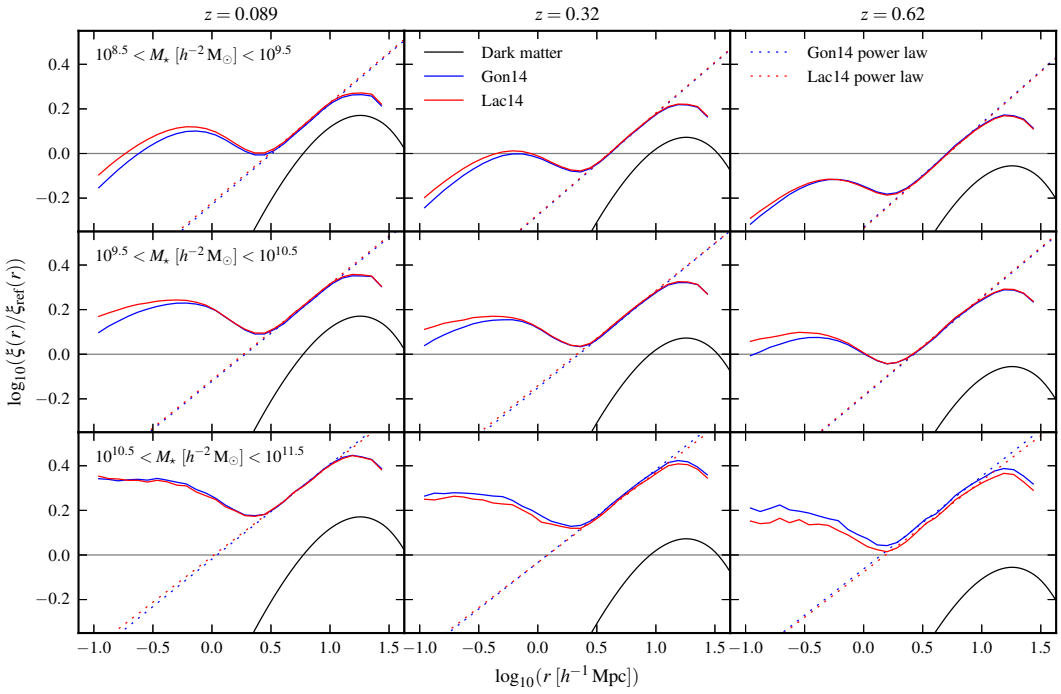

Figure 3.2 – Correlation function of galaxies, $\xi(r)$, as a function of redshift, z (columns), and stellar mass, $M_\star$ (rows), as predicted using the Gon14 and Lac14 GALFORM models. For clarity, each $\xi(r)$ has been divided by a reference power law, $\xi_{\mathrm{ref}}(r)$, with parameters $r_0 = 5\,h^{-1}$ Mpc and $\gamma = 2$ (horizontal line; see equation 3.7). The dotted lines are power law fits to each GALFORM $\xi(r)$, for $3 < r\ [h^{-1}\,\mathrm{Mpc}] < 10$, with the parameters shown in Fig. 3.3. The black curves show $\xi(r)$ for the dark matter at each redshift, computed as the Fourier transform of the MS-W7 linear theory power spectrum (see equation 1.15).

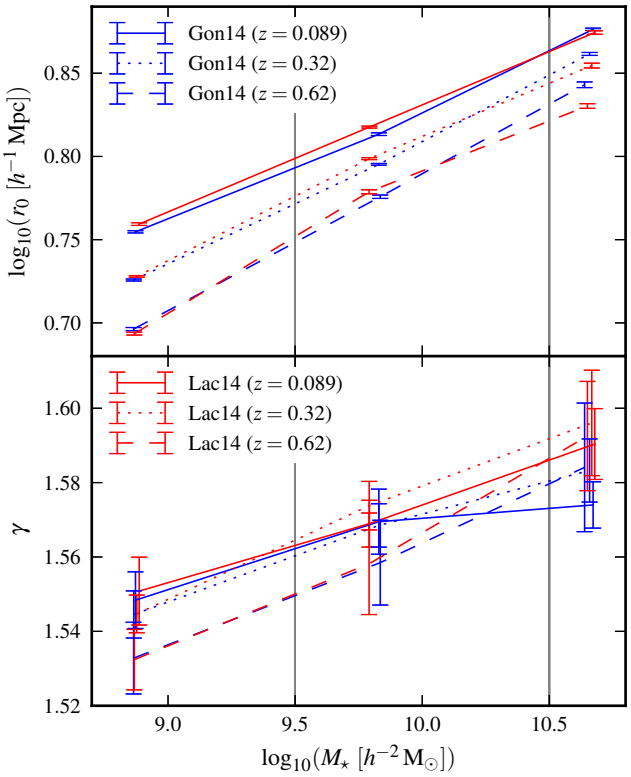

Figure 3.3 – Correlation lengths, r_0 (upper), and slopes, γ (lower), of the power law fitted to the $\xi(r)$ shown in Fig. 3.2 for the Gon14 and Lac14 GALFORM models, as a function of stellar mass, $M_\star$, and redshift, z, (see equation 3.7). The vertical lines indicate the divisions between the mass bins used to measure $\xi(r)$. The $M_\star$ values are the medians for each mass bin, model, and redshift. Each panel uses the same linestyles and colours; the legend is split between panels.

and stellar mass. The clustering in the two models is indistinguishable on large scales (i.e. for pair counts between haloes), but there are differences between the model predictions on smaller scales ($r \lesssim 1\,h^{-1}\,\mathrm{Mpc}$).

There is a general trend of clustering amplitude increasing with stellar mass, on all separation scales, at each redshift. Slight increases in the amplitude of the clustering on large scales are seen with decreasing redshift, for a fixed mass bin. The power law fits shown in Fig. 3.2 ignore clustering on small and very large scales. The fit parameters, as shown in Fig. 3.3, indicate increasing clustering strength, r_0, with both increasing stellar mass and decreasing redshift, where the differences between redshifts become less significant for higher stellar masses. The slope, γ, is fairly constant, but exhibits a weak increase with increasing stellar mass. The redshift evolution of the clustering on small scales in Fig. 3.2 (which is sensitive to pair separations within haloes) is particularly dramatic for low stellar masses.

3.5.1.2 New satellite merger scheme

In the new merger scheme, satellite galaxies track the positions of their associated subhaloes, only making use of an analytically computed merger time-scale once the subhalo can no longer be resolved in the simulation, as described in §3.2.4. Fig. 3.4 shows the stellar mass functions of the Gon14 and Lac14 models at redshift $z = 0.089$, with and without the new merger scheme, decomposed into the contributions from central and satellite galaxies. The same trends are seen at higher redshifts. Changes in the total $\Phi(M_\star)$ due to the new merger scheme are significant for the highest masses (top panel of Fig. 3.4), where for both models the amplitude of the high-mass end increases when switching to the new scheme. Note that the model parameters have not been recalibrated on adopting the new merger scheme. The rise in the galaxy abundance in the Lac14 model at these masses is greater than for the Gon14 model, such that with the new scheme, the two model mass functions become very similar for high masses. The increases in the numbers of the most massive galaxies are due to more mergers taking place in the new scheme, which thus transfer more mass to central galaxies, as can be seen in the middle panel of Fig. 3.4. There is a smaller rise at high masses in the satellite mass function for both models, which again is more significant for the Lac14 model (lower panel). The source of the growth of

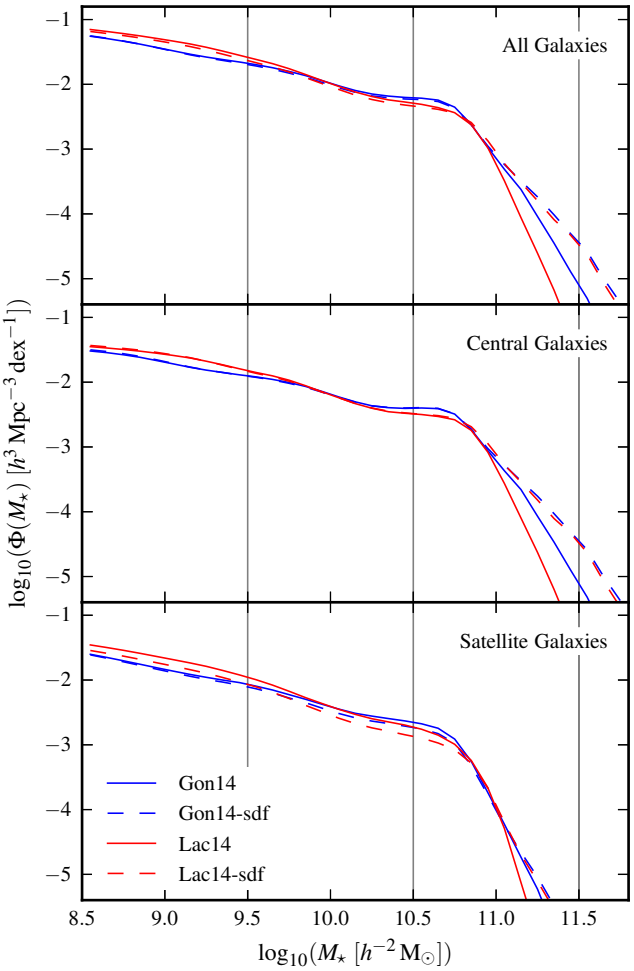

Figure 3.4 – Galactic stellar mass function, $\Phi(M_\star)$, at redshift $z = 0.089$, for all (top), central (middle), and satellite (bottom) galaxies in the Gon14 and Lac14 GALFORM models, and for these models with the alternative subhalo dynamical friction merger scheme (labelled Gon14-sdf and Lac14-sdf). The vertical lines indicate the stellar mass samples used to compute $\xi(r)$ in Fig. 3.5. The lower $\Phi(M_\star)$ axis limit corresponds to 50 model galaxies per bin in the MS-W7 simulation volume. Note that the model parameters have not been recalibrated on adopting the new satellite merger scheme.

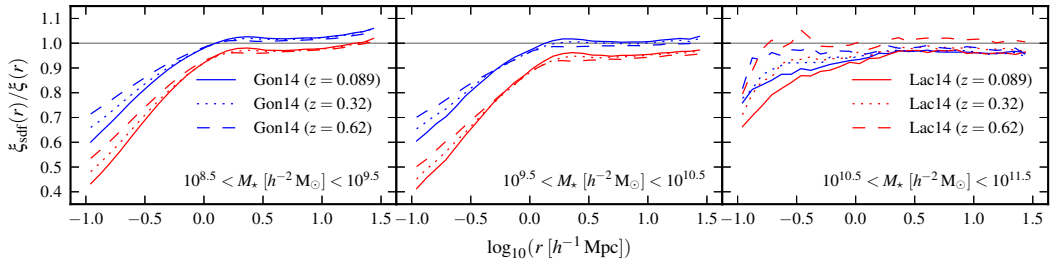

Figure 3.5 – Ratio of the real-space correlation function using the new subhalo dynamical friction merger scheme, $\xi_{\mathrm{sdf}}(r)$, to that obtained with the default satellite merger prescription, $\xi(r)$, computed using the Gon14 and Lac14 GALFORM models, as a function of stellar mass, $M_\star$, and redshift, z, (as labelled). The linestyles show different redshifts, and the colours indicate different models, as labelled. The legend is split across the panels.

higher mass galaxies here is the earlier assimilation of lower mass satellite galaxies. Mergers on to a central galaxy cease when it becomes a satellite of a more massive galaxy.

The changes in $\xi(r)$ for the two GALFORM models when using the new merger scheme are displayed in Fig. 3.5, at redshifts of 0.089, 0.32, and 0.62. The lines show the results with the subhalo dynamical friction merger scheme, divided by the predictions using the original merger scheme, for each model, stellar mass interval, and redshift. The clustering differences with respect to the standard merger scheme are larger for the lower redshifts. Generally there is little change in the clustering on large scales. There are significant decreases in the small-scale clustering amplitude ($r \lesssim 1\,h^{-1}\,\mathrm{Mpc}$) for all but the highest mass, highest redshift, data. The intermediate mass bin shows the largest decreases in $\xi(r)$ on small scales, although these are similar to the reductions in amplitude for the lower mass bin. There is significantly less change in the highest mass bin, relative to that in the other mass bins, at a given redshift. Changes in both the radial distribution and number of satellites contribute to the changes in the small-scale clustering when using the new scheme (see lower panel of Fig. 3.4, and §3.2.4).

3.5.2 Comparisons to observational data

We now present the results of our comparisons to observational measurements of the clustering of galaxy samples selected by stellar mass. First we consider the changes induced in the model predictions through stellar

mass estimation using SED fitting (§3.5.2.1), then discuss rescaling of the estimated model mass functions to reproduce observational mass function measurements (§3.5.2.2), and present the changes in the clustering predictions resulting from using the new subhalo dynamical friction merger scheme in GALFORM (§3.5.2.3). In §3.5.2.4 we comment on the need for SED fitting in order to make a fair comparison to observational clustering measurements, even when carrying out mass function abundance matching. Finally, we discuss some uncertainties relevant to these comparisons between theoretical predictions and the observational clustering results (§3.5.2.5).

3.5.2.1 True and estimated stellar masses

Mitchell et al. (2013) show that for some ranges in stellar mass, the process of SED fitting to the broad-band photometry of model galaxies, as predicted by GALFORM, introduces a roughly constant offset in estimated stellar mass with respect to true model masses, combined with associated scatter (see §3.4.1.1). However, for GALFORM galaxies with strong dust attenuation, SED fitting tends to significantly underestimate the stellar mass, leading to important systematic differences between the distributions of true and estimated masses, which cannot be described by a constant mean offset and scatter. This phenomenon is relevant for model galaxies with true stellar mass $\gtrsim 10^{10}\,h^{-2}\,\mathrm{M}_{\odot}$.

