



UCL

UNIVERSITY COLLEGE LONDON

Faculty of Mathematical and Physical Sciences

Department of Physics & Astronomy

Genetically modified galaxies: performing controlled experiments in cosmological galaxy formation simulations

Martin Pierre Rey

Thesis submitted in partial fulfilment of the requirements for the Degree of
Doctor of Philosophy of University College London

Supervisors:

Dr. Andrew Pontzen

Dr. Amélie Saintonge

Examiners:

Prof. George Efstathiou

Prof. Richard S. Ellis

6th of January, 2020

Martin Pierre Rey

Genetically modified galaxies: performing controlled experiments in cosmological galaxy formation simulations

Cosmology and Galaxy Formation, 6th of January, 2020

Examiners: Prof. George Efstathiou and Prof. Richard S. Ellis

Supervisors: Dr. Andrew Pontzen and Dr. Amélie Saintonge

UNIVERSITY COLLEGE LONDON

Astrophysics Group

Faculty of Mathematical and Physical Sciences

Department of Physics & Astronomy

Gower Street , WC1E 6BT

London, UK

Declaration

I, Martin Pierre Rey, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis. In particular, I would like to point out the following contributions:

- Work presented in Chapter 2 is published in Rey and Pontzen (2018).
- Work presented in Chapter 3 is published in Rey et al. (2019b).
- Work presented in Chapter 4 has been submitted and is available online in Rey et al. (2019a).
- Generating initial conditions for the results presented in Chapter 3, 4 and 5 has been performed using unpublished numerical methods, jointly developed by (alphabetically) Hiranya Peiris, Andrew Pontzen, Nina Roth, Stephen Stopyra and myself. This software will be released publicly in the near future (Stopyra et al. in prep).
- Simulations presented in Chapter 4 and 5 are the result of a collaborative effort with (alphabetically) Oscar Agertz, Matthew Orkney, Andrew Pontzen, Justin Read and myself. A description of our model and numerical methods is available in Agertz et al. (2019) of which I am a co-author.
- I also acknowledge the following people for contributing to this thesis through their hospitality during travels and useful discussions with myself: Lauren Anderson, Vasily Belokurov, Alyson Brooks, Aaron Ludlow, Andrey Kravtsov, Hiranya Peiris, Tjistik Starkenburg, Michael Tremmel, Anna Wright and the anonymous referee of Rey and Pontzen (2018).

London, UK, 6th of January, 2020

Martin Pierre Rey

Abstract

This thesis develops and applies a novel approach to studying the formation of galaxies in our Universe. Galaxies grow through gravitational amplification of early-Universe overdensities, within which gas reaches sufficient densities to trigger star formation. A galaxy's mass growth is therefore seeded randomly, originating from quantum inflationary perturbations. Understanding how this intrinsic stochasticity in histories couples with strongly non-linear astrophysics is key to interpreting the observed diversity of the galaxy population.

To provide new insights to this issue, we clarify and extend the “genetic modification” framework in Chapter 2. This approach generates alternative versions of a simulation's initial conditions, each version with a carefully engineered change to the galaxy's history. This in turn creates controlled experiments allowing us to construct a causal account of the galaxy's response to modifying its merger history. We introduce a new class of variance modifications aiming at improving control over several mergers. We then evolve these variance-modified initial conditions using the simulation code `RAMSES`, first studying dark matter halo formation (Chapter 3). We causally recover the known correlation between halo formation time and concentration when modifying the merger histories of two haloes, and further establish how late major mergers determine concentrations at fixed formation time.

We then turn to the formation of ultra-faint dwarf galaxies with high-resolution hydrodynamical simulations. Scanning through histories, we demonstrate that earlier forming ultra-faints have higher stellar mass today and predict a new class of highly diffuse ultra-faint galaxies which assemble through late mergers (Chapter 4). We finally use a larger suite of objects (Chapter 5) to show how ultra-faints growing sufficiently in dynamical mass after reionization can accrete gas and re-ignite star formation. We conclude that, by transforming cosmological histories into tuneable parameters, “genetically modified” experiments generate new insights on the complexity of dark matter halo and galaxy formation.

Impact Statement

This thesis presents a combination of original analytical and computational works aimed at deepening our understanding of the formation of galaxies in our Universe. The benefits of this corpus are primarily academic, developing and applying a new method to quantify the importance of mergers in shaping the diversity of galaxies. The theoretical underpinnings of the method (Chapter 2) were published in a leading astronomical journal and, together with the planned public release of our computational methods, will provide a reference for future use of the method throughout the astronomical community. Our applications presented in Chapter 3 and 4 have already achieved academic impact. They have been made publicly available through scientific publications, been cited by independent researchers across the world and led to new international collaborations for the author. Furthermore, Chapters 4 and 5 are likely to have long-lasting impact, predicting new classes of galaxies to be discovered in the coming decade by large-scale, hundred-million dollars, observational campaigns such as the Large Synoptic Survey Telescope.

In addition to advancing our academic understanding of galaxies, this thesis adds to the broader context of determining our cosmic origins and the scientific quest to understand the composition of our Universe. Chapter 3, 4, 5 have indirect consequences for pinpointing the nature of dark matter, motivating the author to present some of their content in a public engagement campaign. Visualizations derived from Chapter 3 for example appeared in a public lecture attended by over 200 participants and discussed on local radio and newspapers.

The computational nature of this thesis also reaches further than the academic world. Skills developed to achieve all chapters are applicable in numerous industrial and commercial contexts; they were thus passed along by training fellow undergraduate and postgraduate students through optional teaching, mentoring and tutorials over the past three years. Moreover, Chapter 2 presents an algorithm to efficiently solve constrained optimization problems over a large number of stochastic variables. The long-term impact of this algorithm is hard to quantify but such mathematical problems are routinely applied to inform decision making and cost minimization of commercial activities. Similarly, simulating hydrodynamics with massively-parallelised simulations over hundreds, if not thousands, of computers (Chapter 4 and 5) drives progress in

both high-performance computing and our understanding of fluid dynamics. These advances can then be applied to industrial activities, for example optimizing the cost and environmental efficiency of planes or predicting long-term weather and climate patterns.

Acknowledgements

Many people have participated to the completion of this journey and it is a great pleasure to now acknowledge those to whom I am deeply indebted. First and foremost my two supervisors, Andrew Pontzen and Amélie Saintonge. Your dedication and support has guided me during the past three years, providing me with the best possible conditions to construct this manuscript. You have been incredible mentors, changing the way I reflect about many problems (not just color schemes and fonts) and inspiring a taste for the academic world which might well be long-lasting. As obvious as it sounds, none of this would have been possible without you and I will never be able to thank you enough for it.

I am also in debt to staff members over the world that guided and mentored my evolution through the academic world, offering me support, time, inspiration and hospitality when I needed it. It is a pleasure to thank in particular Oscar Agertz, Lauren Anderson, Alyson Brooks, Richard Ellis, Andreu Font-Ribera, Andrey Kravtsov, Hiranya Peiris, Justin Read, Tjitske Starckenburg and Michael Tremmel.

An academic-style thesis will never reflect how much fun it has been to interact with astronomers, in particular those at UCL. Group A has felt like a home for the past three years largely due to the number of friends I have made there. Thanks goes to Andreu for smuggling spanish ham into our squash sessions, Arthur for introducing me to the definition of right-wing politics, Boris for your always acidic stories, Catarina for sincerely trying to make the world a better place, Chris for our mutual taste of geekyness and puppies, Constance for understanding important needs in life such as Michelin starred food, Felix for your unquenchable thirst, Guido because one does not meet a third English-third Belgian-third Italian space traveler every day, Harry for introducing me to English must-sees like Road Wars, Isabella and Tom J. for hiding amazing wittiness under quietness, Keir for your incomparably dry sense of humour, Koki for your taste of radiatively transferred badminton, Krishna for your resilience in supporting Arsenal over the years, Lewis for teaching me the infamous Southampton-Liverpool crossed accent, Lorne for our mathematically oriented sailing chats, Luisa because nothing is the same

when you are away, Niall for your love of absurd semantic discussions, Nicolas A. for being crazily amazing, Nicolas L. because we are Auvergnats before astronomers, Roger because I am sure you bet a penny you would be higher up this list than Luisa even though it is alphabetical, Selim for your undeniable energy at proving to me that you are exhausted, Teresita for never being politically correct, Tom F. for our private moaning sessions in the white quad, Tom W. because you enjoy rugby as decent people should, Will for being kind to everyone and hence to me, and finally Susan for the careful and brutal proofreading of this manuscript.

Crossing the Channel is not an easy decision to take, but one that I never regretted. With experiencing a new culture, food and weather came the opportunity to meet incredibly friendly individuals. This combination ensured made my time in the UK extremely enjoyable. So special credits goes to the IBN, and in particular Catherine, Han, James, Marta, Ningyu, Soraya, Will and Zita for refreshing my mind with beers when I needed a break from astronomy and to sailors Calum, Dan, David, Ian, Jade, Mark, Panos and Will for refreshing my mind with buckets of seawater when fulfilling my longtime dream of offshore racing. Going away would be easy if no one was waiting for you and it is a pleasure to thank longer time friends, Binh, Cecile, David, Gab, Hardouin, Henri, Hugues, Julian, LouisJ, Lucas, Mariette, Mathias, Max, Nhi and Ugo that supported me in difficult times, visited me and always ensured that going-backs would be met with excitement.

This list would never be complete without mentioning family. Your unconditional support, motivation and love is undeniably the biggest factor in my inspirations and achievements today. Mamouchka, Papadou, Margot, Simon, Zoé, I am incredibly grateful to what you have done for me and it is a pleasure to dedicate this thesis to you.

Contents

List of Figures	15
List of Tables	16
1. Introduction	17
1.1. The cosmological formation of galaxies	18
1.2. The need for improving our understanding of galaxies	22
1.2.1. Galaxies as probes of physics beyond the Standard model	23
1.2.2. Galaxies as probes of astrophysics	25
1.3. Cosmological simulations of galaxy formation	26
1.4. The initial cosmological density field	32
1.5. Thesis overview	35
2. Extending the genetic modification framework	39
2.1. Introduction	40
2.2. Linear constraints and modified fields	41
2.2.1. Constrained ensemble	41
2.2.2. Genetic modifications	44
2.3. Extension to quadratic modifications	46
2.3.1. Linearised solution	48
2.3.2. Joint quadratic and linear modifications	50
2.4. Demonstration	51
2.4.1. Defining an example modification	51
2.4.2. Results	52
2.5. Discussion	53
2.5.1. The advantage of quadratic over linear modifications	53
2.5.2. Multiple quadratic modifications	55
2.6. Conclusions	57

3. Sensitivity of dark matter haloes to their accretion histories	59
3.1. Introduction	59
3.2. Numerical setup	61
3.2.1. Initial conditions	61
3.2.2. N-body evolution	64
3.3. Controlling the smoothness of merger histories	65
3.3.1. Creating modified haloes	65
3.3.2. Results	66
3.4. Response of haloes to changes in their merger history	71
3.4.1. Calculating halo secondary properties	71
3.4.2. The evolution of concentration and formation time	72
3.4.3. Halo environment	75
3.4.4. Halo spin	76
3.5. Conclusion	77
4. The origin of scatter in ultra-faint dwarf stellar masses and surface brightness	81
4.1. Introduction	81
4.2. Genetically modified dwarf galaxies	82
4.2.1. Initial conditions	82
4.2.2. Galaxy formation model	83
4.3. Growing the stellar mass of ultra-faints	86
4.4. Generating scatter in the stellar mass-halo mass relation	88
4.5. The formation of a diffuse ultra-faint	90
4.6. Conclusion	92
5. Rejuvenating star formation in ultra-faint dwarfs	95
5.1. Introduction	95
5.2. Numerical setup	96
5.3. Late time star formation in ultra-faint dwarfs	99
5.4. The fate of gas during the rejuvenation process	100
5.5. Rejuvenation and mass assemblies	104
5.6. Discussion and conclusions	106
6. Conclusions and future prospects	109
6.1. Controlling cosmological merger histories	109

6.2. A new angle on dark matter halo formation	111
6.3. The diversity of ultra-faint dwarf galaxies	112
6.4. Future work	114
A. Appendix: Supplementary interpretations of quadratically modified fields	117
A.1. Constructing constrained ensembles for quadratic constraints	117
A.2. Geometrical interpretation	120
Bibliography	123

List of Figures

1.1.	Growth of structure across the Universe	19
1.2.	Cutting through allowed parameter space by combining independent datasets	23
1.3.	Illustration of the dynamic range probed by galaxy formation	28
1.4.	Trade-off between volume and resolution for modern cosmological simulations	30
2.1.	Overview of genetic modifications, contrasted with constrained realisations .	42
2.2.	Toy-model example of a quadratically modified density field	53
2.3.	Comparison of pure linear against combined linear-quadratic genetic modifications	54
3.1.	Overview on how to generate a genetically modified halo	62
3.2.	The growth of virial mass for Halo 839 and the response of its merger history	67
3.3.	The growth of virial mass for Halo 740 and the response of its merger history	69
3.4.	Response of concentrations and formation times for the two modified halo families	73
3.5.	Response of halo spins and environments for the two modified halo families .	76
4.1.	Growth of dynamical mass and stellar mass assembly over cosmic time for our reference ultra-faint and its genetically modified counterparts	87
4.2.	Scatter in stellar masses at fixed halo mass arising from different histories . .	89
4.3.	Response of an ultra-faint's absolute magnitude and half-light radius	90
5.1.	Growth of dynamical mass and star formation histories for a suite of ultra-faints	98
5.2.	Time evolution of the inner gas mass with increasing final dynamical mass . .	101
5.3.	Phase diagram of gas at two key stages of the rejuvenation process	103
5.4.	Sensitivity of rejuvenation to modifying the growth of dynamical mass	105
A.1.	Geometrical interpretation of genetic modifications	121

List of Tables

3.1. List of simulations and their associated modifications used in Chapter 3	64
5.1. List of simulations and their associated final properties presented in Chapter 5	97

Introduction

Galaxies are cosmological building blocks, the fundamental unit into which stars and gas are collected to form the largest observable structures in the Universe. But do we understand how galaxies form in our Universe? As for any physical question that we wish to solve, we require three ingredients – (i) the initial conditions from which galaxies will eventually appear from – (ii) a physical model evolving these initial conditions to today and (iii) observational evidence to validate the combination of (i) and (ii).

In the context of galaxy formation, the initial conditions are firmly established by the modern concordance model of cosmology, Λ CDM. Though it introduces two poorly understood components, dark matter (CDM) and dark energy (Λ), this model is now verified by numerous independent measurements, thus providing a robust framework in which to anchor a galaxy formation model (Section 1.1).

Great difficulties in understanding the formation of galaxies arise from (ii) the physical model. Galaxies spread across an extended range of scales and their formation is a strongly non-linear process: they are highly sensitive to cosmological environments on super-Mpc scales and astrophysical processes on scales down to black hole event horizons. The tool of choice to model non-linear hydrodynamics is numerical simulations, but this vast dynamic range makes ab initio numerical models intractable. We therefore have to resort to simplifying assumptions in our modelling, adapted to the currently available computing power (Section 1.3), and are thus a long way from understanding how the properties and abundances of galaxies are set from cosmological initial conditions.

Refining our physical understanding of galaxy formation is paramount for the coming decade. The ability of dark matter and dark energy to fit cosmological data belies our lack of understanding as to their nature. Galaxies, through their properties and distribution across the Universe, provide a sensitive probe to these dark components. Moreover, the arrival of new observatories such as JWST and ALMA will further provide key inputs into open astrophysical questions, in particular how stars and black holes co-evolve with galaxies throughout cosmic times (Section 1.2). This ability of galaxies to jointly inform physics beyond the standard model and astrophysics strongly advocates for deepening our understanding of their formation. We

review the current boundaries of our knowledge in this introduction and argue for the need for new methods to transform observations of galaxies into physical inferences.