It is important to note that the IMF assumed in a galaxy formation model influences the mass-to-light ratios of the galaxies. The predicted luminosity functions are directly sensitive to the IMF, and are used to calibrate the models, which in turn affects the stellar masses. Within the machinery of SED fitting, the choice of IMF corresponds to a roughly constant multiplicative offset in the resulting stellar masses. The Chabrier (2003) IMF used in our SED fitting matches the GAMA and VIPERS stellar masses, but not the Kroupa (2001) IMF of the SDSS data (see §3.3 and §3.4.1.1). Therefore we multiplied the mass estimates by 1.1 following the SED fitting to match SDSS. This factor is that found by Davidzon et al. (2013) as the mean offset between the stellar masses derived using Chabrier and Kroupa IMFs in their SED fitting (see also Baldry et al. 2008).

The stellar mass functions of the Gon14 and Lac14 models are shown in Fig. 3.6, at redshifts of 0.089, 0.17, and 0.62, alongside the relevant

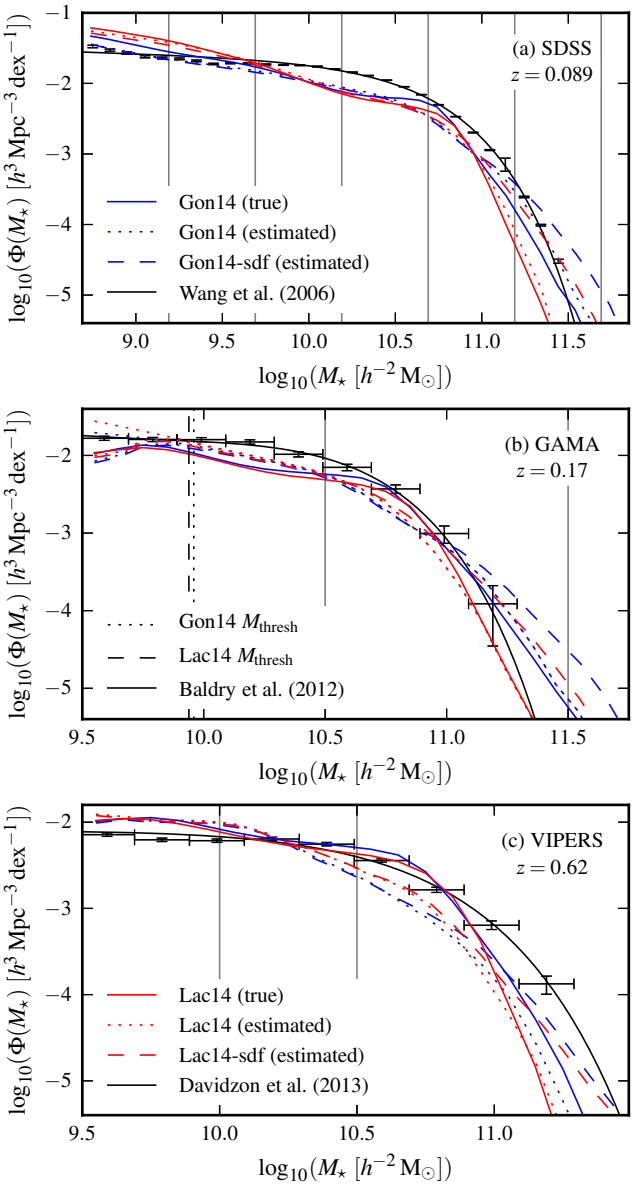

Figure 3.6 – Galactic stellar mass function, $\Phi(M_\star)$, as predicted using the Gon14 and Lac14 GALFORM models, at redshifts (a) 0.089, (b) 0.17, and (c) 0.62. Solid and dotted lines show the true and estimated $\Phi(M_\star)$ respectively for the default models. Dashed lines show the estimated $\Phi(M_\star)$ for the models using the new merger scheme (Gon14-sdf and Lac14-sdf). The formats of the model lines are the same in each plot, and are labelled in (a) and (c). The points with error bars are the measurements of (a) Wang et al. (2006) using SDSS, (b) Baldry et al. (2012) using GAMA, and (c) Davidzon et al. (2013) using VIPERS; shown alongside Schechter (1976) function fits (black lines; see Table 3.1). *continued...*

Figure 3.6 – *continued*... In (b) and (c), the apparent magnitude limits of Farrow et al. (2015) and Marulli et al. (2013) have been imposed on the model galaxies, respectively (see Table 3.1 and §3.4). The solid vertical grey lines in (a), (b), and (c) indicate the samples used to measure $w_{\rm p}(\sigma)$ in Figs. 3.7, 3.8, and 3.9 respectively. The lower $\Phi(M_\star)$ axis limit corresponds to 50 model galaxies per bin in the MS-W7 simulation volume. At $z = 0.17$, we compare to the $z < 0.06$ Baldry et al. measurement under the assumption that there is little evolution in $\Phi(M_\star)$ since $z = 0.17$ (see §3.4.1.2). The Baldry et al. $\Phi(M_\star)$ measurement is complete over the mass range shown in (b), and should not be compared to the magnitude limited model $\Phi(M_\star)$ at the lowest masses. The estimated $\Phi(M_\star)$ for the default models is repeated in (b) without magnitude limiting (with the same formatting), to show the impact of the faint limit at this redshift. To aid comparison, a corresponding threshold mass, $M_{\rm thresh}$, is indicated for each default model in (b) by a dotted or dashed vertical black line (see §3.5.2.1), as labelled in (b).

observationally inferred mass function data from SDSS, GAMA, and VIPERS, respectively. The true and estimated mass functions are shown for each model using solid and dotted lines, respectively (the dashed lines show the estimated mass functions using the new merger scheme, which will be discussed in §3.5.2.3). For the comparison to SDSS (Fig. 3.6a), no magnitude limits have been imposed on the model galaxies, because the clustering measurements of Li et al. (2006) are for samples that are complete in stellar mass (see §3.3.2). When comparing to GAMA and VIPERS (Fig. 3.6b,c), we imposed the apparent magnitude limits relevant for comparison to the clustering data of Farrow et al. (2015) and Marulli et al. (2013) respectively (see Table 3.1 and §3.4).

The Baldry et al. (2012) GAMA stellar mass function fit can be regarded as complete in stellar mass over the range of masses shown in Fig. 3.6b, as it was constrained for $z < 0.06$ (see §3.4.1.2). Therefore we also show the estimated model stellar mass functions *without* applying the GAMA magnitude limit in Fig. 3.6b, as these are appropriate to compare to the Baldry et al. data over the full range of masses shown (these lines are shown with the same formatting as the magnitude limited versions). Thus the differences due to the magnitude limit can be seen where the dotted lines of a given colour do not overlap (the lower, magnitude limited, dotted lines are close to the solid true model stellar mass lines at the lowest masses, following a reduction in amplitude of roughly 0.4 dex due to the magnitude limit). In this way, it can be seen that the faint r-band limit of GAMA

influences the completeness of the estimated model stellar mass functions for $M_\star \lesssim 10^{10}\,h^{-2}\,\mathrm{M}_\odot$ at $z = 0.17$. The vertical (dotted or dashed) black lines in Fig. 3.6b indicate a threshold stellar mass, M_{thresh}, for each model, above which the GAMA magnitude limit causes one percent of the galaxies in the complete model sample to be missed in the simulation volume, computed using the estimated stellar masses.

The true model stellar mass functions can change significantly when carrying out SED fitting. An effect that can be seen at each redshift is that the knee of the estimated mass function is smoothed out and suppressed with respect to the true mass function, for both models. The changes at the knee become larger with increasing redshift, and can be traced to the influence of dust attenuation (see Mitchell et al. 2013). Additionally, scatter in the estimated masses at a given true mass will cause $\Phi(M_\star)$ to rise for masses above this, particularly near the high-mass end, due to the rapid changes in number density as a function of stellar mass. This Eddington bias conspires with the systematic effects of dust attenuation to dictate the shape of the estimated mass function.

The projected correlation functions computed using the model galaxies, for samples selected by stellar mass, are shown in Figs. 3.7, 3.8, and 3.9, alongside the relevant observational measurements, for the comparisons to SDSS, GAMA, and VIPERS, respectively. In each case, column (a) shows the model clustering predictions when using the true stellar masses, and column (b) shows the results when using the estimated masses from SED fitting. These masses correspond to the true and estimated mass functions shown in Fig. 3.6, where the relevant $w_{\mathrm{p}}(\sigma)$ mass bin edges are indicated by vertical lines. The results shown in columns (c) and (d) of Figs. 3.7, 3.8, and 3.9 will be discussed in §3.5.2.2 and §3.5.2.3, respectively.

Switching from true to estimated masses can change the model clustering predictions significantly, comparing columns (a) and (b) in Figs. 3.7, 3.8, and 3.9. More massive haloes are more strongly clustered in hierarchical cosmologies (e.g. Cole & Kaiser 1989), and the stellar mass of a galaxy is closely related to its host halo mass (e.g. Guo et al. 2010). Therefore the changes in the clustering predictions when carrying out SED fitting can be understood in terms of model galaxies in dark matter haloes of certain masses being transferred across the stellar mass bin boundaries. The Gon14

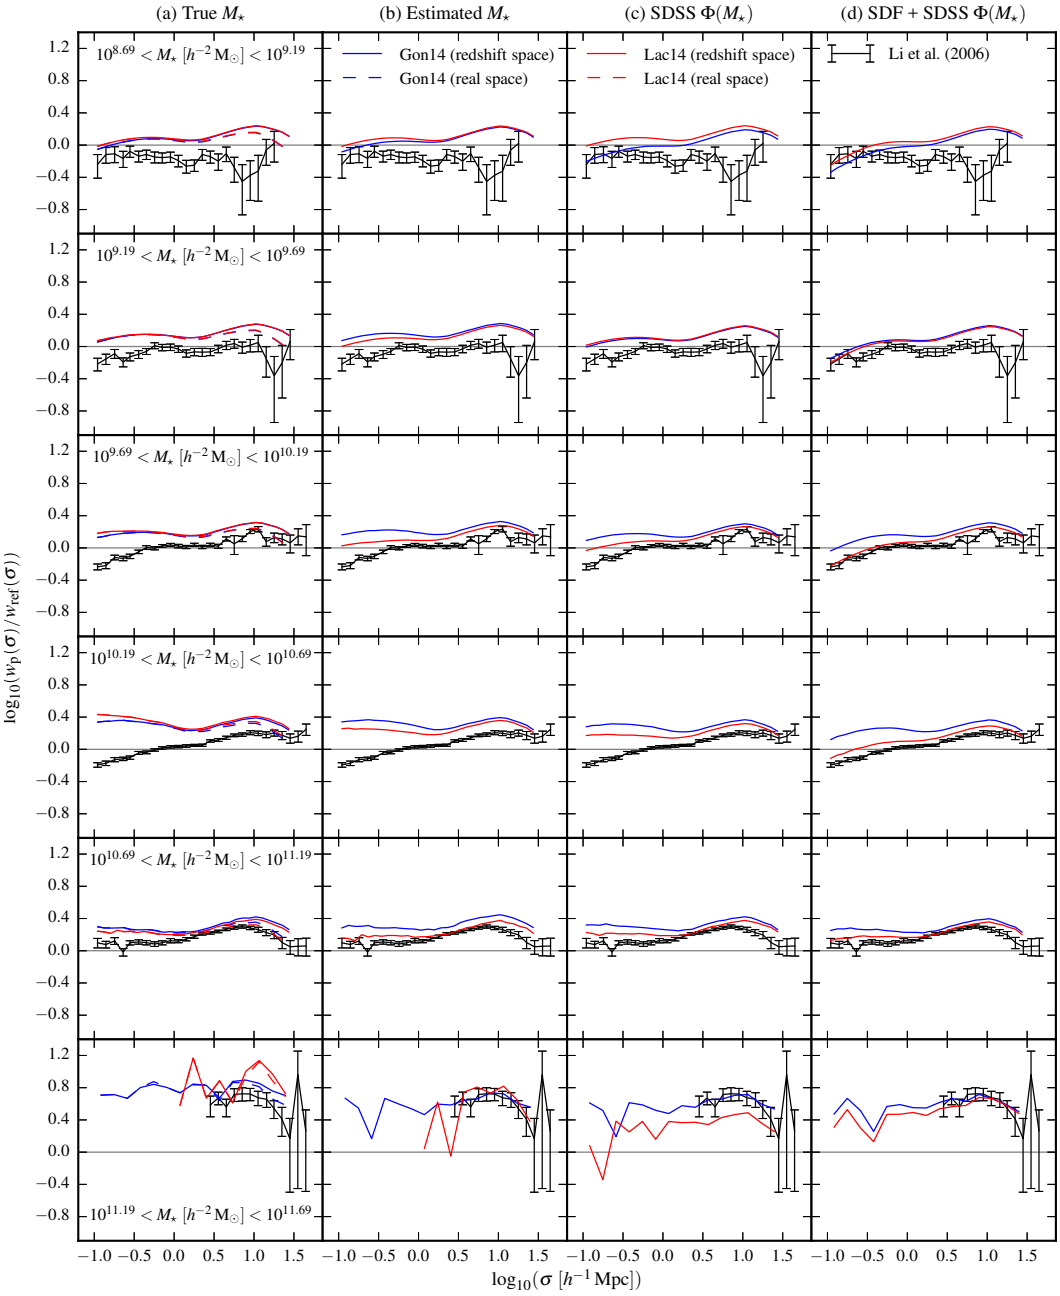

Figure 3.7 – Projected correlation function of galaxies, $w_\mathrm{p}(\sigma)$, as a function of stellar mass, $M_\star$ (as labelled in each row), at redshift $z = 0.089$, computed using the Gon14 and Lac14 GALFORM models, in real (first column only) and redshift space. The SDSS measurements of Li et al. (2006) are shown (black lines with errorbars). Each $w_\mathrm{p}(\sigma)$ has been divided by a reference power law, $w_\mathrm{ref}(\sigma)$, with parameters $r_0 = 5\,h^{-1}$ Mpc and $\gamma = 2$ (horizontal line; see equation 3.9). Half the standard number of σ bins are shown for the highest mass interval, due to the small number of galaxies. Columns (a) and (b) show the default models with the true and estimated masses, respectively. The remaining columns show the results of matching the SDSS stellar mass function (as an adjustment following the SED fitting), for (c) the default models, and (d) the models using the new subhalo dynamical friction (SDF) merger scheme.

and Lac14 models predict relatively similar clustering in most cases when using true stellar masses, but the SED fitting tends to significantly increase the differences between the results for the two models. The changes in the clustering predictions of the Lac14 model when switching to estimated masses are often larger than for the Gon14 model. It is likely that this is due to a combination of different factors, such as the different SPS models assumed and levels of dust extinction calculated in the two models. The SED fitting tends to decrease the clustering amplitude in the Lac14 model, but tends to increase that predicted in the Gon14 model. In general, switching to estimated masses has a larger impact on the clustering on small scales ($\sigma \lesssim 1\,h^{-1}\,\mathrm{Mpc}$), such that the shape of $w_\mathrm{p}(\sigma)$ is modified.

Comparing to SDSS, the agreement with the clustering measurements of Li et al. (2006) is improved or similar when switching from true to estimated stellar masses, for both GALFORM models (see Fig. 3.7a,b). For both models in most mass bins, the clustering amplitudes using true or estimated masses are higher over the range of separations shown than inferred from SDSS. However, there is somewhat better agreement for the highest masses, in particular, when using estimated masses in the second highest (for Lac14) and highest (for Gon14 and Lac14) mass bins.

In Fig. 3.8a, the clustering predictions using true stellar masses are similar for the two models, except that the Lac14 model predicts relatively low clustering amplitude in the lower mass bin. There is reasonable agreement with the large-scale GAMA clustering measurements of Farrow et al. (2015), especially in the higher mass bin. Both models predict stronger clustering on the smallest scales than is inferred observationally; this difference is most significant for the Gon14 model in the lower mass bin. The Lac14 model clustering predictions yield improved agreement with the GAMA measurements when switching to estimated masses. However, the corresponding increases in the clustering amplitudes of the Gon14 model lead to poorer agreement with the GAMA data than found with the true stellar masses (see Fig. 3.8a,b).

Using true stellar masses, the model predictions agree with the VIPERS clustering measurements of Marulli et al. (2013) on intermediate scales ($\sigma \sim 1\,h^{-1}\,\mathrm{Mpc}$) in the two lower mass bins in Fig. 3.9a, where they tend to overpredict the clustering amplitude on smaller and larger scales. In the

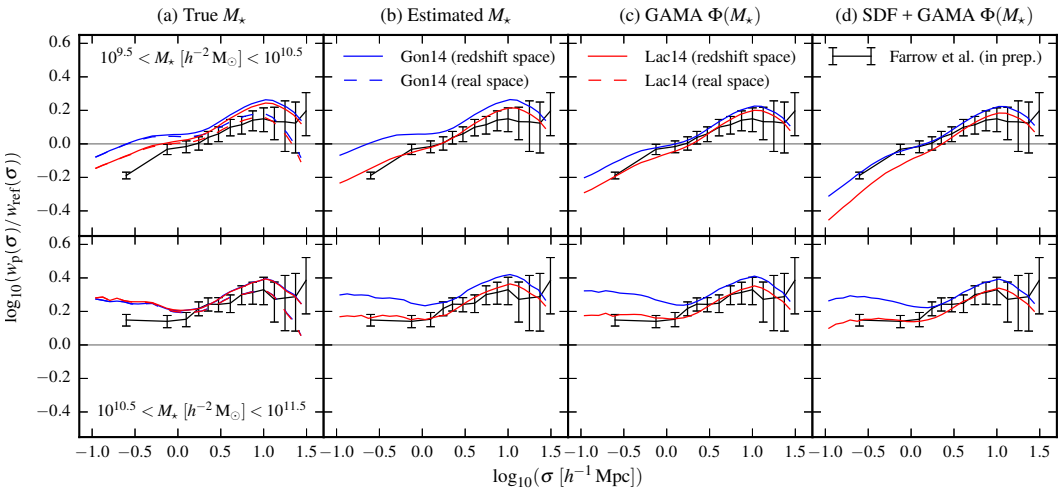

Figure 3.8 – Projected correlation function of galaxies, $w_\mathrm{p}(\sigma)$, for different stellar mass, $M_\star$, bins (as labelled in each row) at redshift $z = 0.17$, computed using the Gon14 and Lac14 GALFORM models, in real (first column only) and redshift space. The GAMA measurements of Farrow et al. (2015) are shown, whose apparent magnitude limit has been imposed on the model galaxies (black lines with errorbars; labelled *in prep.* since the final data published by Farrow et al. 2015 use slightly different mass intervals etc. to those used here; see §3.3.3). For clarity, each $w_\mathrm{p}(\sigma)$ has been divided by a reference power law, $w_\mathrm{ref}(\sigma)$, with parameters $r_0 = 5\,h^{-1}\,\mathrm{Mpc}$ and $\gamma = 2$ (horizontal line; see equation 3.9). Columns (a) and (b) show the default models with the true and estimated masses, respectively. The remaining columns show the results of matching the GAMA stellar mass function (as an adjustment following the SED fitting), for (c) the default models, and (d) the models using the new subhalo dynamical friction (SDF) merger scheme.

highest mass bin, both models agree reasonably well with the observational measurements on small and large scales, but predict lower clustering amplitude than observed on intermediate scales. Switching to the estimated masses (see Fig. 3.9b), improves the large-scale agreement with the VIPERS data for the Lac14 model, in the two lower mass bins. Otherwise, the level of agreement between the model predictions and observational measurements on large scales is similar or reduced when switching to the estimated masses. On small scales, switching to estimated masses worsens the agreement with the VIPERS clustering data for each model and mass bin, except for the Lac14 model in the intermediate bin, where the level of agreement is slightly increased. These changes result in the models with estimated masses yielding clustering predictions that encompass the VIPERS data on small scales, with the clustering amplitude being lower for the Lac14 model, and higher for the Gon14 model.

It is clear from the results shown in columns (a) and (b) of Figs. 3.7, 3.8, and 3.9 that carrying out SED fitting to the broad-band photometry of model galaxies to recover estimates of their stellar masses can have a significant impact on the model clustering predictions for samples selected by stellar mass. Thus it is important to implement this procedure to obtain a reliable comparison between the clustering of galaxies as a function of stellar mass, as measured for observed galaxies and predicted by theoretical models.

3.5.2.2 Abundance matching of stellar mass functions

Figs. 3.7c, 3.8c, and 3.9c show the projected correlation functions resulting from carrying out abundance matching of the estimated model stellar mass functions to Schechter (1976) function fits to the mass function data measured using each survey (SDSS, GAMA, and VIPERS respectively). The mass function fits are shown alongside the model predictions in Fig. 3.6 (see §3.4.1.2 for details of the matching procedure). The differences between the model clustering predictions shown in columns (b) and (c) of Figs. 3.7, 3.8, and 3.9 are thus entirely due to the model mass estimates from SED fitting being rescaled to reproduce the observationally inferred mass function fits.

Looking at the comparison to the SDSS results of Li et al. (2006) in Fig. 3.7b,c, the mass function matching leads to a suppression of the model clustering amplitude for both models (in most mass bins). This leads to

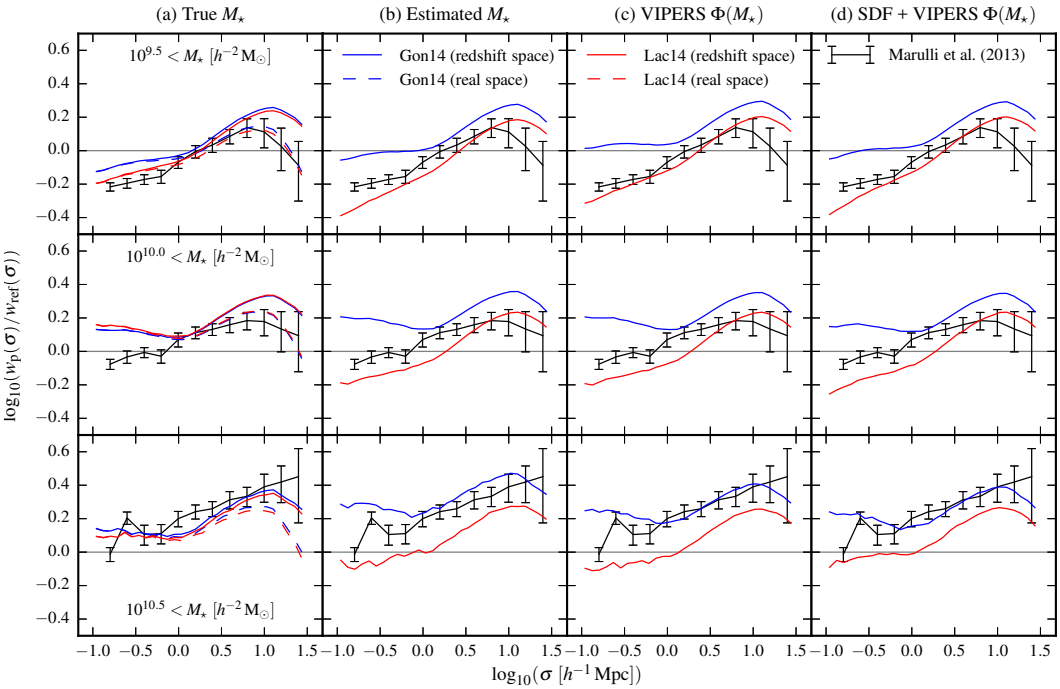

Figure 3.9 – Projected correlation function of galaxies, $w_\mathrm{p}(\sigma)$, as a function of cumulative stellar mass, $M_\star$ (as labelled in each row), at redshift $z = 0.62$, computed using the Gon14 and Lac14 GALFORM models, in real (first column only) and redshift space. The VIPERS measurements of Marulli et al. (2013) are shown (black lines with errorbars), whose apparent magnitude limit has been imposed on the model galaxies. For clarity, each $w_\mathrm{p}(\sigma)$ has been divided by a reference power law, $w_\mathrm{ref}(\sigma)$, with parameters $r_0 = 5\,h^{-1}$ Mpc and $\gamma = 2$ (horizontal line; see equation 3.9). Columns (a) and (b) show the default models with the true and estimated masses, respectively. The remaining columns show the results of matching the VIPERS stellar mass function (as an adjustment following the SED fitting), for (c) the default models, and (d) the models using the new subhalo dynamical friction (SDF) merger scheme.

improved, or unchanged, agreement with the observational results; except for the Lac14 model in both the second highest mass bin (where poorer agreement is found due to a rise in clustering amplitude), and in the highest mass bin (where the model $w_{\rm p}(\sigma)$ becomes lower in amplitude than measured from SDSS). However, the changes in the clustering predictions due to imposing the SDSS mass function are typically relatively minor, and in general there remains excessive clustering on both small and large scales with respect the SDSS measurements in Fig. 3.7c.

Comparing to the measurements of Farrow et al. (2015) using GAMA data (see Fig. 3.8b,c), imposing the observationally inferred mass function causes the clustering amplitude to be suppressed for both models in the lower mass bin (where the most significant impact is on the Gon14 model). This leads to improved agreement with the GAMA clustering results for the Gon14 model, and similar agreement for the Lac14 model, in this mass bin. In the higher mass bin, imposing the mass function has very little influence on the Lac14 model clustering prediction, but slightly boosts the amplitude of the small-scale clustering in the Gon14 model, exacerbating the discrepancy with the GAMA measurement on small scales.

It can be seen in Fig. 3.9b,c that the mass function matching leads to improved agreement with the VIPERS data in the lower mass interval for the Lac14 model, and worsened agreement for the Gon14 model, as the clustering amplitude rises in both models. The matching does not have a significant impact on the clustering predictions of either model in the intermediate mass interval. Both models exhibit suppression in the clustering amplitude due to the mass function matching for the highest masses, improving the agreement with the VIPERS data for the Gon14 model, and reducing the level of agreement for the Lac14 model. Thus there persist notable differences between the model predictions and the VIPERS measurements following the abundance matching, in Fig. 3.9c. The Gon14 model predicts excessive clustering with respect to the VIPERS results on both small and large scales in the two lower mass bins. The Lac14 model predicts lower clustering amplitudes than measured from VIPERS, for small scales in the intermediate mass interval, and on all scales for the highest masses.

The changes in the model clustering predictions due to matching observationally inferred stellar mass function fits reflect the underlying modi-

fications to the distribution of halo masses in a given stellar mass interval. Decreases in the clustering amplitude are seen where there is a significant net influx of galaxies in haloes of relatively low mass, thus diluting the clustering signal (and *vice versa*). The changes due to the stellar mass function matching are typically smaller than the changes that occur when switching from true model stellar masses to estimated masses recovered from SED fitting to the model photometry.

3.5.2.3 New satellite merger scheme

Thus far, we have compared the model clustering predictions to observational results using the default treatment of satellite galaxies in GALFORM. In the default scheme, satellite galaxies are assumed to enter the main halo on random orbits (independently of the orbit of their associated subhalo), and merge with the central galaxy upon the elapse of an analytically determined dynamical friction merger time-scale. We now consider the results obtained when using a new merger scheme, in which satellite galaxies track their associated subhaloes until these are no longer resolved, and only then is an analytic merger time-scale computed (see §3.2.4).

Using the standard merger scheme (looking in particular at Figs. 3.7c, 3.8c, and 3.9c), the Gon14 model tends to predict stronger small-scale clustering than seen in the observational data. The Lac14 model predicts higher small-scale clustering than measured from SDSS in Fig. 3.7c, but is in agreement with (or predicts lower clustering amplitude than) the small-scale results from GAMA and VIPERS (Figs. 3.8c and 3.9c respectively). It is noteworthy that both models tend to predict higher clustering amplitudes on small scales than measured from SDSS.

Stronger small-scale clustering in the models, with respect to observational results, may be interpreted as an excess of satellite galaxies in the models, or too concentrated a radial distribution of satellites (Contreras et al., 2013). In order to reconcile the model predictions with observations, it is possible that a more detailed study of the disruption and mergers of satellite galaxies is needed (Henriques et al., 2008; Kim et al., 2009). Another possibility is to make use of the substructure information available in dark matter simulations, i.e. to employ our new satellite merger scheme, which has an impact on small-scale clustering (see §3.2.4 and §3.5.1.2).