1.1 The cosmological formation of galaxies

The major accomplishment of modern cosmology is establishing a concordance model, pinpointing the contents of the Universe to sub-percent precision. This successful cosmological model provides the framework for galaxy formation through two key components:

1. A mechanism generating small amplitude perturbations in the early Universe.
2. An expanding background cosmology in which these perturbations amplify under gravity, eventually creating the clumpy, “webby” structure of galaxies that we observe throughout the Universe (Figure 1.1).

We start by briefly reviewing how cosmological structure and galaxy formation proceeds from these two ingredients. We wish to provide insights into the relevant physical processes of galaxy formation rather than an exhaustive description and refer to Mo et al. (2010) and Somerville and Davé (2015) for in-depth explorations.

Seeding structure formation

A fundamental assumption of standard Big Bang cosmology is the cosmological principle, i.e. spatial homogeneity and isotropy. However, the very existence of galaxies, and generally structure across the Universe (see Figure 1.1), requires some deviation from perfect uniformity. In the standard picture of structure formation through gravitational amplification, these deviations must be small and therefore generated in the early Universe. Big Bang cosmology does not propose such a seeding mechanism without resorting to specific initial conditions.

The “initial condition” problem of seeding matter perturbations is most elegantly solved by invoking an inflationary mechanism. During this early phase, the Universe undergoes accelerated expansion powered by the vacuum energy or the “slow-roll” of at least one quantum field (Guth 1981; Linde 1982). Fluctuations in the field are then inflated exponentially to cosmological scales with an amplitude set by the value of the field when they exit the horizon (Hawking 1982; Guth and Pi 1982). These random fluctuations provide the early-Universe seeds for future structure formation.

Inflationary mechanisms are somewhat speculative, invoking unknown states of matter in the early Universe. Their key appeal relies in their ability to, at once, solve the classical issues of Big-Bang cosmology (the “horizon”, “flatness” and “monopole” problems), *and* generate

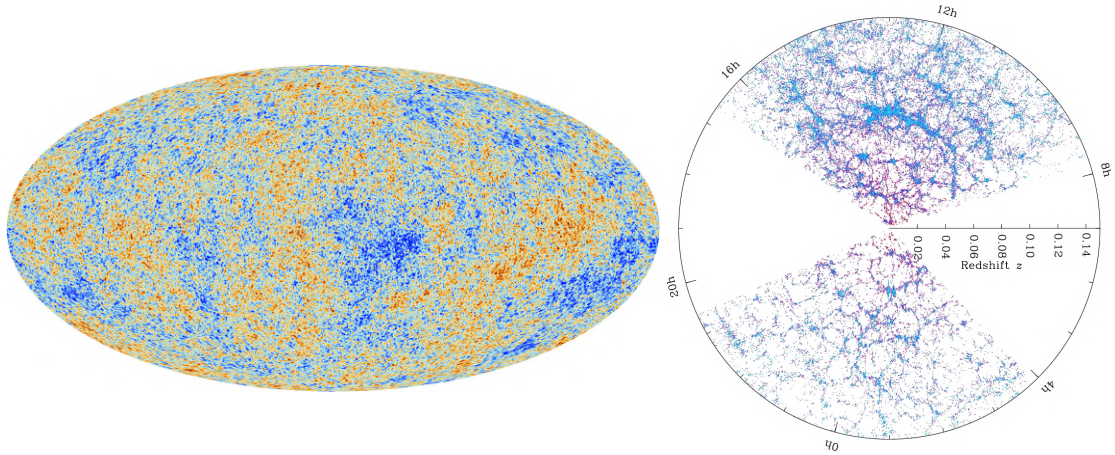


Fig. 1.1.: Density perturbations at the time of the cosmic microwave background (left panel). In this early picture of structure growth, inflationary density perturbations are still small and linear. They will then gravitationally amplify to eventually form the non-linear galaxy distribution observed throughout the Universe (right panel). Statistical patterns in both early and late density perturbations allow us to constrain the constituents of our Universe. Image credit: Planck Collaboration and SDSS collaboration.

Gaussian density perturbations with a near Harrison-Zel'dovich spectrum (Harrison 1970; Zeldovich 1972). The amplitude and spectrum of these predicted perturbations can then be compared to those observed in the cosmic microwave background (hereafter CMB) to constrain inflationary models.

The growth of cosmological structure

Once inflation ends and energy is transferred back into radiation, density fluctuations start evolving in the expanding cosmological background. In particular, non-relativistic density perturbations grow with time through gravitational amplification. The specific rate of growth can be analytically derived in the context of linear cosmological perturbation theory (see Mukhanov et al. 1992; Bernardeau et al. 2002 for extensive reviews) and is sensitive to the cosmological model. This sensitivity in turn allows us to constrain the Universe's content by measuring the amplitude and patterns of linear perturbations.

The first observable signature of structure growth is contained in the CMB, the background radiation produced when the Universe has cooled sufficiently to allow the primordial plasma to recombine and become transparent to photons. Patterns in the observed temperature anisotropies of the CMB, together with their lensing by foreground structure, provide the strongest standalone constraints on our cosmological model. From these measurements, we can fit the present-day cosmological densities of dark energy: $\Omega_\Lambda = 0.6847 \pm 0.0073$, total matter: $\Omega_m = 0.3153 \pm 0.0073$ and ordinary matter: $\Omega_b = 0.0493 \pm 0.00033$ (Planck Collaboration et al. 2018; see Hinshaw et al. 2013; Smoot et al. 1992 for previous measurements) to exquisite

precision. These values, and their error budget, show with high significance that dark energy makes up $\sim 70\%$ of the total content of the Universe and that $\sim 85\%$ of matter is not in the form of visible baryons.

Gravitational amplification is a run-away process: overdense regions attract their surroundings more strongly than a mean-density region, hence becoming ever more overdense. Linear matter perturbations continue growing after recombination, eventually become non-linear, and start gravitationally collapsing. In a cosmological model containing baryonic and dark matter such as our preferred Λ CDM, each matter perturbation starts with a mixture of both states. At the time of non-linear collapse, the evolution of “ordinary matter” (i.e. pressurized primordial gas) and “dark matter” (i.e. pressureless matter) starts to diverge.

Due to its collisionless nature, dark matter does not shock when undergoing collapse but still relaxes to near-equilibrium (Lynden-Bell 1967). Resulting structures, named “dark matter haloes”, set the potential wells into which cosmic gas then flows. Collapsed dark matter structure is therefore the invisible backbone of visible cosmological structure. As we will see in Chapter 3 and onwards, the properties and abundances of dark matter haloes are tightly linked with those of galaxies, leading to the so-called galaxy-halo connection (see Wechsler and Tinker 2018 for a review).

Gas cooling and star formation

Unlike dark matter, gas is supported by thermal pressure and thus shocks under gravitational collapse, raising its temperature (Binney 1977; Rees and Ostriker 1977). Subsequent evolution depends on a crucial ingredient for galaxy formation, the ability for gas to cool.

Processes such as recombination, bremsstrahlung and collisional ionization can decrease thermal support by allowing atoms to radiate kinetic energy away. Rates at which energy is radiated are sensitively dependent on the temperature, density and chemical composition of the gas, making cooling a strong function of local gas conditions (Sutherland and Dopita 1993). Because most of these processes involve interactions between two particles, cooling is increasingly effective at higher densities, scaling with the gas density squared. This in turn creates another runaway process: if gravitationally shocked gas can cool, its internal temperature and hence pressure support reduces. The gas then flows inwards towards the centre of the potential where densities are higher, making cooling even more efficient (see Fabian 1994 for a review on cooling flows).

This condensation of cold gas at the centre of dark matter haloes ultimately leads to densities high enough for gas clouds to collapse under their own gravity and fragment to form protostars (White and Rees 1978). This basic picture provides a strong starting point in our goal to build a bridge between cosmological initial conditions, star formation and the assembly of galaxies.

Feedback

However, the above picture fails catastrophically to reproduce the observed properties of galaxies. In the cold dark matter scenario, small haloes collapse first and then merge hierarchically to create bigger structures. Since cooling is efficient at high densities, dense haloes at high redshift are predicted to transform the near-totality of their gas into stars (White and Rees 1978). This is blatantly in contradiction with observations: only $\sim 20\%$ of the baryon content in the Universe is in stars (Fukugita et al. 1998), most of the rest in gas. We therefore require physical processes to make star formation inefficient and prevent gas from “over-cooling”. These processes, broadly known as feedback, are required to either heat gas or remove it from the dense regions in which cooling is efficient.

An obvious candidate for feedback is provided by massive stars: they emit UV-radiation which heats their surroundings and can also sweep gas away by driving stellar winds and exploding violently as supernovae (Dekel and Silk 1986). Active black holes at the centre of galaxies (hereafter AGNs) also release vast amounts of energy, potentially sufficient to heat the whole gas content of the galaxy (Silk and Rees 1998). More exotic sources of gas pressure such as magnetic fields, cosmic rays or radiative transfer effects are also routinely considered in modern models of galaxy formation. Pinpointing the most important sources of feedback and, more importantly, how they couple to their surrounding gas to regulate star formation, is currently the largest uncertainty in modelling galaxy formation (see Naab and Ostriker 2017 for a review).

It should be apparent at this point that building an *ab initio* model of galaxy formation is complex, involving many intertwined astrophysical and cosmological processes. Before we turn to the details of how this modelling is performed with numerical simulation in Section 1.3, we first argue that improving our understanding of galaxies is paramount as they provide key insights in fundamental physics questions.

1.2 The need for improving our understanding of galaxies

The best-fit Λ CDM model to the CMB is highly unsatisfactory as it invokes two unexplained components, dark energy and dark matter. However, its true strength relies in its ability to jointly fit the spectrum of the CMB and a range of observations mapping the growth of structure across multiple epochs in the Universe. The same model parameters successfully explain the spatial distribution of galaxies (Alam et al. 2017) and gas (Palanque-Delabrouille et al. 2013; Iršič et al. 2017), the distortions of galaxy shapes from gravitational lensing (van Uitert et al. 2018; DES collaboration et al. 2018; Hikage et al. 2019) the dimming of supernovae (Betoule et al. 2014; Scolnic et al. 2018) and the abundance of galaxy clusters (Planck Collaboration et al. 2016b; de Haan et al. 2016). When combined together (e.g. Figure 1.2), these independent datasets provide an overwhelming body of evidence for the existence of dark matter and dark energy. Modern cosmology is no longer in the regime of establishing *if* these components exist but rather pinpointing *how and why*.

A physical understanding of the nature of these two components is indeed lacking. The mathematically simplest form of dark energy is a cosmological constant, with fixed energy density across space and time. A natural interpretation would be for this constant to arise from the quantum vacuum energy. However, the measured density of dark energy is at least 10^{41} smaller than the theoretically expected quantum value (Weinberg 1989). Alternative models have therefore been proposed (see Weinberg et al. 2013 for a review), for example based on anthropic arguments in an inflationary multiverse (Efstathiou 1995), additional scalar fields with effective negative pressures (Ratra and Peebles 1988) or the breakdown of general relativity on cosmological scales (Carroll et al. 2004). So far, no scenario has been favoured over a pure cosmological constant by observational evidence nor theoretical preference. Clarifying the nature of dark energy has motivated extensive funding for the current and next generation of cosmological surveys (Albrecht et al. 2006), which might be bearing fruits by revealing tensions between datasets (e.g. Bernal et al. 2016).

Likewise, dark matter is usually assumed to be a novel particle beyond the standard model of particle physics. The historically motivated weakly-interacting cold particle has, despite extended efforts, escaped detection (LUX collaboration et al. 2017; Xenon collaboration et al. 2018). This lack of detection has triggered a policy of “no stone left unturned” but the range of allowed alternative models is vast. Dark matter could be warmer, fuzzier, interacting with itself or baryons or having an altogether different axionic or black hole nature (see Bertone and




Fig. 1.2.: The power of combining independent datasets based on the dimming of supernovae (green, DES collaboration et al. 2019) and temperature patterns in the CMB (brown, Planck Collaboration et al. 2018) to find overwhelming evidence (red) of a flat Universe made at $\sim 70\%$ of dark energy and $\sim 30\%$ of matter. Tighter constraints could be obtained by further including lensing of the CMB by foreground structures and measurements of the spatial distribution of galaxies. Reproduced from DES collaboration et al. (2019), figure 2.

Tait 2018 for a review). This extensive theoretical landscape strongly advocates for a combined effort from all sensitive dark matter probes to reduce and constrain the allowed model space.

1.2.1 Galaxies as probes of physics beyond the Standard model

Galaxies, and in particular dwarf galaxies, provide just such a sensitive probe on the nature of dark matter. This sensitivity is best illustrated by the historical mismatch between the properties of observed dwarf galaxies and those predicted from a pure CDM simulation (see Bullock and Boylan-Kolchin 2017 for a review):

1. The number of observed satellites around the Milky Way is vastly smaller than the number of dark matter subhaloes found in CDM-only simulations (Klypin et al. 1999; Moore et al. 1999).
2. The central density profiles of observed dwarf galaxies is flat (Flores and Primack 1994; Moore 1994) when CDM-only simulations predict it to be rising towards the centre as $\rho \propto r^{-1}$ (Navarro et al. 1997) .

Early solutions to these problems motivated modifications to the nature of dark matter. For example, making dark matter warmer, i.e. less massive and thus relativistic when decoupling, reduces the abundance of subhaloes and helps to solve Problem 1 (Bode et al. 2001). Similarly, a dark matter with stronger self-interaction, i.e. allowing elastic scattering at larger cross sections, would naturally produce flat density profiles and solve Problem 2 (Spergel and Steinhardt 2000).

Unfortunately, modifying the nature of dark matter is not the only possible solution to these problems. Improvements in our understanding of dwarf galaxy formation have now shown that they can be addressed within the context of pure CDM. For Problem 1, the number of observed satellites around the Milky Way has more than doubled since the original statement following the improvements of optical imaging over the years (SDSS, PANSTARS, DES, HSC and in the future LSST). In addition, one can reduce the tension by either preventing galaxies from forming in small dark matter haloes (Efstathiou 1992; Somerville 2002) or increasing the destruction of dark matter subhaloes by including the additional gravitational potential of the disk (Read et al. 2006). Similarly, for Problem 2, repeated outflows from supernovae explosions couple dynamically with the surrounding dark matter, flattening the central dark matter density profiles (Pontzen and Governato 2012; Di Cintio et al. 2014).

Despite now having candidate solutions within the context of pure Λ CDM, these historical mismatches serve two purposes: (i) they demonstrate the potential of dwarf galaxies to inform us about the nature of dark matter but (ii) that this information is degenerate with our incomplete understanding of dwarf galaxy formation. The advent of ever deeper optical imaging with e.g. LSST will prove revolutionary in the number and quality of dwarf galaxies' observations. For example, they will likely complete the census of dwarf galaxies within the Milky Way radius and start discovering faint dwarfs in the field (Drlica-Wagner et al. 2019) allowing us to re-evaluate the statistical significance of Problem 1 in new unprobed regimes. The need for accurate dwarf galaxy models to interpret these observations strongly motivates the studies of Chapter 4 and 5.

In addition, improving cosmological constraints to pinpoint the nature of dark energy will require next-generation surveys to extract information on scales affected by non-linear structure and galaxy formation. Uncertainties arising from our incomplete knowledge of galaxy formation is becoming a leading source of systematic error, potentially biasing cosmological inferences. For example, the central black hole of a galaxy can launch powerful outflows, redistributing matter on observable scales for weak lensing (Chisari et al. 2018). The orientation of galaxies with respect to their neighbours is also a byproduct of galaxy formation and creates additional

correlations in the shapes of galaxies on the sky (Joachimi et al. 2015). These last two effects are worrying for current weak lensing surveys such as KIDS or DES and will become leading order systematics for future instruments such as HSC, LSST and EUCLID. More generally, cosmological inferences rely on observing galaxies and assuming they track the underlying, invisible dark matter structure. However, galaxy formation proceeds differently in dark matter haloes with varying internal properties, histories and cosmological environments. How to model the galaxy-halo connection and the subsequent galaxy “biasing” is one of the leading question for galaxy clustering surveys (e.g. BOSS, DESI; see Wechsler and Tinker 2018 for a review and Chapter 3).

1.2.2 Galaxies as probes of astrophysics

So far, we have described galaxy formation as a nuisance preventing us from precise inferences on the nature of dark matter and dark energy. However, we have seen in Section 1.1 that galaxies contain stars, black holes and other astrophysical constituents thereby granting us a unique opportunity to study how they grow, evolve and couple with their galactic host.

In particular, we saw that galaxies require additional heating sources to regulate star formation. Explosions from massive stars and activity from central super-massive black holes have become the prime candidates to inject this required energy. However, this scenario opens a series of questions at the boundary of our knowledge: (i) where and how giant molecular clouds turn into stars (Krumholz 2014); (ii) how stellar, intermediate and super-massive black holes form and grow over cosmic time (Volonteri 2010); (iii) how does feedback couple with the surrounding galaxy and shape its color and morphology (Somerville and Davé 2015; Naab and Ostriker 2017).

Leaps of progress towards answering these questions have historically been provided by revolutions in instrumentation techniques, for example with SDSS and multi-object spectroscopy or space imaging with Hubble. This trend is continuing today, with the now widespread use of Integral Field Units (SAMI, Croom et al. 2012; MANGA, Bundy et al. 2015; MUSE, Bacon et al. 2015; and in the future KCWI, Morrissey et al. 2018 or NIRSpec onboard JWST, Bagnasco et al. 2007). These new instruments allow us to access both spatial and kinematic information, resolving the dynamics of gas and stars in the interior of thousands of galaxies. The arrival of interferometry in the sub-millimeter wavelengths (ALMA) has also opened access to mapping the cold gas and dust content of galaxies to unprecedented resolution and depth (e.g. Walter et al. 2016; Laporte et al. 2017). Furthermore, feedback-driven outflows do not stop at the boundaries of optical galaxies but rather expand and deposit metals and energy in the

surrounding circumgalactic medium (CGM). Ambitious programs have started unlocking the information content of the CGM by studying gas in absorption of background quasar light (COS; Tumlinson et al. 2013), direct emission (e.g. Hayes et al. 2016; Lokhorst et al. 2019) and in the future via X-ray wavelengths (ATHENA; Barcons et al. 2017).

The advent of gravitational wave astronomy (LIGO Scientific Collaboration and Virgo Collaboration 2017) has further opened a unique window to study the formation and evolution of black holes. Pulsar timing arrays (e.g. Arzoumanian et al. 2018) and in the future LISA (Klein et al. 2016) will probe the mergers of super-massive black holes from $z \sim 10$ to today, providing invaluable insights into their co-evolution with galaxies across cosmic times. We will gain new clues about how black holes can become sufficiently massive to power the brightest quasars at $z \sim 6$ (Fan et al. 2006; Mortlock et al. 2011; Bañados et al. 2018) and the combination with JWST will shed light on the respective contribution of black holes and galaxies in reionizing the Universe (Natarajan et al. 2017).

This large amount of incoming information is an exciting prospect to deepen our understanding of how galaxies assemble. However, turning this information into a physical understanding of dark matter, dark energy, stars and black holes requires us to interpret, or even better predict, future observations. Making such predictions requires an ab initio, self-consistent model which, as we have seen in Section 1.1, is challenging. We will now review the current state of our best numerical simulations and argue that their limitations motivate finding novel ways to investigate the formation of galaxies.

1.3 Cosmological simulations of galaxy formation

The broad picture of Section 1.1 states that cosmological perturbations made of dark matter and gas grow due to gravity until they reach sufficient densities to ignite star formation. There are therefore three key physical processes to solve in order to model the formation of galaxies over the age of the Universe: gravity for dark matter and gas, hydrodynamics and astrophysics. All three processes are extremely non-linear, intrinsically coupled to each other and can therefore only be modelled self-consistently using numerical simulations. We now briefly describe common modelling strategies for each process.

Gravity

Typical CDM models predict dark matter to be collisionless on astrophysical scales and dark matter is therefore modelled as a collisionless fluid discretised into point masses. We can safely

ignore the detailed particle physics nature of CDM for our applications¹, as the mass of a dark matter particle is typically $\sim \text{GeV}$, i.e. 10^{57} times smaller than a solar mass.

The discrete set of particles is then evolved in a so-called “N-body simulation” with pure Newtonian gravity², determining the force acting on each particle through the Poisson equation. Numerically solving this equation is non-trivial, as Newtonian gravity has an infinite range and diverges when two point masses are close together. We thus need to perform approximations to make the calculation tractable (see Bertschinger 1998; Dehnen and Read 2011 for more extensive reviews), for example by grouping particles in a hierarchical tree and approximating long-range forces through their multipole moments (Barnes and Hut 1986) or computing the gravitational potential on a mesh through the Fourier transform of the density field (Hockney and Eastwood 1988). Though considered cheap in the context of galaxy formation, N-body simulations are already a significant computational challenge with their largest iterations (e.g. Klypin et al. 2011) requiring millions of hours on supercomputers.

Hydrodynamics

Gas follows the collapse of dark matter but with an intrinsic pressure and temperature. Solving for the dynamics and thermal evolution of gas therefore requires the Poisson equation to be coupled with the Navier-Stokes equations of hydrodynamics³. This coupling enormously increases the complexity and cost of the simulation, as hydrodynamics is notoriously non-linear and hard to solve numerically.

To render the calculation tractable, several approximations have been developed over the years. In the Lagrangian approach of “smoothed particle hydrodynamics” (see Monaghan 1992; Springel 2010b for reviews), gas is discretised into particles and every fluid quantity is made local through a kernel average across its neighbouring particles. Examples of widely used codes implementing this method are GADGET (Springel 2005) and GASOLINE (Wadsley et al. 2004). Using a different philosophy, Eulerian methods discretize gas into grid cells and compute the advection of fluid quantities at the cell boundaries (see Teyssier 2015 for a review). The grid is typically adaptatively refined to increase the dynamic range, as in RAMSES (Teyssier 2002; Chapters 4 and 5) and ENZO (Bryan et al. 2014). Recently, attempts at combining these two

¹Alternative dark matter models can require additional modelling of quantum or particle physics effects. See for example Nori and Baldi (2018) for a fuzzy dark matter solver.

²Relativistic corrections are negligible in the context of pure Λ CDM but could be required for extensions with relativistic species such as neutrinos or more general gravity theories. See Adamek et al. (2016) for a relativistic N-body solver.

³Cosmological simulations typically neglect the effect of viscosity (unless artificially introduced) and rather solve the Euler equations.




Fig. 1.3.: Slice through the simulation volume of the Eagle simulation (Schaye et al. 2015) illustrating the dynamic range probed by galaxy formation. Such simulations are small by cosmological standards (~ 50 Mpc) and barely resolve the internal structure of galactic disks and their interstellar medium (inset panel). They are nonetheless amongst the most ambitious projects undertaken by the community, requiring tens of millions of computer hours. Credit Eagle Project.

historical approaches has led to using deformable meshes, following and adapting to the fluid's evolution, such as implemented in AREPO (Springel 2010a; Weinberger et al. 2019).

These different approaches at solving the hydrodynamics equations have respective benefits and drawbacks that are continuously under development. For example, Eulerian methods were traditionally stronger at resolving shocks and discontinuities (e.g. Agertz et al. 2007), but more recent implementations of Lagrangian methods (e.g. Hopkins 2013; Wadsley et al. 2017) have largely closed the gap. More importantly, modern code comparison projects (e.g. Kim et al. 2014; Kim et al. 2016) have demonstrated that uncertainties in modelling the internal astrophysics far outweighs differences in numerical schemes. Galaxy simulations using similar hydrodynamics solvers but different astrophysical implementations yield significantly different results. We will now see that though critical improvements have been made in the modelling of feedback and star formation, this remains the biggest challenge of modern cosmological simulations.

Feedback

We have seen that a successful model of galaxy formation requires a model of internal astrophysical processes. With this task comes the greatest challenge for numerical simulations,

the involved dynamic range (see Figure 1.3 for a visual example). During its formation, a single galaxy accretes low density gas from cosmological scales (\geq Mpc) which then cools and compresses in the circumgalactic medium (~ 100 kpc) then further in the interstellar medium (~ 1 kpc), then further in giant molecular clouds (~ 100 pc) and finally into star clusters (~ 1 pc). Resolving star formation in a single galaxy with tens of resolution elements therefore requires 8 orders of magnitude in dynamic range. This is already outside the reach of modern computers, before we start discussing a study of the galaxy population across a cosmological volume (~ 100 Mpc), modelling the energy input from supernovae ($\sim 10^{-2}$ pc) and black holes ($\sim 10^{-6}$ pc) and adding self-consistent treatments of radiation and magnetic fields.

The direct consequence of this dynamic range is the need to construct effective “subgrid” models to represent relevant astrophysical processes unresolved by current cosmological simulations. Examples range from, but are not limited to, star formation; stellar feedback in various forms (e.g. photo-ionization, photo-heating, winds, supernova explosions); black hole formation, growth, dynamics and feedback; cosmic ray creation, propagation and interaction with gas; metal mixing and turbulence. The huge body of literature on this topic has been extensively reviewed elsewhere (e.g. Somerville and Davé 2015; Naab and Ostriker 2017), so we will restrict ourselves to describing relevant implementations in Chapter 4 and 5. We now provide a more global view of successes and issues of subgrid modelling.

With the first subgrid models (e.g. Katz et al. 1992; Navarro and White 1993) came the realisation that results were highly sensitive to the chosen implementation and that stars were vastly over-produced in the central regions of galaxies (e.g. Sommer-Larsen et al. 1999; Balogh et al. 2001; Thacker and Couchman 2001; Robertson et al. 2004). These two conclusions have stuck to this day: (i) feedback needs to be strong and effective to prevent over-cooling and produce realistic galaxy populations – (ii) different implementations of the same physical process yield diverging results, which explicitly or implicitly depend on the simulation’s resolution. Strong feedback could then be seen as a prediction of cosmological simulations. However, the robustness of this prediction is hindered by (ii) and by the observational difficulties in confirming the widespread efficiency of feedback (Veilleux et al. 2005; Fabian 2012).

Despite this slightly alarmist statement, undeniable progress has been achieved in the past decade by tuning parameters of subgrid models to reproduce chosen observables such as the galaxy stellar mass function. This approach has allowed modern cosmological simulations to create realistic galaxy populations (Vogelsberger et al. 2014; Schaye et al. 2015; Dubois et al. 2016; Nelson et al. 2019b), with a diversity of stellar masses, colours and morphologies.

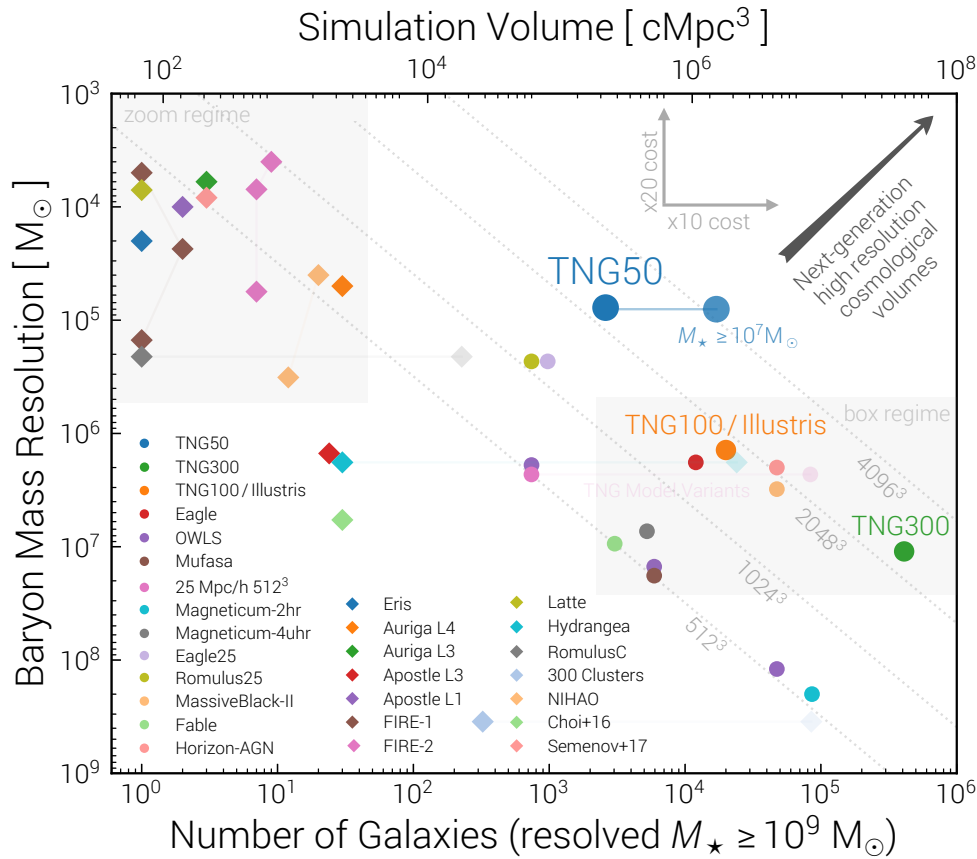


Fig. 1.4.: Illustrating the trade-off of modern cosmological simulations between large numbers of galaxies at low resolution (bottom right corner) and zooming in on single objects at high resolution (top left corner). Increasing particles numbers (dashed lines) by a factor 2^3 usually leads to a 16-fold increase in the simulation’s cost and hence 8 years of improved computing power (Moore 1965). The dynamic range required by galaxy formation is therefore inaccessible in the foreseeable future. Reproduced from Nelson et al. (2019a), figure 1.

However, the physical interpretation of such simulations is delicate and debated, as quantifying the impact of the fitting procedure on final predictions is problematic.

Another route for cosmological simulations is provided by the “zoom” resimulation method, in which all computing power is concentrated into one object. The increased resolution allows for finer subgrid modelling with fewer built-in assumptions. For example, increased resolution allows us to directly resolve the multi-phase structure of the interstellar medium (e.g. Agertz and Kravtsov 2015; Christensen et al. 2016; Hopkins et al. 2018 for disk galaxies) or the cooling radius of supernovae explosions (Wheeler et al. 2018; Agertz et al. 2019; Chapters 4 and 5), thus avoiding modelling it effectively. Even with zoom simulations, however, we are far from achieving the dynamic range required by galaxy formation. The trade-off between resolution, i.e. a more refined subgrid modelling, and volume, i.e. increased statistical power, will continue to drive numerical studies for the coming decades (see Figure 1.4).

An additional complication to the dynamic range arises from the stochastic nature of cosmological initial conditions. For a given galaxy, its growth of mass (i.e. mergers and accretion) and environment are seeded randomly, from density perturbations generated during the inflation era (see Section 1.1). This intrinsic stochasticity is problematic due to its associated number of degrees of freedom: a modern simulated galaxy originates from randomly drawing millions of resolution elements, each of which could be considered a degree of freedom. In practice, this freedom materialises in the wide range of possible histories, assemblies and environments for a galaxy in a Λ CDM universe. This diversity of histories then couples with the internal astrophysics to create the observed diversity in the galaxy population.