The estimated stellar mass functions of the Gon14 and Lac14 models when using the new 'subhalo dynamical friction' merger scheme are shown as dashed lines in Fig. 3.6 (where they are labelled Gon14-sdf and Lac14-sdf). Comparing these lines to the estimated mass functions using the standard merger scheme (dotted lines), it can be seen that the most significant differences are at the highest masses. At each redshift, both models have more high-mass galaxies when using the new scheme, where the differences between the predictions for the two schemes for a given model emerge roughly at the knee of the mass function. For the true masses, the mass functions for both models using the new scheme are very similar to each other at the highest masses, for each survey comparison redshift (to the level shown in Fig. 3.4 at the SDSS comparison redshift); and yet the estimated mass functions shown in Fig. 3.6 are distinct at high masses for the Gon14 and Lac14 models using the new scheme.

Figs. 3.7d, 3.8d, and 3.9d show the projected clustering predictions of the GALFORM models using the new merger scheme. These results use stellar masses estimated through SED fitting, followed by matching to the observational stellar mass function fits shown in Fig. 3.6 (as described in §3.4.1.2). In this way, the results shown in columns (c) and (d) of Figs. 3.7, 3.8, and 3.9 have been required to reproduce the same observationally inferred mass function fits, making use of the ordering of galaxies in stellar mass obtained through SED fitting, but differing in whether they use (c) the default merger scheme, or (d) the new scheme.

The results shown in Fig. 3.5 show that the changes in the model clustering predictions due to the new merger scheme, for samples selected by true stellar mass, are larger for lower stellar masses and lower redshifts. This trend is also seen for the abundance matched estimated masses when comparing columns (c) and (d) in Figs. 3.7, 3.8, and 3.9.

Reductions in the clustering amplitude lead to improved small-scale agreement with the SDSS data for both models in Fig. 3.7c,d. In the SDSS comparison, the Lac14 clustering also becomes stronger in the highest mass bin, resulting in improved agreement with the observational data on large scales. However, in several mass bins in Fig. 3.7d the small-scale clustering amplitude is still high with respect to the SDSS measurements, particularly for the Gon14 model. Additionally, the model clustering predictions on the

largest scales are often excessive with respect to the SDSS results in Fig. 3.7d, particularly so for the two lowest mass bins.

Both models experience suppressed small-scale clustering in Fig. 3.8c,d, which is more significant in the lower mass bin, leading to good agreement with the small-scale GAMA measurements for the Gon14 model in the lower mass bin, and for the Lac14 model in the higher mass bin. However, the Lac14 model clustering amplitude in Fig. 3.8d is now low on small scales with respect to the GAMA result in the lower mass bin. The suppression of the Gon14 clustering in the higher mass bin, due to the new merger scheme, is small compared to the excessive clustering on small scales with respect to the GAMA measurement in Fig. 3.8c.

In Fig. 3.9c,d, suppression of the small-scale clustering tends to push the $w_{\rm p}(\sigma)$ predicted by the Gon14 model closer to the VIPERS data, and further below this for the Lac14 model; except in the highest mass bin, where the Lac14 model clustering amplitude rises. However, the impact of the new merger scheme on the clustering predictions shown in Fig. 3.9d is relatively small, and the overall level of agreement with the VIPERS measurements is almost unchanged from the results shown in Fig. 3.9c.

3.5.2.4 The need for SED fitting

We have used abundance matching of model stellar mass functions to observational measurements as a basis for comparing the clustering predictions of the different GALFORM models to each other, and to observational estimates. The abundance matching has been carried out as a final adjustment to the estimated model galaxy masses. These results are shown in columns (c) and (d) of Figs. 3.7, 3.8, and 3.9, and have been discussed in §3.5.2.2 and §3.5.2.3 above. It is interesting to ask how these results, which have undergone mass function matching following SED fitting, compare to those obtained by simply matching the observationally inferred mass functions using the true model stellar masses. This test examines how important the SED fitting is when making clustering predictions to compare to observations, and in particular how significantly the ordering of galaxies in stellar mass changes when carrying out SED fitting (in terms of the impact on the clustering).

In a small number of cases, the changes due to the SED fitting are insignificant when using imposed mass functions. For example, the results for

the lowest mass GAMA galaxies (as shown in Fig. 3.8c,d) are unperturbed by SED fitting for both the Gon14 and Lac14 models, with or without the new merger scheme. Yet for the higher mass GAMA galaxies, not carrying out SED fitting results in clustering predictions using the new merger scheme that on small scales are midway in $\log_{10}(w_{\rm p}(\sigma))$ between those shown in Fig. 3.8d, for both GALFORM models. Similar trends are obtained in the highest VIPERS mass interval, for example, where for either merger scheme, both models yield small-scale clustering amplitudes midway between the model lines shown in Fig. 3.9c,d.

In general, we find that without SED fitting, the models with matched mass functions produce clustering predictions that are different to those of their counterparts based on estimated masses, particularly for $\sigma \lesssim 1\,h^{-1}\,$Mpc.

3.5.2.5 Sources of uncertainty

We now discuss sources of uncertainty that are relevant to the comparisons to observational data carried out in this work.

1. A potential complication when comparing to observational galaxy clustering measurements is the inability of spectrographs to resolve galaxies in close proximity on the sky (due to fibre collisions). This systematically lowers the clustering amplitude, particularly on small scales (e.g. Zehavi et al. 2002; Pollo et al. 2005; de la Torre et al. 2013).

   For SDSS, spectroscopic fibres on a given plate must be separated by at least 55 arcsec, which means that at the median redshift of Li et al. (2006), two galaxies cannot be observed simultaneously within $\sigma \lesssim 0.1\,h^{-1}\,$Mpc, although this can influence $w_{\rm p}(\sigma)$ out to larger scales. Li et al. carry out a correction to account for this, based on comparing the projected correlation function to that derived using the parent photometric catalogue. It is possible that, despite this correction, there remains some systematic bias in the observational results on the small scales shown in Fig. 3.7.

   GAMA employs a sophisticated 'greedy' tiling strategy, yielding nearly spatially uniform spectroscopic completeness (Robotham et al., 2010; Driver et al., 2011). Thus the small-scale clustering measurements of

Farrow et al. (2015) shown in Fig. 3.8 should be reasonably reliable (i.e. robust against fibre collisions).

A correction is made by Marulli et al. (2013) to account for spectroscopic fibre collisions in VIPERS, motivated by the Munich semi-analytic galaxy formation model of de Lucia & Blaizot (2007). Mock catalogues are constructed from this model, with and without modelling of the spectrograph selection. The relative difference between the correlation functions computed for the two sets of mock galaxies is used to impose a correction on the observed result. This model dependent correction may produce results that are systematically different on small scales to what would be measured for a truly spectroscopically complete sample. de la Torre et al. (2013) note that for VIPERS the fraction of missing galaxy pairs becomes significant below 0.03 degrees, which at the median redshift of the data to which we compare corresponds to $\sigma \lesssim 0.8\,h^{-1}$ Mpc (cf. Fig. 3.9).

2. The differences on large scales between $w_{\rm p}(\sigma)$ computed in real and redshift space for a given model in Figs. 3.7, 3.8, and 3.9 indicate that $\pi_{\rm max} = 30\,h^{-1}$ Mpc is not sufficiently large for the projected correlation function computed in redshift space to converge to the real space result, but this integration limit is appropriate for comparison to the observational data (see §3.4.2.3). Note that we integrate the real and redshift-space pair separations out to the same $\pi_{\rm max}$ when computing $w_{\rm p}(\sigma)$ (see §3.4.2.2). For clarity, the real-space clustering is shown only in column (a) of Figs. 3.7, 3.8, and 3.9. The offsets between the real and redshift-space clustering are similar for the columns where the real-space lines are not shown.

3. Differences exist between the broad-band SED fitting procedure implemented here, and the method of Kauffmann et al. (2003) used to estimate stellar masses in the SDSS data (see §3.3.2 and §3.4.1.1). However, our methodology yields stellar masses that are more appropriate to use than the true model stellar masses when comparing to these observational results.

4. Our SED fitting procedure uses the Calzetti et al. (2000) dust model,

whereas Davidzon et al. (2013) also permit the Prevot-Bouchet model for VIPERS (see §3.3.4). However, Davidzon et al. demonstrate that using just the Calzetti et al. model has only a marginal impact on the recovered stellar mass function.

5. As we make use of simulation snapshots in this chapter (i.e. model outputs at fixed redshifts), rather than constructing lightcones that cover specific redshift intervals (e.g. Merson et al. 2013), the interpretation of a sample of galaxies selected in apparent magnitude is slightly different to that for an observational survey. For example, if the model redshift is chosen to be close to the median redshift of an observational sample (as we do here), intrinsically fainter galaxies can exist in the survey sample than in the model, when the model and observations use the same faint apparent magnitude limit, due to the width of the observational redshift bin. This effect becomes important, for our purposes, for stellar masses at which the mass function becomes highly incomplete, due to the combination of redshift and apparent magnitude limits (see §3.4.1.2). In particular, this mass scale will tend to occur at higher stellar masses in a model snapshot, than in a survey sample with a lower minimum redshift than the snapshot redshift. While the use of lightcones could facilitate a more accurate treatment of the selections in apparent magnitude, the stellar mass ranges over which we compare the model clustering predictions to observational measurements are reasonably close to being complete in stellar mass. That is, we do not consider observational measurements where the survey selections result in a severely incomplete stellar mass function, and so the use of snapshots should not have a significant influence on the reliability of our comparisons. Lightcones must be used, for example, to compare theoretical model predictions to the higher redshift GAMA clustering data of Farrow et al. (2015), which we do not consider here for this reason.

3.6 Conclusions

The stellar masses of real galaxies have to be somehow derived from observables, typically using broad-band SED fitting. We have presented a new methodology for comparing the clustering predictions of galaxy formation models to observational data, for samples selected by stellar mass. The approach is to use estimated masses for the model galaxies, recovered using SED fitting to the model broad-band photometry. This allows us to incorporate the various systematic errors and biases involved in the fitting procedure, which can lead to significant differences with respect to the true stellar masses. These differences cannot be properly accounted for by assuming a mean rescaling, or even a rescaling plus a scatter, between the true and estimated masses, particularly for massive galaxies with strong dust attenuation (Mitchell et al., 2013).

If stellar masses from galaxy formation models and observations can be compared in such a consistent way, the clustering of galaxies as a function of stellar mass can be used to constrain the physical processes implemented in the models, alongside traditional statistics such as the luminosity function. Our methodology demonstrates how to do such a comparison.

We have compared the clustering predictions of the Gon14 and Lac14 GALFORM models to observational measurements from different surveys at redshifts of 0.1, 0.2, and 0.6. The clustering of the model galaxies as a function of stellar mass can change significantly when moving to the estimated masses. This can be understood in terms of the transfer of galaxies between stellar mass bins when estimated masses are used, which changes the distribution of halo masses in a given bin.

Considering the estimated model masses (and in particular following abundance matching to the observed stellar mass functions), we have often found that the models predict higher small-scale clustering amplitude than is inferred observationally (for projected separations $\sigma \lesssim 1\,h^{-1}\,\mathrm{Mpc}$). This tends to be the case for both GALFORM models at the lowest redshift in our comparison (SDSS; see Fig. 3.7c), and for the Gon14 model at higher redshifts (comparisons to GAMA and VIPERS; see Figs. 3.8c and 3.9c). In the higher redshift comparisons, the Lac14 model tends to predict similar or lower small-scale clustering amplitude than is measured using the survey data

(again see Figs. 3.8c and 3.9c).

The small-scale clustering is sensitive to the treatment of galaxy mergers. We have introduced a new scheme for the merging of satellites with their central galaxy, in which satellite galaxies track their associated subhalo in the dark matter simulation until this is no longer resolved. In the standard GALFORM approach, galaxies are assigned analytic merger time-scales as soon as they become satellites. The latter implementation yields higher small-scale clustering, as noted by Contreras et al. (2013), with respect to the Munich LGALAXIES semi-analytic models. The Munich models use a merger scheme that is similar to the new scheme that we have implemented in GALFORM (see §3.2.4 for a full description of the new scheme). Using this new scheme, together with estimated stellar masses abundance matched to observationally inferred stellar mass functions, generally offers improved agreement with the observational data on small scales (particularly in the SDSS and GAMA comparisons; see Figs. 3.7c,d and 3.8c,d), or a similar overall level of agreement where the spread in the models encompasses the observational results (see the VIPERS comparison in Fig. 3.9c,d).

Despite reductions due to the new merger scheme, the small-scale clustering often remains higher than inferred from SDSS data, particularly for the Gon14 model (see Fig. 3.7d). In the same comparison, the predicted clustering on the largest scales is in agreement with the SDSS measurements for only the highest masses considered. The models using the new merger scheme overall agree reasonably well with the GAMA clustering data (see Fig. 3.8d). However, in each mass bin, only one of the two models is in good agreement with the small-scale clustering. Comparing the same models to the VIPERS measurements in Fig. 3.9d, the model that is closest to the observational constraints varies with stellar mass, and there can be significant differences between the modelled and observationally inferred clustering, on both small and large scales.

Kim et al. (2009) have shown that the galaxy clustering as a function of luminosity predicted by GALFORM is stronger on small scales than is measured from observations at low redshift, and thus argue for the inclusion in the models of mergers between satellite galaxies and their tidal disruption. Such adjustments may also improve the agreement with observational measurements of the clustering as a function of stellar mass. We have not

included these extensions here, but their potential implementation in future models, along with the new merger scheme introduced here, will produce galaxy formation model predictions that take better account of the true dynamics of galaxies in dark matter haloes, and are as faithful as possible to the underlying dark matter distribution, allowing more robust constraints on the galaxy formation physics.

In summary, we have found that the model clustering predictions agree reasonably well with measurements from GAMA ($z = 0.2$). However, we have identified some relatively large discrepancies with respect to data from SDSS ($z = 0.1$) and VIPERS ($z = 0.6$), considering in particular the results shown in column (d) of Figs. 3.7, 3.8, and 3.9. The differences between the model predictions and the observations vary as a function of stellar mass, redshift, and projected separation, for each GALFORM model. The models considered in this chapter have not been calibrated to reproduce any observational clustering data (or stellar mass functions). Our new methodology will enable future models to be calibrated to the clustering of galaxies as a function of stellar mass, thus providing more accurate and reliable constraints on how galaxies populate dark matter haloes.