One solution to probe this interplay has been to simulate large numbers of galaxies (Vogelsberger et al. 2014; Schaye et al. 2015; Dubois et al. 2016; Nelson et al. 2019b), sampling many possible histories. But as we have seen, this comes at the cost of coarse subgrid modelling. Furthermore, the inherent nature of statistical sampling makes it hard to causally interpret the emergent correlations in galaxy properties because, by construction, every degree of freedom is modified from one galaxy to the next. Another option, entirely bypassing this issue, is to replace cosmological initial conditions with “idealised” simulations. Rather than starting from inflationary perturbations, idealised simulations are initialised from vacuum, fully controlled configurations (e.g. a stable disk; Hernquist 1993). With full control over their initial conditions, such setups allow us to isolate a particular physical effect of interest, for example studying the response of star formation and central black holes as we vary the parameters of an infalling merger (e.g. Springel and Hernquist 2005; Di Matteo et al. 2005). However, the long-term evolution of galaxies is primarily regulated by the balance between cosmological accretion of gas and feedback, the former being by construction neglected in idealised simulations thus limiting their interpretation.

A direct improvement for simulations would be to create a hybrid setup, providing control over histories and mergers as in idealised simulations, while still using cosmological initial conditions. This goal is the leading theme of this thesis and we will argue throughout that such a setup generates new insights into the formation of galaxies, complementary to more traditional approaches. Achieving control from cosmological initial conditions requires us to understand how, at least approximately, the early linear Universe relates to a halo’s non-linear history. We now turn to describing how we can extract the information contained in the linear initial conditions to engineer specific non-linear objectives.

1.4 The initial cosmological density field

The linear growth of cosmological perturbations through gravitational amplification does not need to be solved numerically and has been amply described analytically using cosmological perturbation theory (Mukhanov et al. 1992; Bernardeau et al. 2002). Because density perturbations originate from random quantum fluctuations, a statistical approach must be undertaken. The observable Universe is then modelled as a stochastic realization drawn from an ensemble of possibilities, leading to the introduction of “cosmological random fields” (e.g. density, velocity or gravitational fields). These fields are stochastic variables, similar in mathematical nature to quantum fields, and their covariances such as the power spectrum or the correlation function define the expected ensemble average to which our realization of the observable Universe can be compared.

The overdensity field, $\delta(r)$ defined at every position of space r as the density divided by the mean density of the Universe, is of particular interest to this thesis. A key prediction of inflation is that δ should be near-Gaussian, while the cosmological principle ensures its statistical homogeneity and isotropy in the linear regime. These assumptions heavily simplify analytical treatments of its properties (Bardeen et al. 1986), allowing us to link features of the initial density field with properties of non-linear dark matter haloes. For example, the “extended Press-Schechter” formalism posits that sufficiently high density peaks will turn into haloes, where the threshold is defined as being higher than the critical overdensity for collapse predicted by a spherical (Mo et al. 2010, chapter 5) or ellipsoidal model (Sheth et al. 2001; Sheth and Tormen 2002). The mass of the halo is then given by the “broadness” of the peak at the time of crossing the overdensity threshold. Using this formalism, one can predict the statistical abundance of dark matter haloes at a given mass (Press and Schechter 1974; Bond et al. 1991), halo merger trees (Somerville and Kolatt 1999) and spatial clustering (Mo and White 1996). Even though these analytical predictions have long been surpassed in precision by numerical simulations (e.g. Governato et al. 1999; Reed et al. 2003; Tinker et al. 2008), the key lesson which we will use in subsequent chapters is that the initial density field encodes valuable information about the later, non-linear structure formation. Extracting this information content provides physical insights into the growth of dark matter haloes (e.g. Maggiore and Riotto 2010; Paranjape and Sheth 2012; Lucie-Smith et al. 2018), as well as bypassing expensive simulations of structure formation.

In addition to information extraction, one can use the statistical properties of the field to engineer this information into achieving a particular goal. For example, one can specify “constraints”

on the density field and find realisations of the field that satisfy these conditions (Bertschinger 1987; Hoffman and Ribak 1991). These constraints can be observationally motivated, using observed galaxy populations to re-construct the initial conditions of our Local Volume (e.g. Mathis et al. 2002; Libeskind et al. 2010; Heß et al. 2013; Wang et al. 2016; Sorce et al. 2016; Hoffman et al. 2017). They can also be used to craft specific environments, for example engineering rare high-density peaks for the study of galaxy clusters (e.g. Frenk et al. 1999; Gnedin 2003; Donnert et al. 2009) or high-redshift quasars (e.g. Romano-Diaz et al. 2011; Romano-Diaz et al. 2014; Huang et al. 2019).

Constrained realizations approach our idea of a controlled setup within a cosmological environment, but they have a key conceptual drawback. When performing experiments, one usually wishes to vary one and only one aspect, while all remaining features are kept constant. In the case of constrained realisations, constraints are specified before creating the first initial conditions (see Chapter 2 for more details). Every unconstrained part of the density field is therefore modified from one realisation of the experiment to the next, leading to one fixed aspect while every other is changed (e.g. Romano-Diaz et al. 2006). There are also no guarantees that the final results of the experiments are natural within a Λ CDM universe as constraints can be entirely artificial.

A solution to this issue was proposed through a new approach, “genetic modifications” (Roth et al. 2016). This method, which is the core of this thesis, allows us to create alternative versions of a chosen galaxy’s cosmological initial conditions. A new genetically-modified initial condition is derived as close as possible to the starting point, while satisfying a set of user-defined modifications. By re-creating situations which are nearly identical except for one controlled aspect, genetic modifications generate our required controlled setup. Intrinsically cosmological properties of a galaxy such as its halo mass (Roth et al. 2016) and merger history (Pontzen et al. 2017) can be considered tuneable parameters, while the cosmological environment remains fixed. Using such “genetically modified” experiments, we can construct a causal account of the response of the galaxy to variations in merger histories (Chapters 3, 4, 5 and Pontzen et al. 2017). Before we turn to clarifying the theoretical framework and underpinnings of genetic modifications in Chapter 2, we briefly review some of the technical challenges associated with generating initial conditions for cosmological simulations.

Generating numerical initial conditions

Generating cosmological initial conditions is typically inexpensive compared to the cost of the following simulation but remains a challenging task conceptually. We wish to create discretised gas and dark matter elements that sample the field of density perturbations seeding structure growth. This field is typically generated deep into the linear regime ($z \geq 100$), to ensure its Gaussian nature. We therefore need to:

1. specify the cosmological model to obtain the required power spectrum of the Gaussian matter field⁴, either computed analytically (Eisenstein and Hu 1998) or numerically (Seljak and Zaldarriaga 1996; Lewis et al. 2000; Lesgourgues 2011).
2. obtain a random Gaussian sample from this specified power spectrum and for modes within the simulation volume. Since this volume is finite and discretised, largest and smallest scales are not well sampled introducing numerical artefacts and unphysical power on these scales. A detailed review of these technical issues is beyond the scope of this thesis but we will stress that they are still an ongoing topic of research (e.g. Bertschinger 2001; Prunet et al. 2008; Hahn and Abel 2011; Garrison et al. 2016; Falck et al. 2017; Stopyra et al. in prep).
3. evolve these modes forward in time, as we only need numerics when structure growth becomes non-linear. This is typically done using Lagrangian perturbation theory at linear (Zel'dovich 1970) or second order (Jenkins 2010).

A key aspect of this thesis is the ability to generate zoom initial conditions in order to increase the dynamic range of galaxy formation simulations (e.g. Katz and White 1993; Oñorbe et al. 2014). After an initial coarse run, a subvolume of the simulation is selected and tracked back to the initial conditions. Particles within this “zoom region” are sampled with finer resolution and the full volume is evolved again. The computational power is thus concentrated inside the high-resolution zoom region and the large-scale tidal field is self-consistently evolved at the lower resolution.

For zoom setups, one therefore needs to self-consistently incorporate new small-scale power (associated with the fine resolution) while conserving existing large-scale modes (associated with the coarse resolution). A naive approach could be to sample the whole volume at fine resolution and degrade everywhere except in the target region. However, the resolution requested by simulations presented in Chapters 4 and 5 would lead to density fields sampled

⁴We acknowledge that baryon and CDM have different cosmological transfer functions but this difference is negligible on the scales relevant to galaxy formation (e.g. Valkenburg and Villaescusa-Navarro 2017).

at a resolution of $32\,768^3$ before being degraded. This would require manipulating vectors with memory imprints of nearly 2 petabytes which is clearly infeasible in practice. These demanding memory requirements lead to a different approach, based on hierarchies of nested grids (Bertschinger 2001; Hahn and Abel 2011). To sample specific regions at higher resolution, one creates a new smaller grid embedded in the existing parent volume. Because this new grid has smaller physical size but the same number of resolution elements as the parent, it effectively samples smaller scales for twice the memory cost rather than 8 times. The key to this method is then to draw a density field on the embedded smaller grid with large-scale power matching that of the parent. Obtaining a draw under constraints can be achieved using the previously mentioned constrained realisation method (Hoffman and Ribak 1991).

However, the existing two implementations which create multi-scale zoom initial conditions, *GRAFIC* (Bertschinger 2001) and *MUSIC* (Hahn and Abel 2011) were hard to reconcile in practice with our need to genetically modify these zoom initial conditions. Both softwares provide support for nested zoom grids and constrained realisations independently but not combined, therefore lacking the foundational infrastructure to implement genetic modifications. Furthermore, implemented numerical methods rely on “ghosts regions” around zoom regions to ensure the correctness of Fourier transforms, effectively doubling again the memory imprint of each field. Genetic modification introduces a new field *per* modification, which would have led to intractable memory requirements when using multiple modifications as in Chapters 3, 4 and 5. These limitations guided us towards an altogether different approach, implementing new numerical methods which provide the ability to modify zoom initial conditions and improve the general memory efficiency of the procedure. We will describe these improvements in forthcoming work (Stopyra et al. in prep) and omit their description from this thesis.

1.5 Thesis overview

We conclude this introduction by outlining the structure and organization of this thesis. The overall goal of this thesis is to study the cosmological diversity of the galaxy population, and in particular how the interaction between merger histories and internal astrophysics shapes the final observables of galaxies. We make use of a new approach, “genetic modification” (Roth et al. 2016), which allows us to construct controlled studies in the context of cosmological numerical simulations. Genetic modifications generate alternative versions of a galaxy or dark matter halo, each version engineering a carefully designed and controlled change to the galaxy’s merger history. This approach enables us to construct a causal account between intrinsically cosmological properties, such as halo masses and mergers, and galaxy properties such as stellar

masses and star-formation rates. As an example, in the first application of the method, Pontzen et al. (2017) focussed on the response of a galaxy's central black hole to incrementally increasing a merger's mass ratio in its history. They demonstrated that the galaxy's star formation rate dropped as the merger ratio increased, allowing a systematic study of how the combination of mergers with black hole feedback leads to the quenching of massive galaxies. Despite its successes, this first application highlighted technical limitations of the existing approach and we therefore start this thesis by strengthening and extending theoretical and technical aspects of the genetic modification approach (Chapters 2 and 3). We then turn to new applications to dark matter halo (Chapter 3) and galaxy formation (Chapters 4 and 5)

On a technical level, genetic modifications achieve control over merger histories by engineering the initial conditions of a chosen galaxy, for example increasing the height of a density peak in the initial conditions to increase the future, non-linear mass of a halo. With the existing description of linear modifications on the density field (Roth et al. 2016), each structure to modify has to be identified manually, strongly increasing the complexity of the procedure for multiple mergers. In Chapter 2, we therefore clarify the theoretical framework of genetic modifications, before extending it to a new class of non-linear modifications aimed at streamlining the control of multiple mergers. We introduce a memory-efficient algorithm to create quadratically modified initial conditions. We demonstrate its effectiveness in a one-dimensional toy-model before implementing it in the context of 3D cosmological simulations. Armed with this implementation, we show in Chapter 3 that variance modifications are capable of controlling the ratios of multiple mergers in a halo's history as envisioned in Chapter 2.

Genetic modifications allows us to gain key insights on the sensitivity of an object to its merger history. We focus in Chapter 3 on dark matter haloes, and in particular to the link between halo secondary properties (e.g. their density profile and spin) and their mass build-up. Traditionally, these relationships have been established through correlations in large halo samples, e.g. the correlation between formation time and concentrations of haloes (Wechsler et al. 2002; Macciò et al. 2007; Ludlow et al. 2013). Turning such correlations into physical interpretations is challenging (e.g. Zhao et al. 2003; Ludlow et al. 2016) – in Chapter 3, we show how the causal setup provided by the genetic modification approach allows us to generate new insights in how these correlations emerge from physical processes. For example, by varying the merger ratios of two haloes, we demonstrate that the topology of a halo's merger tree has a strong impact on its final concentration. Such results highlight the complementarity of the genetic modification approach with more traditional, correlation-based, methods.

Galaxies are shaped by the interaction between their cosmological histories, such as mergers and accretion, with their internal astrophysics, such as star formation and feedback. Due to their low-mass, dwarf galaxies are particularly sensitive to this interaction. Understanding the diversity of dwarf galaxies is key to putting constraints on the nature of dark matter and on galaxy formation astrophysics, providing a strong motivation to apply genetic modifications in this regime. In Chapter 4 and 5, we focus on the least luminous galaxies in the Universe, ultra-faint dwarf galaxies, as their observed number is rapidly growing following advances in deep imaging (see Simon 2019 for a review) but theoretical predictions are challenging due to the large dynamic range required to resolve such small objects (Munshi et al. 2017; Wheeler et al. 2018; Agertz et al. 2019). We therefore combine the genetic modification approach with high-resolution, zoom galaxy formation simulations to study the sensitivity of ultra-faint properties to their past histories.

In Chapter 4, we create a controlled study varying the early dynamical mass of an ultra-faint host halo while fixing its value at $z = 0$. This scan through histories reveals an extended diversity of stellar masses at fixed dynamical mass and predicts the existence of undiscovered faint, low surface brightness dwarfs within the reach of future facilities such as LSST. In Chapter 5, we use a larger suite of ultra-faint dwarf galaxies with increasing dynamical mass at $z = 0$. We show that our most massive objects are able to re-ignite star formation after reionization. We link this rejuvenation to a growth in dynamical mass at late times, allowing gas to collapse into the halo centre and eventually lead to star formation.

We conclude in Chapter 6 that the genetic modification approach provides a new, promising avenue to study galaxy formation. The ability to perform controlled experiments, in which a galaxy's history is varied in a controlled way, provides new insights and is highly complementary to more traditional methods based on cosmological volumes. We highlight several possible, technical extensions to the method, as well as additional applications to galaxy formation which we plan on tackling in the near future.

Extending the genetic modification framework

As we have seen in Chapter 1, modeling the diversity of the galaxy population from first principles is extremely challenging. The variety of possible assemblies couples non-linearly with the internal astrophysics of a galaxy. Furthermore, the random nature of assembly histories, seeded by inflationary perturbations, hinders our ability to pinpoint their role in shaping the final properties of the galaxy. In this Chapter, we extend the theoretical framework of “genetic modifications” (Roth et al. 2016), a method utilizing the properties of the early-Universe density field to control and modify the formation history of a cosmological halo. The genetic modification approach starts from the initial condition of a single, reference object and generates an alternative, modified initial condition which is then evolved to $z = 0$. The goal is to reproduce as much as possible of the reference initial condition, while making a controlled change to a future non-linear property such as the halo mass or merger history. To perform such modifications, one requires an approximate understanding of how properties of the early-Universe density field map onto non-linear final characteristics. Motivated by analytical structure formation arguments, Roth et al. (2016) demonstrated that linear modifications of the density field, e.g. tuning the height of a density peak, were capable of controlling the final halo mass of an object.

Roth et al. (2016) hence formulated genetic modifications as creating initial conditions under linear constraints, where user-defined constraints encode the future modifications to non-linear halo mass. This formulation mathematically and conceptually echoed that of the “constrained realisations” approach (Hoffman and Ribak 1991), leading to a confusion on the respective aims of these two intrinsically different methods (e.g. Porciani 2016). The first aim of this chapter is therefore to re-formulate the theoretical framework of genetic modifications, carefully distinguishing it from existing approaches in the literature.

By granting control over halo mass, linear modifications also offer the ability to modify the merger history of an object. This ability was showcased in the first application of the method (Pontzen et al. 2017), in which the authors increased and decreased the halo mass of a galaxy’s merging body. The approach used in Pontzen et al. (2017) however requires to identify and track to the initial conditions *each* individual merging body of interest. This feature makes linear

modifications attractive for tackling single major mergers but cumbersome for multiple, small accretion events. The second goal of this chapter is then to extend the genetic modification framework to non-linear modifications, with the goal of streamlining the control over multiple mergers. Finally, this chapter also sets the stage for future applications exploring the sensitivity of dark matter halo (Chapter 3) and dwarf galaxies (Chapter 4 and 5) to their mass assemblies.

Work shown in this chapter is presented in Rey and Pontzen 2018, MNRAS, 474.

2.1 Introduction

Mergers and accretion are thought to play a key role in shaping the observed galaxy population; in the prevailing cosmological paradigm merger histories are in turn seeded from random inflationary perturbations. Numerical studies must make inferences about the galaxy population from a finite sample of such histories. Due to the limited computer time available, this generates a tension between resolution (for resolving the interstellar medium) and volume (for adequately sampling histories).

One attempt to sidestep this problem is to create and study a small number of carefully controlled tests of the relationship between a galaxy’s history and its observable properties. This has long been attempted in idealised, non-cosmological settings (e.g. Hernquist 1993; Di Matteo et al. 2005; Hopkins et al. 2012). More recently, Roth et al. (2016) proposed performing such tests within a fully cosmological environment by constructing a series of closely-related initial conditions with targeted “genetic modifications” (hereafter GMs). The formalism resembles that of constrained realisations (Bardeen et al. 1986; Bertschinger 1987; Hoffman and Ribak 1991, hereafter HR91) which generates realisations of Gaussian random fields satisfying user-defined constraints on initial densities, velocities or potentials (e.g. Bertschinger 2001). Simulations based on constrained realisations have been extensively applied to recreating the local universe using observed galaxy distributions as constraints (for recent examples see Heß et al. 2013; Wang et al. 2016; Sorce et al. 2016; Hoffman et al. 2017).

Despite a resemblance, genetically modified simulations are markedly different from constrained simulations. The process of GM involves creating multiple versions of the initial conditions, each with carefully selected small changes. By re-simulating each scenario it becomes possible to study how the changes affect the non-linear evolution of structure. For example, modifications can be chosen such that they enhance or suppress merger ratios in incremental steps and so vary a galaxy’s history in a systematic and controlled way. The

first application of this technique in a hydrodynamic simulation was made by Pontzen et al. (2017); that work focuses on the response of a galaxy’s central black hole and its ability to quench star formation as the merger history is changed gradually. Unlike studies based on fully idealised merger simulations, the GM-based approach is able to capture the effects of gradual gas accretion from filaments which is essential when probing the balance between star formation and black hole feedback.

On a technical level, Pontzen et al. (2017) used multiple linear modifications to alter the merger history. Such a method requires human effort on two fronts: (i) to identify and track particles forming the merging substructures; and (ii) to tune the modifications and understand their effects on one another. For instance, GMs suppressing a merger tend to increase the mass of other nearby substructures, which complicates interpretation of the final results (see section 2.3 and figure 2 of Pontzen et al. 2017). Bypassing this behaviour would be possible by individually identifying all substructures and demanding the algorithm fix each one. However, the spiralling complexity of the setup makes this option unattractive.

Another possibility, which is the primary aim of the present chapter, is to find a new type of modification which automatically suppresses the merger ratios of *all* large substructures in a target galaxy’s history. Such a modification would smooth the expected history while keeping its final mass and overall environment fixed. These modifications must be applicable to cosmological simulations, so our objective is an algorithm that remains tractable even working with fields on multidimensional grids. To achieve this goal, we start by clarifying the formulation of GMs (Section 2.2). We then expand the framework to quadratic modifications (Section 2.3), allowing control over the variance at different scales to tackle the problem of multiple mergers. We demonstrate the feasibility of our method on a one-dimensional model (Section 2.4); in forthcoming work we will demonstrate the implementation for a full 3D zoom simulation (Stopyra et al. in prep). Results are discussed in Section 2.5 and we conclude in Section 2.6.

2.2 Linear constraints and modified fields

In this Section, we contrast the method of constraints (HR91) against that of linear genetic modifications. The aim is to clarify the status of the latter as a building block for non-linear GMs, which are introduced in Section 2.3.

2.2.1 Constrained ensemble

We start by reviewing the construction of constrained ensembles (see bottom panel of Figure 2.1). In this case, constraints must be known *a priori*, i.e. independently of any specific realisation.

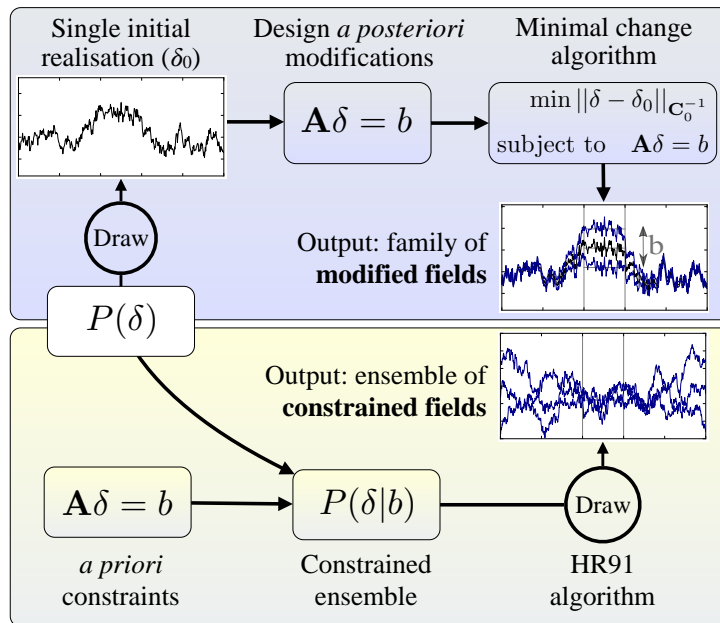


Fig. 2.1.: GMs and constrained ensembles are two techniques to generate targeted initial conditions for numerical simulations. They have markedly different motivations and properties despite sharing similar mathematics. We illustrate the differences by showing the flow of information in the two cases. *Upper Panel:* In the GM case, a single initial realisation (black) is first drawn from the underlying ensemble. Next, modifications are designed to alter chosen properties of this realisation; each modification therefore depends on the specific δ_0 . The modified fields are computed by demanding minimal changes while satisfying the requested modifications. In the illustrated example, we create two modified fields with enhanced and reduced mean values, corresponding to two different values of b inside the target region. *Lower panel:* In the constrained ensemble case, the constraints are independent of any particular realisation and are used to define the ensemble $P(\delta|b)$. This ensemble is efficiently sampled using the HR91 technique. In this example, three fields are drawn and by construction satisfy the same mean value inside the target region.

Constrained ensembles are therefore particularly useful when using observations as external inputs to constrain numerical simulations.

Consider a Gaussian random field sampled at n points to create a vector $\boldsymbol{\delta}$ with covariance matrix $\mathbf{C}_0 = \langle \boldsymbol{\delta} \boldsymbol{\delta}^\dagger \rangle$. The HR91 algorithm allows for an arbitrary number (denoted p) of linear constraints to be placed on $\boldsymbol{\delta}$; these can be expressed as $\mathbf{A} \boldsymbol{\delta} = \mathbf{b}$ where \mathbf{A} is a $p \times n$ matrix and \mathbf{b} is a length- p vector.

We start by constructing the ensemble of all fields $\boldsymbol{\delta}$ satisfying the constraint for a chosen \mathbf{b} , i.e. $P(\boldsymbol{\delta}|\mathbf{b})$. Applying Bayes' theorem, the probability reads

$$P(\boldsymbol{\delta}|\mathbf{b}) = \frac{P(\mathbf{b}|\boldsymbol{\delta}) P(\boldsymbol{\delta})}{P(\mathbf{b})}. \quad (2.1)$$

Using the fact that $P(\boldsymbol{\delta})$ is Gaussian and disregarding normalization, this relation becomes

$$P(\boldsymbol{\delta}|\mathbf{b}) \propto \delta_D(\mathbf{A} \boldsymbol{\delta} - \mathbf{b}) \exp\left(-\frac{1}{2} \boldsymbol{\delta}^\dagger \mathbf{C}_0^{-1} \boldsymbol{\delta}\right), \quad (2.2)$$

where δ_D is the (p -dimensional) Dirac delta function.

This expression suggests a brute force sampling solution: we could draw many trial $\boldsymbol{\delta}$ s from the original ensemble and keep only the ones satisfying the constraints (within some tolerance). This solution is, however, computationally inefficient. Making use of the fact that the Dirac delta function can be represented as the zero-variance limit of a Gaussian, we can instead derive the following results (Bertschinger 1987):

$$\begin{aligned} P(\boldsymbol{\delta}|\mathbf{b}) &\propto e^{-\frac{1}{2}(\boldsymbol{\delta}-\bar{\boldsymbol{\delta}})^\dagger \mathbf{C}^{-1}(\boldsymbol{\delta}-\bar{\boldsymbol{\delta}})}, \text{ with} \\ \bar{\boldsymbol{\delta}}(\mathbf{b}) &= \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} \mathbf{b}, \text{ and} \\ \mathbf{C} &= \mathbf{C}_0 - \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} \mathbf{A} \mathbf{C}_0, \end{aligned} \quad (2.3)$$

where $\bar{\boldsymbol{\delta}}$ and \mathbf{C} are the expectation and the covariance of the Gaussian distribution $P(\boldsymbol{\delta}|\mathbf{b})$. By construction, all fields drawn from this distribution will satisfy the constraints ($\mathbf{A} \boldsymbol{\delta} = \mathbf{b}$).

Hoffman and Ribak (1991) pointed out a convenient shortcut for efficiently sampling from the distribution specified by Equation (2.3). Starting from a draw of the unconstrained ensemble,

δ_0 , we calculate $\mathbf{b}_0 = \mathbf{A} \delta_0$. One can then rewrite δ_0 as the sum of the mean field $\bar{\delta}(\mathbf{b}_0)$ from Equation (2.3) and a residual term δ_{residual} , defined by:

$$\begin{aligned} \delta_{\text{residual}} &\equiv \delta_0 - \bar{\delta}(\mathbf{b}_0) \\ &= \delta_0 - \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} \mathbf{b}_0. \end{aligned} \quad (2.4)$$

From here, a draw from the constrained ensemble δ_1 can be generated by recombining the residuals with the corrected mean $\bar{\delta}(\mathbf{b})$:

$$\delta_1 = \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} \mathbf{b} + \delta_{\text{residual}}. \quad (2.5)$$

To verify this procedure draws samples δ_1 from the constrained distribution, one first writes the mapping from δ_0 to δ_1 in a single step:

$$\delta_1 = \delta_0 - \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} (\mathbf{A} \delta_0 - \mathbf{b}). \quad (2.6)$$

Then, by calculating $\langle \delta_1 \rangle$ and $\langle \delta_1 \delta_1^\dagger \rangle$, it is possible to check that the ensemble has the correct mean and covariance from Equation (2.3). The fact that δ_1 is Gaussian follows from its construction as a linear transformation of δ_0 . The underlying efficiency of this method is that the covariance matrix in Equation (2.3), does not depend on the value of \mathbf{b} , allowing the δ_{residual} term to be the same for both expressions.

In summary, the HR91 algorithm creates a draw from the constrained ensemble in two steps, using the realisation δ_0 as an intermediate construction tool. It provides a computationally efficient way of generating Gaussian constrained fields.

2.2.2 Genetic modifications

We now turn to GMs (see upper panel of Figure 2.1) to contrast their formulation with that of constrained fields. The GM procedure can be summarized as follows:

1. Draw the unmodified realisation δ_0 .
2. Define the modifications by choosing which properties of δ_0 are to be modified. Unlike in the constrained field case, this is accomplished with reference to specific features of the δ_0 realisation (e.g. the location and properties of particular haloes). This reflects how GMs are intended for constructing numerical experiments rather than for recreating observationally motivated scenarios. We focus first on linear modifications, i.e. of the form $\mathbf{A} \delta = \mathbf{b}$.

3. Create the modified field (or multiple modified fields with different values of \mathbf{b}). We require changes between fields to be as small as possible, which relies on the definition of a distance in field space. In the context of Gaussian fields, the only available metric is defined by the χ^2 distance,

$$\chi^2 \equiv \|\boldsymbol{\delta}\|_{C_0^{-1}}^2 = \boldsymbol{\delta}^\dagger C_0^{-1} \boldsymbol{\delta}. \quad (2.7)$$

Consequently, GMs can be formulated as finding the modified field solution of the following optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\delta}} \quad & \|\boldsymbol{\delta} - \boldsymbol{\delta}_0\|_{C_0^{-1}}^2, \\ \text{subject to} \quad & \mathbf{A} \boldsymbol{\delta} = \mathbf{b}. \end{aligned} \quad (2.8)$$

The problem is solved by minimising the Lagrangian

$$\mathcal{L} \equiv (\boldsymbol{\delta} - \boldsymbol{\delta}_0)^\dagger C_0^{-1} (\boldsymbol{\delta} - \boldsymbol{\delta}_0) + \boldsymbol{\lambda}^\dagger (\mathbf{A} \boldsymbol{\delta} - \mathbf{b}), \quad (2.9)$$

where $\boldsymbol{\lambda}$ is a vector of size p containing the Lagrange multipliers for each modification.

By differentiating to find critical points with respect to $\boldsymbol{\delta}$ and $\boldsymbol{\lambda}$, we obtain a system of two vector equations with the solution

$$\boldsymbol{\delta}_1 = \boldsymbol{\delta}_0 - C_0 \mathbf{A}^\dagger (\mathbf{A} C_0 \mathbf{A}^\dagger)^{-1} (\mathbf{A} \boldsymbol{\delta}_0 - \mathbf{b}), \quad (2.10)$$

where $\boldsymbol{\delta}_1$ is the modified field.

Equation (2.10) has regenerated Equation (2.6) using a different motivation and derivation.

To summarise:

- In the case of (2.6), $\boldsymbol{\delta}_0$ is an intermediate construct that is never used in a simulation; it only exists to aid finding $\boldsymbol{\delta}_1$, which is a sample from the distribution (2.3).
- In the case of (2.10), $\boldsymbol{\delta}_0$ and $\boldsymbol{\delta}_1$ are put on equal footing. They are both initial condition fields drawn from the original, underlying ensemble $P(\boldsymbol{\delta})$. The fact that the modifications (choice of \mathbf{A} and \mathbf{b}) depend on $\boldsymbol{\delta}_0$, as emphasised by Porciani (2016), does not impact this interpretation.
- We show in Appendix A.1 that there is no joint expression for GMs and a HR91-like procedure when generalising to non-linear constraints, hence formalising their intrinsic difference.