Chapter 4

Conclusions

In the standard ΛCDM model of cosmology, the energy density of the Universe is currently dominated by dark energy in the form of a cosmological constant (Λ), which acts to accelerate the rate of the expansion of space. The remaining energy density is dominated by cold dark matter particles (CDM), leaving only a few percent of the total cosmic energy density for 'ordinary' baryonic matter. Following a period of rapid inflation after the Big Bang, the curvature (geometry) of space is assumed to be perfectly flat (Euclidean). Inflation not only ensures spatial flatness, but also provides a mechanism to generate the initial perturbations in the cosmic matter density field needed to initiate the formation of cosmic structure through gravitational instability, via the vast amplification of quantum mechanical fluctuations.

Theoretical predictions derived assuming the ΛCDM model have been highly successful in reproducing the observed properties of the Universe on large scales. By measuring the amplitudes of minute anisotropic temperature fluctuations in the cosmic microwave background (CMB), we are able to infer the statistical properties of Gaussian perturbations in the cosmic matter density field at very high redshift, long before structure formation became non-linear. Simulations with initial conditions consistent with the CMB measurements can be used to study the evolution of cosmic structure from small density perturbations at early times into the intricate cosmic web of

the present Universe, joining together dark matter haloes covering a vast dynamic range in mass. Simulations including baryons are able to also follow the formation and evolution of luminous galaxies within the gravitational potential wells of the dark matter haloes.

In Chapter 2 of this thesis, we have used the APOSTLE cosmological hydrodynamical simulations to investigate the structural and stellar dynamical properties of galaxies found in systems like the Local Group. In particular, we have focused on how the accuracy of simple dynamical mass estimators of the form $M(< \lambda R_{\rm e}) \propto \langle \sigma_{\rm los} \rangle^2 R_{\rm e}$ depends on various important galaxy properties (where $R_{\rm e}$ is the projected stellar half-light radius, $\langle \sigma_{\rm los} \rangle$ is the stellar line-of-sight velocity dispersion, and $M(< \lambda R_{\rm e})$ is the total mass enclosed within a sphere of radius $\lambda R_{\rm e}$, for some constant λ). The dynamical masses and density profiles of the satellite galaxies of the Milky Way have played a particularly prominent role in various debates on the validity of the ΛCDM model on small scales, and so it is crucial that we have a clear understanding of the uncertainties involved in estimating these quantities. We have found that for the dispersion-dominated simulated galaxies, the simple mass estimators are almost unbiased, but with an overall 1σ scatter of between 23 and 25 percent for the analytically derived estimators in the literature, which rely on the spherical Jeans equation, combined with additional simplifying assumptions (Walker et al., 2009; Wolf et al., 2010). The estimator error is particularly sensitive to the 3D shape of the stellar mass distribution, such that galaxies that are closer to being spherically symmetric yield more accurate mass estimates; yet the 3D shape of a galaxy is difficult to constrain observationally. The uncertainties that we identify in the estimated masses are large enough to be significant for those Milky Way dwarf spheroidal (dSph) satellite galaxies that have relatively small observational errors on their stellar velocity dispersions and half-light radii. Using a simple model to emulate the method of Walker & Peñarrubia (2011) for measuring the density profile slopes of the Fornax and Sculptor dSphs, using simple mass estimators applied to chemo-dynamically distinct stellar subpopulations, we have shown that taking the uncertainties on the enclosed masses into account can relax the significance with which such methods can exclude NFW (Navarro, Frenk & White, 1996b, 1997) inner density profile slopes, depending on the degree to which the relevant properties of the different

stellar populations are correlated. We have also derived a new unbiased optimal mass estimator, where the stellar velocity dispersion is measured within a radius close to the projected half-light radius, in contrast to the typical definition where the dispersion is integrated over the whole galaxy. This new estimator reduces the overall scatter for our dispersion-dominated galaxies to 20 percent, facilitating more accurate and precise mass estimates. In the future, if chemo-dynamically distinct stellar populations can be identified within galaxies in the APOSTLE simulations, applying simple mass estimators to the separate populations would provide a direct test of the accuracy of the methodology of Walker & Peñarrubia (2011).

In Chapter 3, we have investigated the galaxy clustering predictions of two recently developed complementary versions of the GALFORM semi-analytic galaxy formation model, as a function of stellar mass and redshift. The stellar mass is a fundamental galaxy property, that reflects the intricate evolution history of the host dark matter halo and the baryonic gas, including the history of gas cooling, star formation, and energetic feedback from supernovae and active galactic nuclei. Measurements of the clustering of galaxies as a function of stellar mass thus provide an important constraint on how galaxies populate dark matter haloes. We have carried out broad-band spectral energy distribution (SED) fitting to estimate the stellar masses of the simulated galaxies, tailoring the estimation procedure to closely match the methods used in certain observational studies, to which we then compare the predicted galaxy clustering signals (projected correlation functions). The errors involved in SED fitting have a significant impact on the recovered stellar mass estimates. As a result the galaxy clustering as a function of stellar mass predicted by the models is different when using the true model stellar masses versus the masses estimated via SED fitting. Thus we argue that such an approach, where the model stellar masses are estimated using SED fitting, is vital when comparing model clustering predictions as a function of stellar mass to observationally inferred correlation functions. When such an approach is implemented, observational clustering measurements as a function of stellar mass can be used as a powerful constraint on the predictions of theoretical galaxy formation models, in addition to traditional statistics such as the luminosity function of galaxies. We have also introduced a new merger scheme for the merging of satellite galaxies with the central

galaxy of their halo in GALFORM, which implements a more realistic treatment of the satellite orbits than in the default merger scheme, making greater use of the underlying dark matter only N-body simulation. The new merger scheme influences the galaxy clustering, particularly on small scales (one-halo term), where the clustering is suppressed, such that similar or better agreement is typically found with the observational data. The inclusion of the new merger scheme in future galaxy formation models will allow more realistic predictions of galaxy properties to be generated, including for the clustering as a function of intrinsic properties, such as stellar mass.

The main theme in this thesis has been the accuracy of techniques used to infer the masses of galaxies, via detailed Jeans modelling or simple estimators to recover dynamical masses, or by fitting population synthesis models to spectra of galactic starlight to infer stellar masses. Significant uncertainties lurk in the practical application of such methods, where the error in the recovered mass is likely to be sensitive to various galaxy properties, convolved with the assumptions and modelling choices required to estimate the masses. The errors on the estimated masses may have serious consequences, such as perturbing the apparent halo masses of dispersion-supported galaxies, or implying dark matter cores in dSph galaxies with dual stellar populations where there may in fact be dark matter cusps, or modifying the perceived clustering of galaxies as a function of stellar mass. It is clearly important that the amplitudes and variations of the mass errors be well understood, particularly when the mass estimation techniques are used to produce results from which conclusions are drawn on the accuracy of certain theoretical models or even the validity of the ΛCDM paradigm.

Appendix A

Statistics

This appendix describes certain statistical functions and techniques that are used elsewhere in this thesis. For further details, see for example the books by Riley et al. (2006) and Hughes & Hase (2010).

A.1 Probability distributions

The *Gaussian* (or *normal*) *distribution* has the *probability density function*

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right], \tag{A.1}$$

where μ is the *mean* and σ is the *standard deviation* (σ^2 is the *variance*). The random variable, x, can take any real value, between $-\infty$ and ∞. The normalisation of $P(x)$ satisfies $\int_{-\infty}^{\infty} P(x)\,\mathrm{d}x = 1$. The cumulative probability for a Gaussian distribution, $\int_{-\infty}^{x} P(x')\,\mathrm{d}x'$, is given by the *error function*.

If some quantity is assumed to follow a Gaussian distribution, the '1σ error' is often quoted, which is just the standard deviation (σ) of the distribution. Given μ and σ for some quantity that is Gaussian distributed, we can expect that 68 percent ($\sim 2/3$) of randomly drawn values will fall in the range $\mu \pm \sigma$ (since $\mu - \sigma$ and $\mu + \sigma$ are approximately the 16^{th} and 84^{th} percentiles of

the Gaussian distribution, respectively). Similarly, the range $\mu \pm 2\sigma$ contains 95 percent of values, and the range $\mu \pm 3\sigma$ contains 99.7 percent of values.

It is important not to confuse the standard deviation of the parent probability distribution, σ, with the *standard error on the mean*, which is the *standard deviation of the distribution of sample means*, for many samples. For samples containing N measurements, the standard error on the mean is $\alpha = \sigma/\sqrt{N}$, where σ is still the standard deviation of the underlying probability distribution as above, which can be estimated as the standard deviation of a sample of data drawn from the distribution (this result happens to be true for any distribution, not just the Gaussian distribution).

The Gaussian distribution is symmetric around the mean (and the mean is also the median). However, for asymmetric distributions, we often examine the 16^{th} and 84^{th} percentiles, by way of analogy with the standard deviation of a Gaussian distribution. The standard deviation can always be computed, of course, but its usefulness is likely to diminish for highly asymmetric distributions that are very different in shape to the Gaussian distribution.

According to the *central limit theorem*, the sum of many arbitrary probability distributions, for independent random variables, with different means and variances, tends towards a Gaussian distribution (and so the distribution of sample means tends towards a Gaussian). Therefore the Gaussian distribution is uniquely important, since many physical quantities can be assumed to follow Gaussian statistics, even if the precise underlying details of the relevant physical processes are not fully specified or understood.

The *Poisson distribution* relates to the statistics of counting discrete independent random events. In this case, the probability density function for counting x events in a given time interval, t, is

$$P(x) = \frac{e^{-\lambda t}(\lambda t)^x}{x!}\,, \tag{A.2}$$

where λ is the mean event rate. The expectation (mean) value for the number of events within the time interval is $\bar{x} = \lambda t$. The standard deviation for the Poisson distribution is given by $\sigma = \sqrt{\bar{x}}$. For an experiment that counts N events, it is thus standard practice to quote the uncertainty (1σ error) on the count as $\sqrt{N}$ (known as the *Poisson error*). Although the Poisson distribution is asymmetric, as $\bar{x}$ increases the distribution actually grows

closer to a symmetric shape around the mean. In fact, for high $\bar{x}$, the Poisson distribution can be closely approximated by a Gaussian distribution with the same mean ($\mu = \bar{x}$) and standard deviation ($\sigma = \sqrt{\bar{x}}$).

In this thesis, all quoted uncertainties and plotted errorbars are 1σ errors (or 16^{th} and 84^{th} percentiles), unless stated otherwise.

A.2 Model fitting

Given some set of measurements containing N data points (x_i, y_i), the parameters of some appropriate model function $\tilde{y}(x)$ can be constrained by minimising the χ^2 statistic, defined such that

$$\chi^2 = \sum_{i=1}^{N} \left(\frac{y_i - \tilde{y}(x_i)}{\alpha_i} \right)^2 , \tag{A.3}$$

where α_i is the uncertainty on the y_i value for a given x_i. In the case where α_i is the same for all i, the χ^2 minimisation problem reduces to an *unweighted least-squares* fit (minimisation of squared residuals). If the y_i values are individually Gaussian distributed, with Gaussian width (σ_i proportional to) α_i, then the set of model parameters that minimises the χ^2 is the *maximum likelihood* solution (and the typical appropriateness of assuming Gaussianity follows from the central limit theorem).

For a model with n free parameters $(\psi_1, \psi_2, \ldots, \psi_n)$, the χ^2 minimisation takes place in n-dimensional parameter space. In order to quantify the uncertainty on the determined optimum value of a certain parameter, ψ_j, we can explore the sensitivity of the value of χ^2 to the value of ψ_j in the region of parameter space around the optimum χ^2 value (χ^2_{min}). By varying ψ_j in this region, while continuously *re-optimising* all of the other $n-1$ parameters, the values of ψ_j at which $\chi^2 = \chi^2_{\text{min}} + 1$ can be found; these values of ψ_j correspond to $\pm 1\sigma$ variations in ψ_j, while all $n-1$ other parameters are always optimised to minimise χ^2, given the chosen value of ψ_j. Following the same procedure, different values of $\Delta\chi^2 = \chi^2 - \chi^2_{\text{min}}$ correspond to different confidence intervals for ψ_j. For example, $\Delta\chi^2$ values of 1, 4, and 9 correspond to $\pm 1\sigma$, $\pm 2\sigma$, and $\pm 3\sigma$ variations in ψ_j, respectively. In order to compute e.g. the 1σ uncertainty for each of the n individual parameters

$(\psi_1, \psi_2, \ldots, \psi_n)$, the process just described needs to be repeated for each of the ψ_j in turn (i.e. for each choice of j), where in each case all of the other $n-1$ parameters must be re-optimised with each change in the value of ψ_j. Quoting only the set of optimum parameters along with their associated e.g. 1σ uncertainties can hide important correlations between the different parameters in n-dimensional space (i.e. the precise shape of the n-dimensional error surface could be important, introducing strong correlations between different model parameters). In order to investigate possible correlations between any two parameters, ψ_j and ψ_k, the value of χ^2 can be computed over the 2D $\psi_j - \psi_k$ plane, where at each point all of the remaining $n-2$ parameters are re-optimised to minimise χ^2 given the chosen values of ψ_j and ψ_k. Contours of constant χ^2 in the $\psi_j - \psi_k$ plane then indicate how the two parameters vary with respect to each other within the fitting process.

Markov Chain Monte Carlo (MCMC) methods provide an efficient way to explore large high-dimensional parameter spaces in order to constrain model parameters, using Bayesian statistics (e.g. Lewis & Bridle 2002; Walker & Peñarrubia 2011). The posterior probability distributions of the model parameters can be computed from the MCMC *chains*, i.e. histories of the repeated exploration of the parameter space by the MCMC algorithm. Optimum parameter values and confidence intervals can then be derived from the posterior distributions.

A.3 Bootstrap sampling

In many situations, the properties of the unknown underlying probability distribution of the measured data of interest are not obvious or well defined. Additionally, while the error on the mean (standard deviation of the distribution of sample means) can be estimated as $\sigma/\sqrt{N}$ for a sample of size N with standard deviation σ, the uncertainties in other statistics derived from the data are unlikely to have such an analytical form.

Bootstrap sampling is a powerful technique for estimating the confidence intervals of interest for *any* statistic computed for *any* dataset, regardless of the complexity of the statistic of interest or the underlying probability distribution of the measured data. Given a sample of measured data of size N, a large number of synthetic data samples are generated by drawing

random entries (numbers, vectors, particles, or whatever) from the true, measured, data. The synthetic samples each have size N, like the true sample, and are randomly chosen *with replacement*, so any given datum from the true sample can in theory appear any number of times (from zero to N) in a given synthetic sample. The statistic of interest is then computed for each randomly generated synthetic sample. The resulting distribution of the values of the statistic is treated as an estimate of the sampling distribution for that statistic. Properties of the sampling distribution can then be computed, e.g. the standard deviation, or certain percentiles. For example, given the standard deviation, σ_X, measured from the sampling distribution obtained via bootstrap sampling and repeatedly computing the statistic X, the value of X can be stated as $X = X_0 \pm \sigma_X$, where X_0 is the optimum value of X, as obtained from analysis of the original data sample.

Appendix B

Wolf et al. Estimator in 3D

In Chapter 2 we made use of the estimator of Wolf et al. (2010) expressed in terms of observable quantities: i.e. the projected stellar half-light (half-mass) radius, R_e, and line-of-sight velocity dispersion, $\langle\sigma_{los}\rangle$ (equation 2.9). However, a more fundamental version of the estimator is given in equation (2.8), from which Wolf et al. (2010) derive equation (2.9) by simply assuming that the 3D half-light radius, r_{half}, is equal to $4R_e/3$, which is a reasonable approximation for a range of spherically symmetric density profiles. Yet the appropriateness of this assumption could vary significantly for some galaxies, depending on the shapes of their stellar density profiles. We now investigate the change in the accuracy of the estimator of Wolf et al. (2010), with respect to the standard case given in equation (2.9), when: (i) a value for r_{half} is inferred from the projected stellar density distribution, instead of assuming $r_{half} = 4R_e/3$, and (ii) typically unobservable 3D information on the galaxy size and velocity dispersion is used in the mass estimation.

B.1 Deprojection of the stellar mass profile

In order to assess whether the assumption $r_{half} = 4R_e/3$ is optimal for the dispersion-dominated galaxies in our simulated sample, which have a broad

range of 3D shapes, we carry out an Abel inversion of the projected mass distribution for all 1536 projections of each galaxy. We make use of splines to model the cumulative projected stellar mass profile, $M_{\rm star}(< R)$. This choice allows the model profile to be quite general in form, in the sense that the inference of $r_{\rm half}$ should not suffer from any overly restrictive parameterisation that may bias the results. In detail, we find the quartic[1] spline that interpolates exactly through points of even spacing, ΔR, on the $M_{\rm star}(< R)$ profile, starting from the galactic centre. The profile is reflected about the centre before fitting, so that the model is well behaved in the innermost regions. Additionally, we extend $M_{\rm star}(< R)$ for several ΔR beyond the furthest star particle (where it equals the total stellar mass, $M_{\rm star}$). Starting with an initial spacing of $\Delta R = R_{\rm e}/2$, we recursively multiply ΔR by 0.95 and refit the spline until the following condition is satisfied:

$$\frac{1}{\sigma_{\rm fit}^2} \sum_i [M_{\rm model}(< R_i) - M_{\rm star}(< R_i)]^2 < n_{\rm star} \,, \tag{B.1}$$

where the summation is over all star particles in the galaxy, with projected radii R_i. $M_{\rm model}(< R)$ is the stellar mass enclosed within R according to the spline fit, and $n_{\rm star}$ is the number of star particles. The choice of the parameter $\sigma_{\rm fit} = 0.01 M_{\rm star}$ represents a compromise between the smoothness and accuracy of the fit.

Given a model for $M_{\rm star}(< R)$, and assuming spherical symmetry, the 3D stellar density profile, $\rho(r)$, follows from the Abel integral given in equation (2.5), where

$$\frac{{\rm d}\Sigma(R)}{{\rm d}R} = \frac{1}{2\pi R}\left[\frac{{\rm d}^2 M_{\rm star}(< R)}{{\rm d}R^2} - \frac{1}{R}\frac{{\rm d}M_{\rm star}(< R)}{{\rm d}R}\right]. \tag{B.2}$$

The equations can then be solved numerically for $r_{\rm half}$.

In Fig. B.1, we show the distribution of the ratio of the estimated 3D half-mass radius from our deprojection procedure, $r_{\rm deproj}$, to the true radius, $r_{\rm half}$, for all galaxies in our dispersion-dominated sample, alongside the corresponding distribution from assuming that $r_{\rm half}$ is simply $4R_{\rm e}/3$. The distributions are remarkably similar, and show that overall there is no improvement in the approximation of the 3D size when switching from using

[1] So that the second derivative is both continuous *and* smooth.

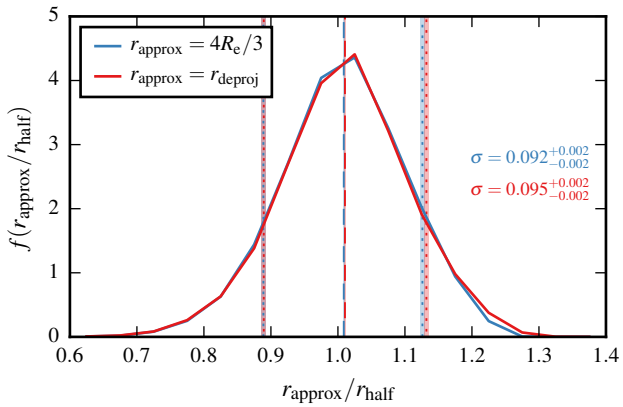

Figure B.1 – Distributions, $f(r_{\rm approx}/r_{\rm half})$, of the ratio of approximations, $r_{\rm approx}$, for the 3D half-mass radius, $r_{\rm half}$, to the true value of the radius for galaxies in our dispersion-dominated sample ($\kappa_{\rm rot} < 0.5$). The $r_{\rm approx}$ values are computed for 1536 evenly distributed lines of sight, assuming either that $r_{\rm half} = 4R_{\rm e}/3$, or using $r_{\rm deproj}$ from our Abel deprojection procedure (different colours). Each projection of each galaxy contributes to the distributions with equal weight, and each distribution is normalised to have unit area. The vertical dashed lines show the median ratios, and the dotted lines show the $10^{\rm th}$ and $90^{\rm th}$ percentiles. The standard deviation, σ, of each distribution is given in the same colour as the lines. The shaded regions around the vertical lines and the σ errors are the $16^{\rm th} - 84^{\rm th}$ percentile confidence limits, derived from 10^4 bootstrap samples of the galaxies for each distribution.

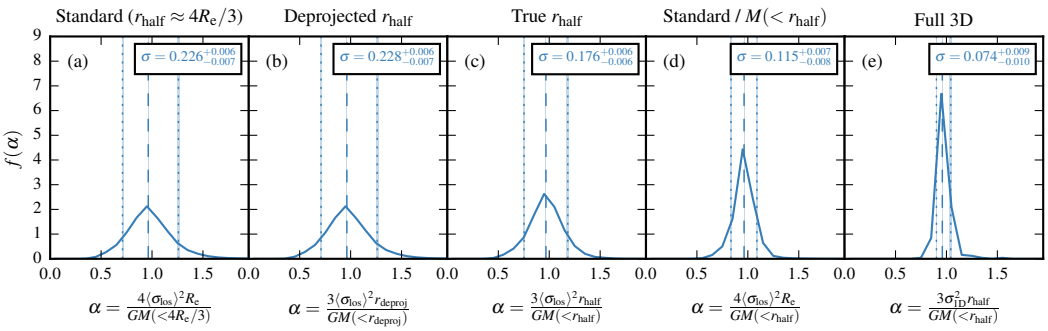

Figure B.2 – Distributions, $f(\alpha)$, of the estimated to true mass ratio, α, for all dispersion-dominated galaxies ($\kappa_{\rm rot} < 0.5$), obtained from projecting over 1536 evenly distributed lines of sight, using different versions of the estimator proposed by Wolf et al. (2010). The definition of α is different in each panel, and is given explicitly in each horizontal axis label. The definitions are: (a) the ratio of the estimated mass using equation (2.9) to the mass enclosed within the *assumed* 3D half-mass radius ($4R_{\rm e}/3$) – this is the standard definition of $\alpha_{\rm Wolf}$ as given in equation (2.14), (b) the estimated to true mass ratio from equation (2.8) using the deprojected 3D half-mass radius, $r_{\rm deproj}$, from our Abel integration procedure in place of $r_{\rm half}$, (c) the same as (b) but using the true 3D half-mass radius, $r_{\rm half}$, (d) the ratio of the estimated mass from equation (2.9) to the mass within $r_{\rm half}$ (i.e. this is the ratio of the estimated to true mass within *different spheres*, unless $r_{\rm half} = 4R_{\rm e}/3$), and (e) the estimated to true mass ratio if we replace $\langle\sigma_{\rm los}\rangle$ in equation (2.8) with the mean one-dimensional stellar velocity dispersion, $\sigma_{\rm 1D} = \sigma_{\rm 3D}/\sqrt{3}$ (so no projected information is used in this panel). Each projection (if applicable) of each galaxy contributes to the distributions with equal weight, and each distribution is normalised to have unit area. The vertical dashed lines show the median α values, and the dotted lines show the $10^{\rm th}$ and $90^{\rm th}$ percentiles. The standard deviation, σ, of each distribution is given in each panel. The shaded regions around the vertical lines and the σ errors are the $16^{\rm th}$ – $84^{\rm th}$ percentile confidence limits, derived from 10^4 bootstrap samples of the galaxies for each distribution.

the factor of four thirds to using our deprojection procedure (the scatter actually increases slightly for $r_{\rm deproj}$).

In Fig. B.2(a) we show the distribution of the estimated to true mass ratio, α, for all dispersion-dominated galaxies in our simulated sample, obtained using the standard projected version of the Wolf et al. (2010) estimator from equation (2.9) (as in the upper panel of Fig. 2.9). Fig. B.2(b) shows the distribution if instead we use the 3D version of the estimator as given in equation (2.8), making the replacement $r_{\rm half} \rightarrow r_{\rm deproj}$. There is no significant difference between the distributions shown in panels (a) and (b), and so we conclude that assuming $r_{\rm half} = 4R_{\rm e}/3$ is already optimal, despite the scatter shown in Fig. B.1.

B.2 Using true 3D galaxy properties

It is interesting to ask how the accuracy of the estimator of Wolf et al. (2010) changes if intrinsic 3D, i.e. typically observationally inaccessible, information is available on the galaxy sizes and kinematics. Fig. B.2(c) shows the distribution of the estimated to true mass ratio when using the true 3D half-mass radius in equation (2.8), so the only projected quantity used to compute the distribution shown is $\langle\sigma_{\rm los}\rangle$. The removal of the noise due to the variation of the projected size reduces the scatter in the distributions, comparing to panels (a) and (b).

In Fig. B.2(d) we show the distribution of the ratio of the estimated mass from equation (2.9), i.e. the version of estimator that uses $R_{\rm e}$, to the true mass within $r_{\rm half}$. Thus, here α is *not* the ratio of the estimated mass to the true mass within the *same* sphere, as everywhere else in Chapter 2 and this appendix. The scatter in this distribution is lower than in panel (c), where the only difference in the definition of α is the replacement $3r_{\rm half} \rightarrow 4R_{\rm e}$ from (c) to (d). That is, equation (2.9) is a better estimator of the mass within the true 3D half-mass radius than equation (2.8). This result may seem counterintuitive at first, as the former makes use of only projected information to estimate the mass, while the true 3D half-mass radius is used in the latter case. This effect can be understood in terms of the coupled variation of $R_{\rm e}$ and $\langle\sigma_{\rm los}\rangle$ over lines of sight which view the galaxy from many orientations, such that the variation in the projected size counteracts

that in the (squared) dispersion, so that the product $4\langle\sigma_{\rm los}\rangle^2 R_{\rm e}$ has a smaller scatter than $3\langle\sigma_{\rm los}\rangle^2 r_{\rm half}$ for the vast majority of galaxies (see Fig. 2.13 and also Laporte et al. 2013a). The result that equation (2.9) is a better estimator of $M(< r_{\rm half})$ than equation (2.8) is only attractive if for some reason one wanted to know $M(< r_{\rm half})$ without any associated estimate of $r_{\rm half}$ itself! When full 3D information is available on the stellar velocities, so that we can replace $\langle\sigma_{\rm los}\rangle$ in equation (2.8) with the mean one-dimensional dispersion, $\sigma_{\rm 1D} = \sigma_{\rm 3D}/\sqrt{3}$, the scatter in the estimated masses reduces significantly, as shown in Fig. B.2(e). Note that no projected information is used in this panel, so each galaxy contributes a single value to the distribution.

Appendix C

GALFORM Model Parameters

As outlined in §3.2.3, the Gon14 and Lac14 GALFORM models are each based on the model of Lagos et al. (2012). This appendix gives details of the parameter differences between the Gon14 and Lac14 models; beyond the choice of IMF, SPS model, and default satellite merger scheme (these are described in §3.2.3 and §3.2.4). Table C.1 gives the values of the relevant parameters in each model. We refer the reader to Cole et al. (2000), Bower et al. (2006), and Lagos et al. (2012) for further details of the model physics.

The total star formation rate of a galaxy, ψ, is the sum of the contributions from quiescent star formation in the disc, $\psi_{\rm disc}$, and from bursts, $\psi_{\rm burst}$. The quiescent component is computed as

$$\psi_{\rm disc} = \nu_{\rm SF} f_{\rm mol} M_{\rm cold,disc} \,, \tag{C.1}$$

where $M_{\rm cold,disc}$ is the mass of cold gas in the disc, a fraction, $f_{\rm mol}$, of which is molecular. $\nu_{\rm SF}$ is the inverse of the star formation time-scale for molecular gas (see Table C.1).