GMs should therefore be seen as a mapping between fields of the same ensemble. A family of modified fields is generated by choosing multiple values for \mathbf{b} ; the resulting mapping between members of the family is continuous and invertible. These properties are highly valuable for providing controlled tests, allowing for systematic exploration of the effects of formation history on a galaxy.

While the algorithm makes the minimal changes to the field, δ_1 may still not be a particularly likely draw from $P(\delta)$ if the modifications are too extreme. To quantify the level of alteration, the relative likelihood of the two fields is given by $\exp(-\Delta\chi^2/2)$ with

$$\Delta\chi^2 = \delta_1^\dagger \mathbf{C}_0^{-1} \delta_1 - \delta_0^\dagger \mathbf{C}_0^{-1} \delta_0. \quad (2.11)$$

As long as $\Delta\chi^2$ stays small, we can regard the modified and unmodified fields as similarly likely draws from Λ CDM initial conditions.

Turning $\Delta\chi^2$ into a precise quantitative statement about the relative abundance of a particular galactic history is a vast undertaking and remains a topic for future research. It relies on knowing the detailed mapping between the initial conditions and the final, non-linear properties of a galaxy. The complexity of this mapping is however the sheer reason for using cosmological galaxy formation simulations – if we could predict galaxy properties from cosmological initial conditions easily, no simulations would be required. In addition, there are likely multiple possible modifications (i.e. choices of \mathbf{A} and \mathbf{b}) leading to a given effect in the target galaxy history (Porciani 2016); some will carry a smaller $\Delta\chi^2$ cost than others. Finding the minimum-cost route to a given change in the non-linear universe is not the aim of GMs; to perform galaxy formation experiments, we only need to find one choice of modification producing *acceptably* likely initial conditions. In practice (e.g. Chapter 3 and 4), we estimate the likelihood of genetic modifications by comparing the non-linear properties of modified objects with those of a large statistical sample. The control, population sample is itself obtained with numerical simulations, albeit with lower resolution or simplifying assumptions (e.g. dark-matter only).

2.3 Extension to quadratic modifications

The main aim of this chapter is to formulate modifications that control the variance of a field. The variance on scales smaller than the parent halo scale relates to the number of substructures in haloes (Press and Schechter 1974; Bond et al. 1991), and is therefore a proxy for the overall importance of mergers. It is important to distinguish variance modifications of a region from alterations to the power spectrum. The power spectrum defines only the average variance over

the entire box, and over all possible realisations. We propose on the other hand to modify the local variance, targeting *one* region of interest and making minimal changes to the remaining structures. Another way to picture this goal is as follows. In any one stochastic ensemble, two realisations might by chance have enhanced or reduced variance in an area. Our procedure aims to map between such realisations rather than to modify the underlying power spectrum.

Variance is quadratic in the field value and therefore the approach in Section 2.2.2 cannot be applied directly. One natural formulation of the problem is through a new minimisation problem (analogous to the original linear case):

$$\begin{aligned} \min_{\boldsymbol{\delta}} \quad & \|\boldsymbol{\delta} - \boldsymbol{\delta}_0\|_{C_0^{-1}}^2, \\ \text{subject to} \quad & \boldsymbol{\delta}^\dagger \mathbf{Q} \boldsymbol{\delta} = q, \end{aligned} \tag{2.12}$$

where \mathbf{Q} is a $n \times n$ matrix and q is a scalar. We can assume without loss of generality that \mathbf{Q} is Hermitian. For a suitable choice of \mathbf{Q} (see Section 2.4), q specifies the variance of a chosen region.

Following a similar approach to the linear modifications, we introduce the Lagrangian

$$\mathcal{L} = (\boldsymbol{\delta} - \boldsymbol{\delta}_0)^\dagger C_0^{-1} (\boldsymbol{\delta} - \boldsymbol{\delta}_0) + \mu (\boldsymbol{\delta}^\dagger \mathbf{Q} \boldsymbol{\delta} - q), \tag{2.13}$$

where μ is a scalar Lagrange multiplier associated with the quadratic modification. Searching for critical points, we obtain two equations relating the modified field $\boldsymbol{\delta}_1$ and the multiplier:

$$\boldsymbol{\delta}_1 = (\mathbf{I} + \mu C_0 \mathbf{Q})^{-1} \boldsymbol{\delta}_0, \text{ and} \tag{2.14}$$

$$\boldsymbol{\delta}_0^\dagger (\mathbf{I} + \mu C_0 \mathbf{Q})^{-1} \mathbf{Q} (\mathbf{I} + \mu C_0 \mathbf{Q})^{-1} \boldsymbol{\delta}_0 = q. \tag{2.15}$$

Equation (2.14) and (2.15) provide a closed system for μ and $\boldsymbol{\delta}_1$ given a target q . Unlike the linear case, the system cannot be solved analytically. A possibility would be to solve Equation (2.15) numerically for μ but direct matrix inversions are prohibited due to their computational cost. One would therefore need to perform approximate matrix inversion at each step of a root-finding scheme for μ , making the worst-case complexity of such method infeasible.

There are moreover deeper reasons why such procedures cannot be straightforwardly adapted to GMs. In the linear case discussed above, we defined GMs as a continuous and invertible mapping. Both of these properties are lost when looking at Equations (2.14) and (2.15). First, it

is not clear that Equation (2.15) has a real solution for μ . Consequently a real-valued δ_1 may not exist¹ for any chosen value of q .

Second, the relationship between δ_0 and δ_1 is asymmetric: if a new field δ' is constructed by taking q back to its original value q_0 , we will have

$$\delta' = \left(\mathbf{I} + (\mu + \mu')\mathbf{C}_0\mathbf{Q} + \mu\mu'(\mathbf{C}_0\mathbf{Q})^2 \right)^{-1} \delta_0, \quad (2.16)$$

for suitable choices of μ and μ' . To obtain a solution μ' allowing recovery of the initial field ($\delta_0 = \delta'$), it must hold that $\mathbf{C}_0\mathbf{Q} \propto (\mathbf{C}_0\mathbf{Q})^2$. This will not generally be the case for our applications, and so we conclude that in general $\delta' \neq \delta_0$. Such asymmetry would be problematic for GM; the sense of a unique ‘family’ of fields is lost.

The combination of computational intractability and loss of key properties for GMs lead us to focus on an alternate method. We describe next a Newton-like method which efficiently approximates a solution to the optimization problem, Equation (2.12), while reinstating the desired properties of the GM mapping.

2.3.1 Linearised solution

In this section, we restate the quadratic problem in a way that has a guaranteed solution and that generates a single family as a function of q . The trick is to make only infinitesimally small changes to the value of q , building up finite changes by following a path through field space that is *locally* minimal. This leads to an iterative procedure for quadratic genetic modifications, which we will demonstrate is both unique and computationally tractable.

One infinitesimal step We start by defining the displacement ϵ from the unmodified field $\delta = \delta_0 + \epsilon$; for sufficiently small changes we may then neglect $\mathcal{O}(\epsilon^2)$ terms. We will discuss next how to practically decompose a macroscopic change into a series of such minor modifications.

At first order, the updated variance (or other quadratic property) is given by

$$\delta^\dagger \mathbf{Q} \delta = \delta_0^\dagger \mathbf{Q} \delta_0 + 2 \delta_0^\dagger \mathbf{Q} \epsilon + \mathcal{O}(\epsilon^2), \quad (2.17)$$

¹We note in passing that, since variance is a positive quantity, \mathbf{Q} is a positive semi-definite matrix. By definition, \mathbf{C}_0 is positive definite. These conditions ensure that δ_1 is unique if it exists – but they do not guarantee existence.

where we have assumed δ is real and made use of the previously stated Hermitian assumption, $Q^\dagger = Q$. Having linearised the modification, we can now find an analytic solution for the displacement and the multiplier μ :

$$\epsilon = \mu C_0 Q \delta_0, \text{ with} \quad (2.18)$$

$$\mu = \frac{1}{2} \frac{q - \delta_0^\dagger Q \delta_0}{\delta_0^\dagger Q C_0 Q \delta_0}. \quad (2.19)$$

Equation (2.19) does not involve matrix inversions and can therefore be efficiently evaluated, even in a 3D cosmological simulation context.

Building finite changes by successive infinitesimal updates We now want to construct a macroscopic change in the field by iterating the infinitesimal steps of Equation (2.18). Performing a finite number of steps N , the modified field reads:

$$\delta_1 = \prod_{j=0}^N (I + \mu_j C_0 Q) \delta_0, \quad (2.20)$$

where μ_j is the Lagrange multiplier at step j . The value of each μ_j depends on how the fixed interval is divided, i.e. implicitly on N . In the limit of increasing number of steps, each individual μ_j becomes infinitesimally small and the final solution is

$$\begin{aligned} \delta_1 &= \lim_{\substack{N \rightarrow \infty \\ \mu_j \rightarrow 0}} \prod_{j=0}^N (I + \mu_j C_0 Q) \delta_0 \\ &= \prod_{j=0}^{\infty} \exp(\mu_j C_0 Q) \equiv \exp(\alpha C_0 Q) \delta_0, \end{aligned} \quad (2.21)$$

where $\alpha = \sum_{j=0}^{\infty} \mu_j$ is the overall displacement and is finite. The right-hand side of Equation (2.21) defines the matrix exponential operator, which is guaranteed to exist and is invertible.

The matrix exponential is a useful formal expression to show that there is a unique result, but does not help computationally since the required value of α to reach the objective $\delta_1^\dagger Q \delta_1 = q$ is unknown. In practice, we use the finite approximation, Equation (2.20). The μ_j at each step are chosen by targeting N intermediate modifications linearly spaced between the starting value $q_0 \equiv \delta_0^\dagger Q \delta_0$ and the target q . At each step, μ_j is calculated using Equation (2.19); ϵ_j is deduced with Equation (2.18); and the field is updated, $\delta \rightarrow \delta + \epsilon_j$.

Step choice for a practical algorithm When calculating Equation (2.20) as an approximation to Equation (2.21), the accuracy will increase with the number of steps N . One must choose a minimal N (for computational efficiency) while ensuring that linearly approximating the modification at each step is sufficiently accurate.

We first perform the calculation with a fixed number of steps N_{initial} . This gives rise to an initial estimate for the modified field that we denote $\delta_{1,\text{initial}}$. The error on the resulting modification can be characterised by the magnitude of η_{initial} , where

$$\eta_{\text{initial}} \equiv \delta_{1,\text{initial}}^\dagger \mathbf{Q} \delta_{1,\text{initial}} - q. \quad (2.22)$$

Because second-order terms are neglected in the modification, the error term η_{initial} should scale inverse-quadratically with the number of steps N_{initial} . We verified this behaviour numerically for a variety of fields and modifications. If η_{initial} is smaller than a desired precision, η_{target} , we retain the initial estimate as our final output field. Otherwise, the calculation must be repeated; the required number of steps to achieve the target precision is inferred from the quadratic scaling as

$$N = N_{\text{initial}} \sqrt{\frac{\eta_{\text{initial}}}{\eta_{\text{target}}}}, \quad (2.23)$$

Note that N_{initial} should be kept small to avoid unnecessary iterations; $N_{\text{initial}} = 10$ has been chosen for our test scenarios below.

The final algorithm has a worst-case complexity of $\mathcal{O}(\eta_{\text{target}}^{-1/2} n^3)$, where n is the number of elements in the field δ . The n^3 arises from matrix multiplications required to compute each step; in practice the matrices will be sparse either in Fourier space (for the covariance matrix) or in real space (for the variance \mathbf{Q} matrix). Therefore, one can speed up the matrix multiplications by transforming back and forth from real to Fourier space, improving the complexity to $\mathcal{O}(\eta_{\text{target}}^{-1/2} n \log n)$.

The final procedure shares numerous similarities with Newton methods, used in large-scale optimization (see Nocedal and Wright 2006 for a comprehensive review). It retains quadratic information in the objective and linear information in the modification at each step and has a quadratic rate of convergence to the solution.

2.3.2 Joint quadratic and linear modifications

The algorithm above can be generalised to the case where we have both a quadratic modification and p linear modifications of the form $\mathbf{A} \delta = \mathbf{b}$. We first apply the linear modifications using Equation (2.8), then turn to the iterative quadratic modifications. However if Equation (2.20)

is applied directly, the linear objective will no longer be satisfied; in other words we need to enforce $\mathbf{A} \epsilon = \mathbf{0}$ at each step. Constructing and solving the appropriate minimisation, expression (2.18) is replaced by

$$\epsilon = -\mu \mathbf{C}_0 \mathbf{Q} \delta + \mu \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} \mathbf{A} \mathbf{C}_0 \mathbf{Q} \delta, \quad (2.24)$$

where

$$\mu = \frac{1}{2} \frac{q - \delta^\dagger \mathbf{Q} \delta}{\delta^\dagger \mathbf{Q} \mathbf{C}_0 \mathbf{A}^\dagger (\mathbf{A} \mathbf{C}_0 \mathbf{A}^\dagger)^{-1} \mathbf{A} \mathbf{C}_0 \mathbf{Q} \delta - \delta^\dagger \mathbf{Q} \mathbf{C}_0 \mathbf{Q} \delta}. \quad (2.25)$$

These results can be iterated to achieve the final modified field, in exactly the same way as for the pure-quadratic modification.

Despite the complexity of these expressions, the evaluation will remain $\mathcal{O}(\eta_{\text{target}}^{-1/2} n \log n)$ for reasons discussed previously. To help interpret the method, there is a clear geometric meaning for each term, which we present in Appendix A.2.

2.4 Demonstration

In this Section we demonstrate our algorithm in a n -pixel, one-dimensional setting as a proof of concept and as a reference for future implementation on cosmological simulations. We choose an example red power spectrum, as typically encountered on the scales from which galaxies collapse. Specifically, we adopt $P(k) = P_0 (k_0 + k)^{-2}$, where P_0 is an arbitrary normalisation and $k_0 = 2\pi/n$, an offset that prevents divergence of $P(k)$ at $k = 0$.

2.4.1 Defining an example modification

The framework developed in Section 2.3 can alter any property that is quadratic in the field by suitable choice of \mathbf{Q} . We now specialise to the case that \mathbf{Q} corresponds to the variance of a length- R region of the field. We start by defining the windowing operator \mathbf{W} as a rectangular matrix picking out the desired R entries from the n pixels in δ . To calculate the variance of the region, one then calculates $\delta^\dagger \mathbf{Q}_{\sigma^2} \delta$ where \mathbf{Q}_{σ^2} can be written

$$\mathbf{Q}_{\sigma^2} = \frac{1}{R^2} \mathbf{W}^\dagger (R \mathbf{I} - \mathbf{1} \otimes \mathbf{1}) \mathbf{W}. \quad (2.26)$$

Here, \mathbf{I} is the $R \times R$ identity matrix and $\mathbf{1}$ is a length- R vector of ones. Expression (2.26) is readily verified by constructing $\delta^\dagger \mathbf{Q}_{\sigma^2} \delta$ and seeing that it does boil down to the variance of the chosen region.

We wish to consider the field variance only on scales smaller than the region size (corresponding to substructures with mass lower than that of the parent halo). To achieve this, \mathbf{Q}_{σ^2} can be high-pass filtered; we use a standard Gaussian high-pass filter \mathbf{F} where in Fourier space the elements of $\tilde{\mathbf{F}}$ are given by

$$\tilde{F}_{lm} = \delta_{lm} \left(1 - \exp \left[-\frac{1}{2} \left(\frac{k_l}{k_f} \right)^2 \right] \right). \quad (2.27)$$

Here, $k_l = 2\pi l/n$ is the wavenumber of the l -th Fourier series element and k_f , the filtering scale, is defined in our case by $k_f = 2\pi/R$. The most appropriate choice of filtering scales and shapes in the context of cosmological simulations will be discussed in a forthcoming work.

In real space the matrix \mathbf{F} is defined by $\mathbf{F} = \mathbf{U}^\dagger \tilde{\mathbf{F}} \mathbf{U}$ where \mathbf{U} is the unitary Fourier transform matrix. Finally, to localise the target modification fully, we can re-window the matrix after smoothing. The operator $\mathbf{W}^\dagger \mathbf{W}$ achieves this by setting pixels outside the target window to zero. With this set of choices, the final quadratic objective is set by

$$\begin{aligned} \mathbf{Q} &\equiv \mathbf{W}^\dagger \mathbf{W} \mathbf{F}^\dagger \mathbf{Q}_{\sigma^2} \mathbf{F} \mathbf{W}^\dagger \mathbf{W} \\ &= \frac{1}{R^2} \mathbf{W}^\dagger \mathbf{W} \mathbf{F}^\dagger \mathbf{W}^\dagger (\mathbf{R}\mathbf{I} - \mathbf{1} \otimes \mathbf{1}) \mathbf{W} \mathbf{F} \mathbf{W}^\dagger \mathbf{W}. \end{aligned} \quad (2.28)$$

In practice, we never calculate the matrix \mathbf{Q} explicitly but rather implement a routine to efficiently calculate $\mathbf{Q} \delta$ for any field δ , which is then used by the algorithm described in Section 2.3. The ability to bypass storing or manipulating \mathbf{Q} is essential to permit the computation to operate on a 3D cosmological simulation.

2.4.2 Results

Figure 2.2 shows examples of modified fields obtained with our algorithm. We alter the variance of a region of width $R = 100$ pixels enclosed by vertical lines, showing two quadratic modifications with the variance reduced by a factor 3 (light grey) and a factor 10 (dark blue). In both cases, the mean of the field is held fixed at the unmodified value (horizontal line). In the setting of a cosmological simulation, we expect to be able to fix the parent halo mass (through the mean value) while modifying the smoothness of accretion (through the variance).

We verified that these fields achieve the linear modification $\mathbf{A} \delta_1 - \mathbf{b}$ to within numerical accuracy and the quadratic modification $\delta_1^\dagger \mathbf{Q} \delta_1 - q$ to $\eta_{\text{target}} = 10^{-6}$ accuracy. The heights of small-scale peaks inside the enclosed region are successfully reduced and brought closer to the mean value for the modified fields. Visually, it can be seen that the changes to the field




Fig. 2.2.: Example genetic modification of a Gaussian random field with power spectrum $P(k) \propto (k_0 + k)^{-2}$. The unmodified and modified fields are shown respectively by dashed and solid curves. The region targeted for alteration is enclosed by vertical lines. We use simultaneous linear and quadratic modifications to conserve the mean value of this region (horizontal line) while reducing the small-scale variance by a factor 3 (grey) and 10 (blue). In the context of galaxy formation, this would maintain the mass of a galaxy and make its formation history smoother, while making minimal alterations to the large-scale environment.

are minimal, maintaining as much as possible of the structure of the unmodified field in the modified versions. This underlines how the analytic minimisation, Equation (2.12), and its refinement to a linearised procedure (Section 2.3.1) agrees with the intuitive sense of making minimal changes. The different versions of the field form a continuous family as illustrated by the smooth deformation when reducing the variance by different factors.

Despite the modification objective Q being strictly confined to the target region, modifications can be seen to “leak” outside (beyond the vertical lines). This effect, which is also seen in linear GMs, is an intentional aspect of the minimisation construction – any sharp discontinuities in the field value or its gradients would give rise to a power spectrum inconsistent with the ensemble. In this specific example, the leakage appears more significant to the left than to the right of the target region. In general, the algorithm is spatially symmetric but its effect in any given case is not.

2.5 Discussion

2.5.1 The advantage of quadratic over linear modifications

Pontzen et al. (2017) showed that using multiple linear modifications was sufficient to change the merger ratios in the history of a galaxy; substructures can be diminished or enhanced by manually modifying individual peak heights.

Nonetheless, we expect the new quadratic approach to bring considerable benefits when making such manipulations; the advantages are illustrated in Figure 2.3. The top panel shows a field representing the density in initial conditions expected for a halo. The field has a broad overdensity enclosed by vertical lines and two narrower peaks labelled (1) and (2). According

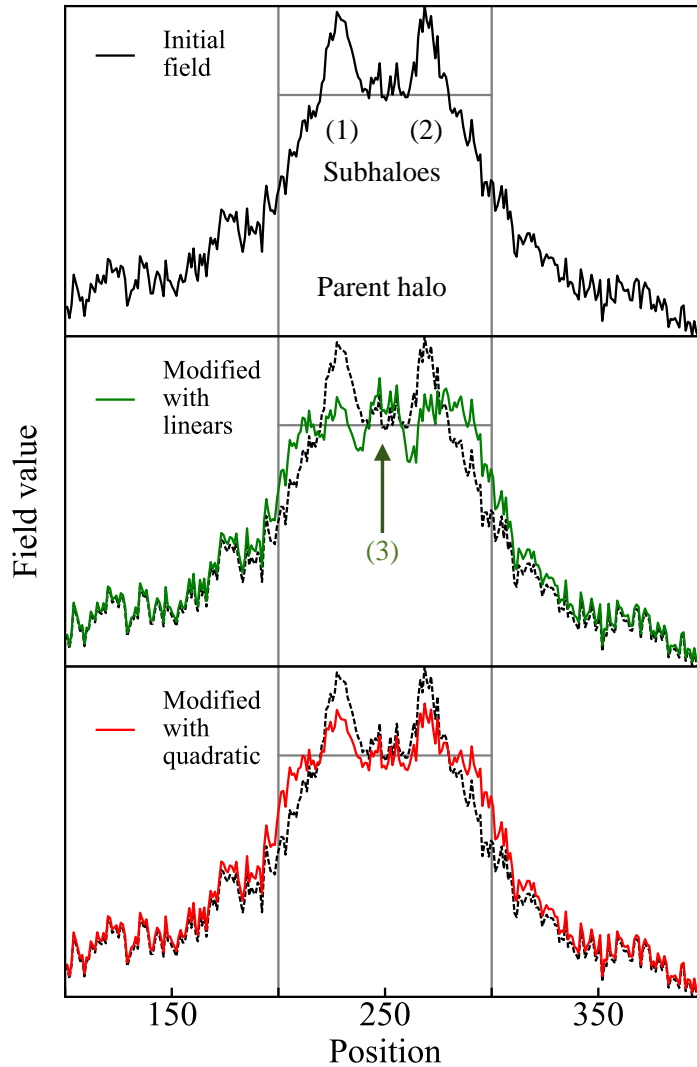


Fig. 2.3.: Comparison of pure linear against combined linear-quadratic GMs. *Top panel:* The unmodified field contains three distinct features: a broad overdensity that would generate a parent halo (enclosed by vertical lines) as well as two localised substructures labelled (1) and (2) that would lead to a merger during the formation history. The objective is to reduce the peak heights of these substructures while conserving the mean height of the parent (horizontal line). *Middle panel:* a GM field with linear modifications designed to bring peaks (1) and (2) to the mean value of the broad region. This approach has successfully smoothed the peak structure. However, as explained in the text, it suffers from the creation of an artificial substructure (3). *Bottom panel:* The same objective has been achieved through a variance modification. This quadratic modification does not require identifying individual subhaloes and by construction prevents unhelpful compensations such as (3).

to the excursion-set formalism (Bond et al. 1991), (1) and (2) will collapse to form two separate haloes that later merge. This, together with smooth accretion, will form the final halo.

Suppose we wish to generate a smoother accretion history by reducing the heights of peaks (1) and (2) while maintaining the large-scale overdensity. In the original approach, we use linear GMs to set the mean values of the peaks to the mean value of the broad overdensity (horizontal line). The middle panel of Figure 2.3 presents the resulting field. However, a number of problems arise when performing the alteration using this approach.

1. We had to identify (1) and (2) as the most interesting substructures and define specific modifications for each. In the context of N-body simulations, this requires manually identifying which particles of the initial conditions constitute each individual subhalo.
2. More importantly, spatially neighbouring modifications interact and create new substructures (peak labelled (3) in our example). One solution to prevent the appearance of new substructures could be to add a new linear objective forcing problematic regions such as (3) to remain unmodified. Identifying and mitigating side effects in this way adds a layer of complexity to the linear GM procedure. Depending on the specific problem and the number of modifications at play, time spent at this tuning phase can rise steeply.

On the other hand, a single quadratic modification can avoid these problems by defining a variance target across the region. The third panel of Figure 2.3 shows the same field with variance reduced by a factor 10 (using the method from Section 2.4). The two local peaks are successfully reduced in amplitude while conserving the remaining small-scale structure of the parent halo. By construction, the variance modification naturally avoids compensation problems inherent to linear GMs. For this reason, quadratic GMs provide a cleaner, streamlined way to control merger histories.

2.5.2 Multiple quadratic modifications

The formalism discussed so far applies a single quadratic modification to a field (possibly in combination with linear objectives). Simultaneously applying *multiple* quadratic modifications would allow one to act concurrently on two separate haloes, or to further fine-tune the merger history of a single object. For instance, decreasing the variance on intermediate scales while increasing on small scales should increase the frequency of minor mergers.

To study this generalisation, we introduce $i = 1, \dots, P$ quadratic modifications, each with matrix \mathbf{Q}_i . For an infinitesimal update, the change in the field ϵ is then given by

$$\epsilon = \sum_i \mu_i \mathbf{C}_0 \mathbf{Q}_i \delta, \quad (2.29)$$

$$\text{with } \delta^\dagger \mathbf{Q}_i \delta + 2\delta^\dagger \mathbf{Q}_i \epsilon = q_i \quad \text{for all } i, \quad (2.30)$$

where μ_i are the Lagrange multipliers associated with each modification. Equation (2.30) defines a system of P equations to be solved. The resulting value of a specific μ_i depends on the whole set of q_i and \mathbf{Q}_i , i.e. modifications are interdependent.

In the same way as Section 2.3.1, the update (2.29) can be iterated to create finite changes. Performing N steps, the modified field reads

$$\delta_1 = \prod_{j=0}^N \left(\mathbf{I} + \sum_{i=0}^P \mu_{ij} \mathbf{C}_0 \mathbf{Q}_i \right) \delta_0, \quad (2.31)$$

where μ_{ij} is the multiplier μ_i at step j . However in the limit that the number of steps $N \rightarrow \infty$, convergence to the matrix exponential,

$$\delta_1 = \exp \left(\sum_i \alpha_i \mathbf{C}_0 \mathbf{Q}_i \right) \delta_0, \quad (2.32)$$

is only guaranteed if either the \mathbf{Q}_i commute with respect to \mathbf{C}_0 (i.e. $\mathbf{Q}_i \mathbf{C}_0 \mathbf{Q}_j = \mathbf{Q}_j \mathbf{C}_0 \mathbf{Q}_i$) or each μ_{ij} is directly proportional to α_i . Because α_i s are not known in advance, the latter option is hard to arrange; the previously noted interdependence of the μ_i s on all q_i and \mathbf{Q}_i exacerbates the difficulty.

With our current algorithms, convergence to the matrix exponential is therefore only assured when the \mathbf{Q}_i matrices commute. The easiest way to arrange for the commutation is to use orthogonal modifications, i.e.

$$\mathbf{Q}_i \mathbf{C}_0 \mathbf{Q}_j \approx 0. \quad (2.33)$$

Physically, this requirement can be achieved by imposing modifications that are spatially separated by a sufficient number of correlation lengths or address distinct Fourier modes. This condition even allows one to apply the formalism of Section 2.3.1 to each modification one-by-one and still converge to the correct overall matrix exponential of Equation (2.32). We leave the case of non-orthogonal multiple quadratic modifications to further work.

2.6 Conclusions

We have presented an efficient algorithm to modify the variance in a particular region of a Gaussian random field realisation, with the aim of manipulating initial conditions for cosmological simulations. The modification produces a field that is as close as possible to the original realisation. In this way it provides a route for controlled tests of galaxy formation where multiple versions of the same galaxy are simulated within a fixed cosmological environment, but with altered accretion history.

We argued that quadratic controls, as developed here, offer a useful complement to the existing linear technique (Roth et al. 2016). In particular, variance on different filtering scales relates to dark matter halo substructure and merging history (Press and Schechter 1974; Bond et al. 1991). The new algorithm can construct GM fields with simultaneous control on the mean value and filtered variance of a region (Figure 2.2). This provides a route to altering merger history and accretion over the lifetime of a given halo in a way that is more streamlined than modifying individual substructures (see Figure 2.3).

In both linear and quadratic GM, the algorithm searches for fields which are nearby in the sense of the χ^2 distance measure. In the quadratic case, this definition is further refined: for large shifts in the control parameter q (which represents the variance in our test cases), the path through field space is defined by following a series of small shifts. Each of these individually minimize the χ^2 distance traveled. We demonstrated a formal convergence property for this series and argued that the approach is desirable for (a) returning a continuously-deforming field δ as a function of the changing target variance q ; (b) being reversible, so that returning the variance to its initial value also returns the field to its initial state; (c) being numerically tractable even for 3D zoom simulations.

In the process, we clarified the mathematical formulation of GM, carefully distinguishing it from the constrained ensemble of HR91 (see Figure 2.1 for an overview). The status of fields constructed in the two approaches is distinct – unlike constrained realisations, GMs should be seen as a mapping between two fields from the same ensemble. In the case of quadratic objectives such as variance, even the cosmetic similarities between constraints and modifications are lost (Appendix A.1).

The next step is an implementation of the new algorithm in a full N -body initial conditions generator, including on varying-resolution grids appropriate to zoom simulations. We will present in the next chapter a successful implementation and evaluate the effectiveness of

quadratic GMs (alongside the existing linear technique) for controlling the mass accretion histories of dark matter haloes.

Sensitivity of dark matter haloes to their accretion histories

In the previous chapter, we extended the “genetic modification” framework to allow for a finer control of a galaxy’s mass accretion history, based on the local variance of the density field. In this chapter, we implement and apply this algorithm in the context of cosmological dark matter only simulations. This chapter serves both as a proof of concept demonstrating the effectiveness of variance modifications before an application to hydrodynamical simulations (Chapters 4 and 5), as well as providing new insights in the relationship between a halo’s history and its final properties.

In order to apply variance modifications in the context of cosmological simulations, the algorithm and toy model described in Chapter 2 needs to be implemented within a 3D cosmological initial condition generator. A major difficulty in this process is to combine multiscale Gaussian random fields (i.e. zoom initial conditions) with the genetic modification algorithm, ensuring that modifications are propagated self-consistently across both the parent volume and the zoom region (see Section 1.4 for more technical insights). This combined ability is currently unavailable in the literature, whether for linear or quadratic constraints, which has led us to develop a purpose-built software to generate zoomed, modified cosmological initial conditions. We omit a description of our new approach in this thesis, as the theoretical underpinnings and numerical methods will be publicly released to the community in Stopyra, Pontzen, Peiris, Roth and Rey, in prep.

Work shown in this chapter is presented in Rey et al. 2019b, MNRAS, 485.