The rate, $\dot{M}_{\rm reheat}$, at which cold gas is reheated by supernova feedback is

$$\dot{M}_{\rm reheat} = \psi \left(\frac{V_{\rm circ}}{V_{\rm hot}} \right)^{-\alpha_{\rm hot}} \,, \tag{C.2}$$

Table C.1 – Parameters that differ between the Gon14 and Lac14 models. A description of each parameter is given in the main text.

Parameter	Gon14	Lac14
$\nu_{\rm SF}$ [Gyr^{-1}]	0.5	0.74
$V_{\rm hot}$ [km s^{-1}]	425	320
$\alpha_{\rm reheat}$	1.2603	0.64
$\alpha_{\rm cool}$	0.6	0.8
$\epsilon_{\rm SMBH}$	0.0398	0.01
$f_{\rm orbit}$	1	0
$f_{\rm burst}$	0.1	0.05
$f_{\rm gasburst}$	0.1	0
$f_{\rm dyn}$	10	20
$\tau_{\rm min}$ [Gyr]	0.05	0.1
$\epsilon_{\rm disc}$	0.8	0.9

where $V_{\rm circ}$ is the circular velocity of the galaxy, at the half-mass radius. $V_{\rm hot}$ is a parameter with velocity units (see Table C.1), and $\alpha_{\rm hot} = 3.2$ for both models. Gas heated by supernova feedback is assumed to be ejected beyond the virial radius, and then reincorporated into the hot halo gas reservoir at a rate proportional to $\alpha_{\rm reheat}$ (see Table C.1).

The onset of AGN feedback in GALFORM requires the cooling time, $t_{\rm cool}$, and free-fall time, $t_{\rm ff}$, of the gas to satisfy

$$t_{\rm cool} > t_{\rm ff}/\alpha_{\rm cool} \,, \tag{C.3}$$

that is, the halo is assumed to be in a state of quasi-hydrostatic cooling. This criterion is controlled by the dimensionless parameter $\alpha_{\rm cool}$ (see Table C.1). Additionally, the cooling rate of the gas must be less than a fraction, $\epsilon_{\rm SMBH}$, of the Eddington luminosity (see Table C.1). If both conditions are satisfied, AGN feedback is assumed to suppress the cooling of the halo gas.

The orbital energy of merging galaxies is assumed to be proportional to $f_{\rm orbit}$, and is used to compute the size of the spheroid formed in the merger. $f_{\rm burst}$ is the minimum galaxy mass (stars plus cold gas) ratio in a minor merger required to trigger bursts of star formation (major mergers occur for mass ratios above 0.3). Bursts are suppressed in minor mergers if the gas fraction in the disc of the main galaxy is less than $f_{\rm gasburst}$. The star formation time-scale for a burst is given either by $f_{\rm dyn}$ times the bulge dynamical time,

or by the parameter $\tau_{\rm min}$, whichever is larger (see Table C.1).

Disc instabilities, causing bursts of star formation, are triggered when

$$\frac{V_{\rm circ}}{\sqrt{GM_{\rm disc}/r_{\rm disc}}} < \epsilon_{\rm disc} \,, \tag{C.4}$$

where $M_{\rm disc}$ and $r_{\rm disc}$ are the mass of the disc, and its half-mass radius, respectively, and $\epsilon_{\rm disc}$ is a dimensionless parameter (see Table C.1).

Bibliography

Abazajian K., et al., 2004, AJ, 128, 502

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

Adams J. J., Gebhardt K., Blanc G. A., Fabricius M. H., Hill G. J., Murphy J. D., van den Bosch R. C. E., van de Ven G., 2012, ApJ, 745, 92

Albrecht A., Steinhardt P. J., 1982, Physical Review Letters, 48, 1220

Amorisco N. C., Evans N. W., 2011, MNRAS, 411, 2118

Asztalos S. J., et al., 2010, Physical Review Letters, 104, 041301

Babcock H. W., 1939, Lick Observatory Bulletin, 19, 41

Baldry I. K., Glazebrook K., Driver S. P., 2008, MNRAS, 388, 945

Baldry I. K., et al., 2010, MNRAS, 404, 86

Baldry I. K., et al., 2012, MNRAS, 421, 621

Balogh M. L., Morris S. L., Yee H. K. C., Carlberg R. G., Ellingson E., 1999, ApJ, 527, 54

Barber C., Starkenburg E., Navarro J. F., McConnachie A. W., 2015, MNRAS, 447, 1112

Barnes J., Hut P., 1986, Nature, 324, 446

Battaglia G., Helmi A., Tolstoy E., Irwin M., Hill V., Jablonka P., 2008, ApJ, 681, L13

Baugh C. M., 2006, Reports on Progress in Physics, 69, 3101

Baugh C. M., Lacey C. G., Frenk C. S., Granato G. L., Silva L., Bressan A., Benson A. J., Cole S., 2005, MNRAS, 356, 1191

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

Bell E. F., de Jong R. S., 2001, ApJ, 550, 212

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

Benítez-Llambay A., Navarro J. F., Abadi M. G., Gottlöber S., Yepes G., Hoffman Y., Steinmetz M., 2016, MNRAS, 456, 1185

Benson A. J., 2010, Phys. Rep., 495, 33

Benson A. J., Bower R., 2010, MNRAS, 405, 1573

Benson A. J., Cole S., Frenk C. S., Baugh C. M., Lacey C. G., 2000a, MNRAS, 311, 793

Benson A. J., Baugh C. M., Cole S., Frenk C. S., Lacey C. G., 2000b, MNRAS, 316, 107

Berger M. J., Colella P., 1989, Journal of Computational Physics, 82, 64

Berlind A. A., Weinberg D. H., 2002, ApJ, 575, 587

Bertone S., De Lucia G., Thomas P. A., 2007, MNRAS, 379, 1143

Bett P., 2012, MNRAS, 420, 3303

Bielby R. M., Gonzalez-Perez V., McCracken H. J., Ilbert O., Daddi E., Le Fèvre O., Hudelot P., Kneib J.-P., Mellier Y., Willott C., 2014, A&A, 568, A24

Binney J., Mamon G. A., 1982, MNRAS, 200, 361

Binney J., Tremaine S., 2008, Galactic Dynamics, Second Edition. Princeton University Press

Binney J. J., Davies R. L., Illingworth G. D., 1990, ApJ, 361, 78

Binney J., Gerhard O., Silk J., 2001, MNRAS, 321, 471

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

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

Booth C. M., Schaye J., 2009, MNRAS, 398, 53

Bouchet P., Lequeux J., Maurice E., Prevot L., Prevot-Burnichon M. L., 1985, A&A, 149, 330

Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, MNRAS, 370, 645

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

Brooks A. M., Kuhlen M., Zolotov A., Hooper D., 2013, ApJ, 765, 22

Bruzual G., Charlot S., 1993, ApJ, 405, 538

Bruzual G., Charlot S., 2003, MNRAS, 344, 1000

Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, ApJ, 533, 682

Campbell D. J. R., Baugh C. M., Mitchell P. D., Helly J. C., Gonzalez-Perez V., Lacey C. G., Lagos C. d. P., Simha V., Farrow D. J., 2015, MNRAS, 452, 852

Campbell D. J. R., Frenk C. S., Jenkins A., Eke V. R., Navarro J. F., Sawala T., Schaller M., Fattahi A., Oman K. A., Theuns T., 2017, MNRAS, 469, 2335

Cappellari M., 2008, MNRAS, 390, 71

Cappellari M., et al., 2006, MNRAS, 366, 1126

Chabrier G., 2003, ApJ, 586, L133

Charlot S., Fall S. M., 2000, ApJ, 539, 718

Chisari N. E., Zaldarriaga M., 2011, Phys. Rev. D, 83, 123505

Cole S., Kaiser N., 1989, MNRAS, 237, 1127

Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, MNRAS, 319, 168

Coleman M. G., Da Costa G. S., Bland-Hawthorn J., Freeman K. C., 2005, AJ, 129, 1443

Coles P., Lucchin F., 2002, Cosmology: The Origin and Evolution of Cosmic Structure, Second Edition. Wiley

Conroy C., Gunn J. E., White M., 2009, ApJ, 699, 486

Contreras S., Baugh C. M., Norberg P., Padilla N., 2013, MNRAS, 432, 2717

Cooray A., Sheth R., 2002, Phys. Rep., 372, 1

Crain R. A., et al., 2015, MNRAS, 450, 1937

Cretton N., van den Bosch F. C., 1999, ApJ, 514, 704

Croton D. J., Springel V., White S. D. M., De Lucia G., Frenk C. S., Gao L., Jenkins A., Kauffmann G., Navarro J. F., Yoshida N., 2006, MNRAS, 365, 11

Dalla Vecchia C., Schaye J., 2012, MNRAS, 426, 140

Davidzon I., et al., 2013, A&A, 558, A23

Davis M., Peebles P. J. E., 1983, ApJ, 267, 465

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

Dicke R. H., Peebles P. J. E., Roll P. G., Wilkinson D. T., 1965, ApJ, 142, 414

Diemand J., Kuhlen M., Madau P., Zemp M., Moore B., Potter D., Stadel J., 2008, Nature, 454, 735

Dolag K., Borgani S., Murante G., Springel V., 2009, MNRAS, 399, 497

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

Dutton A. A., Courteau S., de Jong R., Carignan C., 2005, ApJ, 619, 218

Efstathiou G., Davis M., White S. D. M., Frenk C. S., 1985, ApJS, 57, 241

Einasto J., Saar E., Kaasik A., Chernin A. D., 1974, Nature, 252, 111

Ellis J., Hagelin J. S., Nanopoulos D. V., Olive K., Srednicki M., 1984, Nuclear Physics B, 238, 453

Evans N. W., An J., Walker M. G., 2009, MNRAS, 393, L50

Farrow D. J., et al., 2015, MNRAS, 454, 2120

Fattahi A., et al., 2016, MNRAS, 457, 844

Ferrara A., Bianchi S., Cimatti A., Giovanardi C., 1999, ApJS, 123, 437

Font A. S., Bower R. G., McCarthy I. G., Benson A. J., Frenk C. S., Helly J. C., Lacey C. G., Baugh C. M., Cole S., 2008, MNRAS, 389, 1619

Frenk C. S., White S. D. M., 2012, Annalen der Physik, 524, 507

Frenk C. S., et al., 1999, ApJ, 525, 554

Fryxell B., Olson K., Ricker P., Timmes F. X., Zingale M., Lamb D. Q., MacNeice P., Rosner R., Truran J. W., Tufo H., 2000, ApJS, 131, 273

Garilli B., et al., 2014, A&A, 562, A23

Gebhardt K., et al., 2003, ApJ, 583, 92

Gil de Paz A., et al., 2007, ApJS, 173, 185

Gingold R. A., Monaghan J. J., 1977, MNRAS, 181, 375

Gonzalez-Perez V., Baugh C. M., Lacey C. G., Kim J.-W., 2011, MNRAS, 417, 517

Gonzalez-Perez V., Lacey C. G., Baugh C. M., Frenk C. S., Wilkins S. M., 2013, MNRAS, 429, 1609

Gonzalez-Perez V., Lacey C. G., Baugh C. M., Lagos C. d. P., Helly J., Campbell D. J. R., Mitchell P. D., 2014, MNRAS, 439, 264

Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759

Governato F., et al., 2010, Nature, 463, 203

Grand R. J. J., Springel V., Gómez F. A., Marinacci F., Pakmor R., Campbell D. J. R., Jenkins A., 2016, MNRAS, 459, 199

Grand R. J. J., Gómez F. A., Marinacci F., Pakmor R., Springel V., Campbell D. J. R., Frenk C. S., Jenkins A., White S. D. M., 2017, MNRAS,

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

Guo Q., White S., Boylan-Kolchin M., De Lucia G., Kauffmann G., Lemson G., Li C., Springel V., Weinmann S., 2011, MNRAS, 413, 101

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

Guth A. H., 1981, Phys. Rev. D, 23, 347

Guth A. H., 2000, Phys. Rep., 333, 555

Guth A. H., Pi S.-Y., 1982, Physical Review Letters, 49, 1110

Guzzo L., et al., 2014, A&A, 566, A108

Hatton S., Devriendt J. E. G., Ninin S., Bouchet F. R., Guiderdoni B., Vibert D., 2003, MNRAS, 343, 75

Hawking S. W., 1982, Physics Letters B, 115, 295

Henriques B. M., Bertone S., Thomas P. A., 2008, MNRAS, 383, 1649

Henriques B., Maraston C., Monaco P., Fontanot F., Menci N., De Lucia G., Tonini C., 2011, MNRAS, 415, 3571

Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R. E., Guo Q., Lemson G., Springel V., 2013, MNRAS, 431, 3373

Hockney R. W., Eastwood J. W., 1988, Computer simulation using particles. CRC Press

Hopkins P. F., 2013, MNRAS, 428, 2840

Hopkins P. F., 2015, MNRAS, 450, 53

Hopkins P. F., Raives M. J., 2016, MNRAS, 455, 51

Hughes I. G., Hase T. P. A., 2010, Measurements and their Uncertainties.

Oxford University Press
Hunter C., 1977, AJ, 82, 271
Iovino A., Shaver P. A., 1988, ApJ, 330, L13
Irwin M., Hatzidimitriou D., 1995, MNRAS, 277, 1354
Jeans J. H., 1922, MNRAS, 82, 122
Jenkins A., 2010, MNRAS, 403, 1859
Jiang C. Y., Jing Y. P., Faltenbacher A., Lin W. P., Li C., 2008, ApJ, 675, 1095
Jiang C. Y., Jing Y. P., Lin W. P., 2010, A&A, 510, A60
Kahn F. D., Woltjer L., 1959, ApJ, 130, 705
Kaiser N., 1987, MNRAS, 227, 1
Kauffmann G., Colberg J. M., Diaferio A., White S. D. M., 1999, MNRAS, 307, 529
Kauffmann G., et al., 2003, MNRAS, 341, 33
Kennicutt Jr. R. C., 1983, ApJ, 272, 54
Kim H.-S., Baugh C. M., Cole S., Frenk C. S., Benson A. J., 2009, MNRAS, 400, 1527
Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, ApJ, 522, 82
Koch A., Kleyna J. T., Wilkinson M. I., Grebel E. K., Gilmore G. F., Evans N. W., Wyse R. F. G., Harbeck D. R., 2007, AJ, 134, 566
Komatsu E., et al., 2011, ApJS, 192, 18
Kowalczyk K., Łokas E. L., Kazantzidis S., Mayer L., 2013, MNRAS, 431, 2796
Kravtsov A. V., Klypin A., Hoffman Y., 2002, ApJ, 571, 563
Kravtsov A. V., Berlind A. A., Wechsler R. H., Klypin A. A., Gottlöber S., Allgood B., Primack J. R., 2004, ApJ, 609, 35
Kriek M., et al., 2010, ApJ, 722, L64
Kroupa P., 2001, MNRAS, 322, 231
Lacey C., Cole S., 1993, MNRAS, 262, 627
Lacey C. G., Baugh C. M., Frenk C. S., Benson A. J., 2011, MNRAS, 412, 1828
Lacey C. G., Baugh C. M., Frenk C. S., Benson A. J., Bower R. G., Cole S., Gonzalez-Perez V., Helly J. C., Lagos C. D. P., Mitchell P. D., 2016, MNRAS, 462, 3854
Lagos C. d. P., Lacey C. G., Baugh C. M., Bower R. G., Benson A. J., 2011, MNRAS, 416, 1566
Lagos C. d. P., Bayet E., Baugh C. M., Lacey C. G., Bell T. A., Fanidakis N., Geach J. E., 2012, MNRAS, 426, 2142
Laporte C. F. P., Walker M. G., Peñarrubia J., 2013a, MNRAS, 433, L54
Laporte C. F. P., White S. D. M., Naab T., Gao L., 2013b, MNRAS, 435, 901
Lawrence A., et al., 2007, MNRAS, 379, 1599

Lewis A., Bridle S., 2002, Phys. Rev. D, 66, 103511
Lewis A., Challinor A., 2011, CAMB: Code for Anisotropies in the Microwave Background, Astrophysics Source Code Library (ascl:1102.026)
Li C., White S. D. M., 2009, MNRAS, 398, 2177
Li C., Kauffmann G., Jing Y. P., White S. D. M., Börner G., Cheng F. Z., 2006, MNRAS, 368, 21
Li C., et al., 2012, MNRAS, 419, 1557
Liddle A., 2003, An Introduction to Modern Cosmology, Second Edition. Wiley
Lidsey J. E., Liddle A. R., Kolb E. W., Copeland E. J., Barreiro T., Abney M., 1997, Reviews of Modern Physics, 69, 373
Lilly S. J., et al., 2007, ApJS, 172, 70
Linde A. D., 1982a, Physics Letters B, 108, 389
Linde A. D., 1982b, Physics Letters B, 116, 335
Linde A. D., 1983, Physics Letters B, 129, 177
Linde A., Linde D., Mezhlumian A., 1994, Phys. Rev. D, 49, 1783
Liske J., et al., 2015, MNRAS, 452, 2087
Lonsdale C. J., et al., 2003, PASP, 115, 897
Lucy L. B., 1977, AJ, 82, 1013
Lyskova N., Thomas J., Churazov E., Tremaine S., Naab T., 2015, MNRAS, 450, 3442
Lyth D. H. D. H., Riotto A. A., 1999, Phys. Rep., 314, 1
MacArthur L. A., McDonald M., Courteau S., Jesús González J., 2010, ApJ, 718, 768
Magorrian J., 2006, MNRAS, 373, 425
Maraston C., 2005, MNRAS, 362, 799
Marchesini D., van Dokkum P. G., Förster Schreiber N. M., Franx M., Labbé I., Wuyts S., 2009, ApJ, 701, 1765
Marigo P., Girardi L., 2007, A&A, 469, 239
Martinez V. J., Portilla M., Jones B. J. T., Paredes S., 1993, A&A, 280, 5
Marulli F., et al., 2013, A&A, 557, A17
Mashchenko S., Wadsley J., Couchman H. M. P., 2008, Science, 319, 174
McConnachie A. W., 2012, AJ, 144, 4
Meneux B., et al., 2009, A&A, 505, 463
Merson A. I., et al., 2013, MNRAS, 429, 556
Mitchell P. D., Lacey C. G., Baugh C. M., Cole S., 2013, MNRAS, 435, 87
Mo H., van den Bosch F. C., White S., 2010, Galaxy Formation and Evolution. Cambridge University Press

Monaco P., Fontanot F., Taffoni G., 2007, MNRAS, 375, 1189

Monaghan J. J., 1992, ARA&A, 30, 543

Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJ, 524, L19

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

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

Muñoz R. R., et al., 2006, ApJ, 649, 201

Navarro J. F., Eke V. R., Frenk C. S., 1996a, MNRAS, 283, L72

Navarro J. F., Frenk C. S., White S. D. M., 1996b, ApJ, 462, 563

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493

Neyman J., Scott E. L., 1952, ApJ, 116, 144

Norberg P., et al., 2001, MNRAS, 328, 64

Norberg P., Baugh C. M., Gaztañaga E., Croton D. J., 2009, MNRAS, 396, 19

O'Shea B. W., Bryan G., Bordner J., Norman M. L., Abel T., Harkness R., Kritsuk A., 2004, ArXiv Astrophysics e-prints,

Odenkirchen M., et al., 2001, AJ, 122, 2538

Oh S.-H., de Blok W. J. G., Walter F., Brinks E., Kennicutt Jr. R. C., 2008, AJ, 136, 2761

Oh S.-H., de Blok W. J. G., Brinks E., Walter F., Kennicutt Jr. R. C., 2011, AJ, 141, 193

Oke J. B., 1974, ApJS, 27, 21

Orsi A., Lacey C. G., Baugh C. M., Infante L., 2008, MNRAS, 391, 1589

Ostriker J. P., Peebles P. J. E., 1973, ApJ, 186, 467

Ostriker J. P., Peebles P. J. E., Yahil A., 1974, ApJ, 193, L1

Pakmor R., Bauer A., Springel V., 2011, MNRAS, 418, 1392

Peñarrubia J., McConnachie A. W., Navarro J. F., 2008a, ApJ, 672, 904

Peñarrubia J., McConnachie A. W., Navarro J. F., 2008b, ApJ, 687, 1460

Peacock J. A., Smith R. E., 2000, MNRAS, 318, 1144

Peebles P. J. E., 1980, Phys. Scr, 21, 720

Penzias A. A., Wilson R. W., 1965, ApJ, 142, 419

Perlmutter S., et al., 1999, ApJ, 517, 565

Pforr J., Maraston C., Tonini C., 2012, MNRAS, 422, 3285

Planck Collaboration 2016, A&A, 594, A13

Pollo A., et al., 2005, A&A, 439, 887

Pontzen A., Governato F., 2012, MNRAS, 421, 3464

Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14

Pozzetti L., et al., 2010, A&A, 523, A13

Preskill J., Wise M. B., Wilczek F., 1983, Physics Letters B, 120, 127

Prevot M. L., Lequeux J., Prevot L., Maurice E., Rocca-Volmerange B., 1984, A&A, 132, 389

Puget P., et al., 2004, in Moorwood A. F. M., Iye M., eds, SPIE Conference Series Vol. 5492, Ground-based Instrumentation for Astronomy. pp 978–987, doi:10.1117/12.551097

Richstone D. O., 1980, ApJ, 238, 103

Richstone D. O., Tremaine S., 1988, ApJ, 327, 82

Riess A. G., et al., 1998, AJ, 116, 1009

Riley K. F., Hobson M. P., Bence S. J., 2006, Mathematical Methods for Physics and Engineering, Third Edition. Cambridge University Press

Rivolo A. R., 1986, ApJ, 301, 70

Rix H.-W., de Zeeuw P. T., Cretton N., van der Marel R. P., Carollo C. M., 1997, ApJ, 488, 702

Robotham A., et al., 2010, PASA, 27, 76

Rosas-Guevara Y. M., Bower R. G., Schaye J., Furlong M., Frenk C. S., Booth C. M., Crain R. A., Dalla Vecchia C., Schaller M., Theuns T., 2015, MNRAS, 454, 1038

Rubin V. C., Ford Jr. W. K., 1970, ApJ, 159, 379

Sales L. V., Navarro J. F., Theuns T., Schaye J., White S. D. M., Frenk C. S., Crain R. A., Dalla Vecchia C., 2012, MNRAS, 423, 1544

Salpeter E. E., 1955, ApJ, 121, 161

Sawala T., et al., 2016, MNRAS, 457, 1931

Scannapieco C., et al., 2012, MNRAS, 423, 1726

Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, MNRAS, 454, 2277

Schaye J., 2004, ApJ, 609, 667

Schaye J., Dalla Vecchia C., 2008, MNRAS, 383, 1210

Schaye J., et al., 2015, MNRAS, 446, 521

Schechter P., 1976, ApJ, 203, 297

Schwarzschild M., 1979, ApJ, 232, 236

Schwarzschild M., 1993, ApJ, 409, 563

Seljak U., Zaldarriaga M., 1996, ApJ, 469, 437

Shaposhnikov M., 2008, Journal of Physics Conference Series, 136, 022045

Silva L., Granato G. L., Bressan A., Danese L., 1998, ApJ, 509, 103

Simha V., Weinberg D. H., Conroy C., Dave R., Fardal M., Katz N., Oppenheimer B. D., 2014, preprint, (arXiv:1404.0402)

Simon J. D., Bolatto A. D., Leroy A., Blitz L., 2003, ApJ, 596, 957

Simon J. D., Bolatto A. D., Leroy A., Blitz L., Gates E. L., 2005, ApJ, 621, 757

Smolčić V., Zucker D. B., Bell E. F., Coleman M. G., Rix H. W., Schinnerer E., Ivezić Ž., Kniazev A., 2007, AJ, 134, 1901

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

Spergel D. N., Steinhardt P. J., 2000, Physical Review Letters, 84, 3760

Springel V., 2005, MNRAS, 364, 1105

Springel V., 2010, MNRAS, 401, 791

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726

Springel V., et al., 2005, Nature, 435, 629

Springel V., Wang J., Vogelsberger M., Ludlow A., Jenkins A., Helmi A., Navarro J. F., Frenk C. S., White S. D. M., 2008, MNRAS, 391, 1685

Starobinsky A. A., 1982, Physics Letters B, 117, 175

Stone J. M., Norman M. L., 1992, ApJS, 80, 753

Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, ApJS, 178, 137

Strigari L. E., Bullock J. S., Kaplinghat M., Diemand J., Kuhlen M., Madau P., 2007, ApJ, 669, 676

Strigari L. E., Bullock J. S., Kaplinghat M., Simon J. D., Geha M., Willman B., Walker M. G., 2008, Nature, 454, 1096

Strigari L. E., Frenk C. S., White S. D. M., 2010, MNRAS, 408, 2364

Strigari L. E., Frenk C. S., White S. D. M., 2014, preprint, (arXiv:1406.6079)

Subramanian S., Subramaniam A., 2012, ApJ, 744, 128

Swaters R. A., Verheijen M. A. W., Bershady M. A., Andersen D. R., 2003, ApJ, 587, L19

Syer D., Tremaine S., 1996, MNRAS, 282, 223

Taylor E. N., et al., 2011, MNRAS, 418, 1587

Teyssier R., 2002, A&A, 385, 337

Tonini C., Maraston C., Devriendt J., Thomas D., Silk J., 2009, MNRAS, 396, L36

Valenzuela O., Rhee G., Klypin A., Governato F., Stinson G., Quinn T., Wadsley J., 2007, ApJ, 657, 773

Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, MNRAS, 436, 3031

Vogelsberger M., Genel S., Springel V., Torrey P., Sijacki D., Xu D., Snyder G., Bird S., Nelson D., Hernquist L., 2014, Nature, 509, 177

Wadsley J. W., Stadel J., Quinn T., 2004, New A, 9, 137

Walker M. G., Peñarrubia J., 2011, ApJ, 742, 20

Walker M. G., Mateo M., Olszewski E. W., Peñarrubia J., Wyn Evans N., Gilmore G., 2009, ApJ, 704, 1274

Walker M. G., Mateo M., Olszewski E. W., Peñarrubia J., Wyn Evans N., Gilmore G., 2010, ApJ, 710, 886

Wang L., Li C., Kauffmann G., De Lucia G., 2006, MNRAS, 371, 537

Weinberg M. D., Nikolaev S., 2001, ApJ, 548, 712

Weinmann S. M., Neistein E., Dekel A., 2011, MNRAS, 417, 2737

Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, ApJ, 827, L23

White S. D. M., 1996, in Schaeffer R., Silk J., Spiro M., Zinn-Justin J., eds, Cosmology and Large Scale Structure. p. 349

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., Frenk C. S., Davis M., 1983, ApJ, 274, L1

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

Wolf J., Martinez G. D., Bullock J. S., Kaplinghat M., Geha M., Muñoz R. R., Simon J. D., Avedo F. F., 2010, MNRAS, 406, 1220

Worthey G., Ottaviani D. L., 1997, ApJS, 111, 377

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

Zehavi I., et al., 2002, ApJ, 571, 172

Zehavi I., et al., 2011, ApJ, 736, 59

Zel'dovich Y. B., 1970, A&A, 5, 84

Zhu Q., Marinacci F., Maji M., Li Y., Springel V., Hernquist L., 2016, MNRAS, 458, 1559

Zibetti S., Gallazzi A., Charlot S., Pierini D., Pasquali A., 2013, MNRAS, 428, 1479

Zwicky F., 1933, Helvetica Physica Acta, 6, 110

de Blok W. J. G., Walter F., Brinks E., Trachternach C., Oh S.-H., Kennicutt Jr. R. C., 2008, AJ, 136, 2648

de Lorenzi F., Debattista V. P., Gerhard O., Sambhus N., 2007, MNRAS, 376, 71

de Lucia G., Blaizot J., 2007, MNRAS, 375, 2

de la Torre S., et al., 2013, A&A, 557, A54

van der Marel R. P., Cretton N., de Zeeuw P. T., Rix H.-W., 1998, ApJ, 493, 613