3.1 Introduction

In a Λ CDM universe, galaxies form and evolve embedded inside dark matter haloes. The mass of the halo is believed to be the primary driver of most galaxy properties; for example, empirical models of galaxy formation have often relied on a parametrized mapping between dark matter halo mass and galaxy stellar mass (see Wechsler and Tinker 2018 for a review). Dark matter halo mass is also the main parameter of halo clustering models used to recover cosmological information from galaxy surveys (Cooray and Sheth 2002).

However, dark matter haloes grow over time through hierarchical merging: smaller building-blocks merge together to assemble larger haloes. The same mass at a given time could have been assembled in many different ways: through accretion of numerous small bodies or through a smaller number of more significant events. This diversity of possible mass accretion histories (hereafter MAHs) at fixed halo mass is thought to generate scatter on both the galaxy-halo relationship (Moster et al. 2013; Rodriguez-Puebla et al. 2016) and in halo clustering bias (Gao et al. 2005; Wechsler et al. 2006; Wetzel et al. 2007). The scatter can be further characterised by investigating the role of a halo’s “secondary” properties (e.g. density concentration, spin, or age). The evolution of these secondary properties is shaped by the response of dark matter haloes to external factors, such as mergers and large-scale environment.

For a given halo, mergers and large-scale environment are seeded stochastically from inflationary perturbations. This poses a challenge in studying any physical processes related to secondary properties; the most common solution to date has been to simulate large numbers of haloes to sample possible MAHs and cosmological environment at a given mass scale (Bullock et al. 2001b; Wechsler et al. 2002; Macciò et al. 2007; Ludlow et al. 2013; Klypin et al. 2016, though see also Zhao et al. 2003). Emergent correlations in the halo population have been characterised with such methods, e.g. the relationship between halo concentration and mass (Ludlow et al. 2014; Diemer and Kravtsov 2015; Klypin et al. 2016). But statistical sampling intrinsically makes it hard to construct causal models because, by definition, every degree of freedom changes from one halo to the next. For example, halo concentrations have been found to be correlated with halo formation time (Wechsler et al. 2002), halo environment (Avila-Reese et al. 2005; Macciò et al. 2007; Maulbetsch et al. 2007; Lee et al. 2017) or halo spin (Macciò et al. 2007) but the interpretation of these correlations is still under debate. These uncertainties propagate into empirical and semi-analytical models of galaxy formation that rely on a physical account of the link between dark matter halo properties and galactic properties.

Recently, genetic modification (hereafter GM, Roth et al. 2016) was introduced as a method to study the response of a halo or galaxy to a controlled change in its merger history. GMs create different versions of the same halo, each with carefully specified modifications, while maintaining the same cosmological large-scale structure. One can then compare a range of scenarios for the formation of a particular halo or galaxy in a simulation, keeping all degrees of freedom fixed except those specifically targeted. The first application to galaxy formation was made by Pontzen et al. (2017) who studied the response of a galaxy’s star formation history to increased or decreased merger activity.

In practice, modifications are made in the initial conditions, and the simulation is then performed again. Control over MAHs is achieved by modifying the height and broadness of density peaks in the linear universe, motivated by analytical structure formation theories (Press and Schechter 1974; Bardeen et al. 1986; Bond et al. 1991). This was achieved in Pontzen et al. (2017) by manually tracking each merging substructure to the initial conditions and modifying each region to obtain the required merger history. Tackling multiple mergers with such a method risks spiralling complexity. To simplify this procedure, we presented in Chapter 2 an extension to the existing framework targeting the local variance of the density field. Variance encodes the height of multiple peaks and troughs compared to a mean value, so should allow us to obtain direct control over the importance of multiple mergers.

The aim of this chapter is two-fold: to demonstrate that variance modifications provide the expected control over the overall smoothness of merger histories, and to show how GMs can develop a causal account of the role of merger histories in shaping halo secondary properties. We will show that a study of a small number of GM objects complements existing large population studies. This chapter is organized as follows. In Section 3.2, we outline the procedure to generate modified initial conditions and evolve them to $z = 0$. In Section 3.3, we construct specific families of modified haloes and demonstrates that variance modifications directly control MAHs at fixed mass in the way anticipated in Chapter 2. In Section 3.4, we show that the continuous range of GM scenarios causally links the details of MAHs to population-level variations in halo secondary properties. We summarize our results and conclude in Section 3.5.

3.2 Numerical setup

We start by reviewing the necessary ingredients to generate a family of genetically modified haloes. The theoretical underpinnings of GMs are described in Roth et al. (2016) and Chapter 2; we focus here on the practical aspects of the procedure. We describe how we define modifications in Section 3.2.1 and in Section 3.2.2 how we evolve initial conditions to $z = 0$ in a cosmological context.

3.2.1 Initial conditions

In this section, we describe the general method of selecting haloes for re-simulation with the zoom technique (Katz and White 1993) and generating a family of modified initial conditions using the quadratic algorithm of Chapter 2.

To set up the zoom regions, we start from a simulation with uniform resolution and select a halo, which will become the reference halo of the modified family. (The specific two haloes

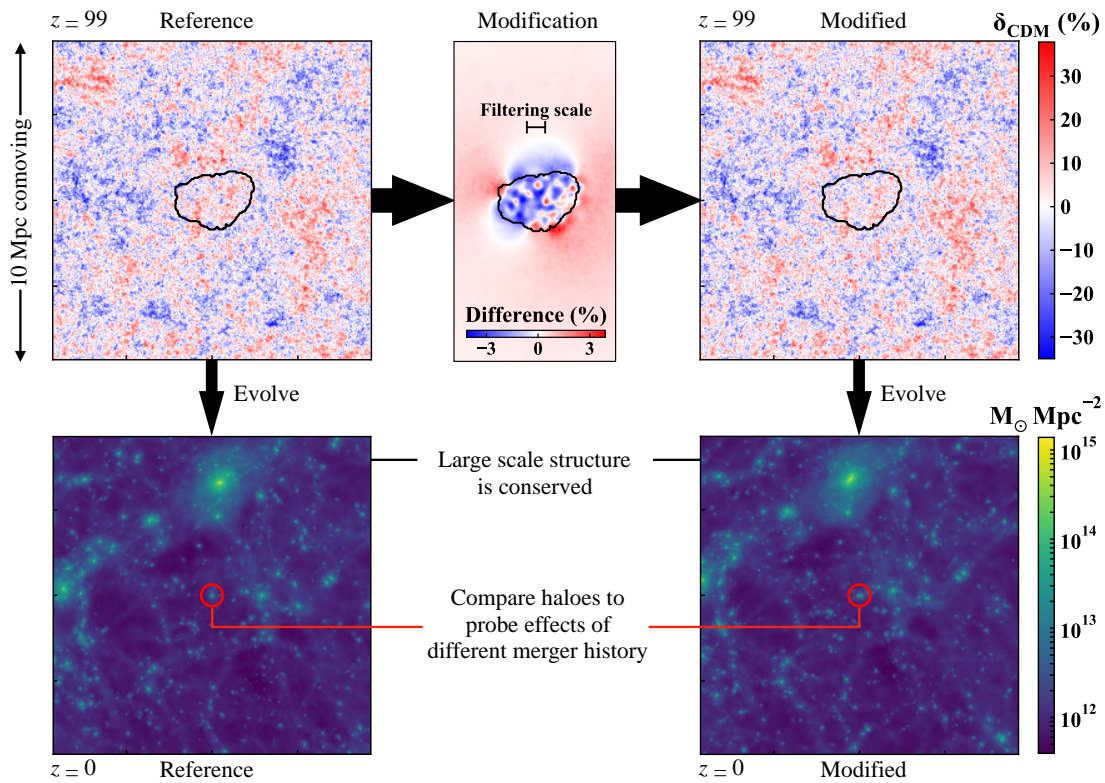


Fig. 3.1.: Workflow for generating a genetically modified halo from a reference halo. The top left panel shows a slice through the density field at $z = 99$, i.e. the initial conditions for the reference halo. We evolve the density field to $z = 0$ (bottom left) to obtain the properties of this halo, highlighted by a red circle of 3 virial radii. We then construct a genetically modified initial condition (top middle and right panels) with variance modifications (see Section 3.2.1). We modify the smoothness of merger histories by redistributing the density structure inside the targeted region (black contour enclosing all particles tracked back from the reference halo at $z = 0$). The large-scale structure (bottom right) matches that of the reference run (bottom left) but the merger histories of the two haloes are different (e.g. Figure 3.2) as well as their structural properties (e.g. Figure 3.4).

chosen for this chapter are described in Section 3.3.1.) We track back the centre of this halo to the initial condition and, for this chapter, open a spherical region of radius comoving $3 h^{-1} \text{Mpc}$ in which the mass resolution is refined. The large extent of the zoom region is a simple way to ensure that the halo of interest is not contaminated by heavy, i.e. low-resolution, particles in our present context where computing time is not a limiting factor. We verified that all haloes are contaminated to less than 0.1 per cent at $z = 0$, including after modifications which can lead to different particles falling into a halo. The initial conditions are defined from a flat ΛCDM cosmology (Planck Collaboration et al. 2016a), with $h = 0.6727$, $\Omega_m = 0.3139$, $\Omega_b = 0.04916$, $\Omega_\Lambda = 0.686095$, $\sigma_8 = 0.8440$, $n_s = 0.9645$ and evolved using the Zel'dovich approximation (Zel'dovich 1970) to $z = 99$. All simulations have a box size of $50 h^{-1} \text{Mpc} \approx 74 \text{Mpc}$ and distances are stated in comoving units unless otherwise stated.

From these zoom initial conditions, we now generate modified initial conditions. Our goal is to modify the merger history of a given halo while keeping its final mass constant. Modifying the mean density in the Lagrangian region of a halo has a direct impact on the final halo mass (see Roth et al. 2016), but we expect variance modifications predominantly to redistribute mass inside the region, hence acting on the merger history.

Several ingredients are required to define these modifications:

1. the spatial windowing. We target the Lagrangian region of the unmodified halo by tracing back to $z = 99$ all particles found in the halo at $z = 0$. We define the boundary of a halo as the spherical radius, r_{200c} , at which the mass density is equal to 200 times the critical density of the Universe.
2. a mass scale for merging substructures. Variance modifications are designed to act on the peaks and troughs of the density field *at a given scale*. To motivate the choice of the spatial filtering scale, we use the approximation that in the initial conditions, mass M is linked to a scale R via $M \approx 4\pi \bar{\rho} R^3/3$, where $\bar{\rho}$ is the average density of the Universe. In practice, we choose the spatial scale corresponding to the mass at infall of the targeted merging substructures. Note that it is not necessary to identify specific substructures before generating modifications, unlike in the linear GM case.
3. a control for the final halo mass. We adjust the mean overdensity inside the targeted region to fix the halo mass of the modified halo at $z = 0$.

Choosing the values of these generic parameters depends both on the halo in hand and the requested modifications to its MAH. We will describe in Section 3.3.1 the details of the halo families used in this chapter.

Simulation name	Mean changes (%)	Variance changes (%)
Volume-L50N512	N/A	N/A
Halo 740 family	{=, +10, =, =, =, -5}	{=, +10, -5, -8, -10, -20}
Halo 839 family	{=, =, =, =, =}	{=, -10, +10, +20, +30}

Tab. 3.1.: Description of the modifications made in this chapter. From a uniform resolution volume, we select two target haloes (740 and 839) as reference for halo families. We create modified initial conditions from these reference haloes: changes in the mean density of the Lagrangian region control the halo mass at $z = 0$, while changes in the variance control the smoothness of MAHs (see Section 3.3.1) around a targeted mass scale. Changes are quoted with respect to the value of the reference halo.

3.2.2 N-body evolution

Once the reference and modified initial conditions have been generated, we need to evolve them to $z = 0$. All simulations presented in this chapter are dark matter only; we use `RAMSES` (Teyssier 2002) which follows N-body evolution with a particle-mesh method and cloud-in-cell interpolation. The mesh on which the forces are calculated is adaptively refined over the course of the simulation. We allow mesh refinement when 8 dark matter particles are inside the same cell, up to a maximum spatial resolution of 1.1 kpc. The mass of dark matter particles is $1.5 \times 10^7 M_\odot$ and $1.2 \times 10^8 M_\odot$ inside and outside the zoom regions, respectively. We save 50 snapshots, equally spaced in scale factor between $z = 99$ and $z = 0$.

We identify dark matter haloes using the HOP halo finder as described in Eisenstein and Hut (1998), and discard haloes with fewer than 100 particles. To calculate halo properties and merger trees, we make use of the `PYNBODY` and `TANGOS` software packages (Pontzen et al. 2013; Pontzen and Tremmel 2018). A key aspect of this chapter is to compute the build-up of mass in a halo. `TANGOS` uses the unique ID carried by each dark matter particle to match a halo with its successor in time, based on the fraction of common particles between two structures. Repeating this procedure for each snapshot constructs halo merger trees which are stored in a database. In our `GENETIC` initial conditions generator and modifier, IDs are generated self-consistently between modified and unmodified simulations, allowing us also to match haloes across different simulations.

We use the shrinking-sphere algorithm (Power et al. 2003) to determine halo centres. At each timestep, we define the virial mass of a halo M_{200c} as the mass enclosed within the spherical radius, r_{200c} at the output redshift.

3.3 Controlling the smoothness of merger histories

We now describe the specific modification setup used in this chapter. Figure 3.1 presents a graphical summary of the procedure described in Section 3.2, with an example reference and variance-modified halo. The top row focusses on the initial conditions of the two haloes, while the bottom row shows the resulting evolution at $z = 0$. We show slices of the overdensity field for the reference (top left) and modified (top right) initial conditions, as well as a slice through the difference between these fields (top middle). The two initial conditions are similar except in the target region (black contour). The filtering scale defines the scale at which density changes are targeted inside the target region.

Leakage outside the targeted region can be observed in the difference field. This is a feature of the GM algorithm: density perturbations are correlated on all scales in a Λ CDM universe. Since the algorithm is maintaining the correct power spectrum, it requires changes to have some level of non-local impact. In this case, the leakage is sufficiently minimal to visually recover near-identical large-scale environments at $z = 0$ (bottom panel), and we therefore leave it unconstrained. We quantitatively discuss the impact of this choice in Section 3.4.3.

3.3.1 Creating modified haloes

We now turn to the construction of two specific halo families that we will use in the remainder of this chapter. Since the greatest number of applications will eventually be in hydrodynamical galaxy formation problems, we select two haloes at the peak of star formation efficiency (Behroozi et al. 2013), i.e. $M_{200c} \approx 10^{12} M_{\odot}$ at $z = 0$. These two haloes, halo 740 and halo 839, are selected to have different MAHs and environment, so that we can explore the effect of variance modifications in different regimes. The modifications performed on these two haloes are summarized in Table 3.1.

Halo 839 has a quiet merger history with a series of small events prior to $z \sim 2$ and steady accretion of mass thereafter (see Figure 3.2). We generate a family of modified haloes targeting the merging structures around $z \sim 2$, enhancing their significance by following the procedure described in Section 3.2.1. The merging structures have a mass of $\sim 1.5 \times 10^{10} M_{\odot}$ at first infall, motivating a filtering scale of $0.30 h^{-1} \text{Mpc}$. To enhance merger significance, we increase incrementally the variance by 10, 20 and 30 per cent relative to the value in the reference run. We also decrease the variance by 10 per cent to explore whether the merger history can be made even smoother. We emphasize that each member of the family sits in a large-scale

environment minimally modified from the reference halo and that conserving the mean density in the Lagrangian region ensures that halo masses match at $z = 0$ to within 5 per cent.

Halo 740 has a more complex merger history, dominated by two equal-mass mergers in less than 1 Gyr around $z \sim 1$ (see Figure 3.3). We wish to make the overall history smoother by decreasing the local variance. The merging structures around $z \sim 1$ have masses $\sim 2 \times 10^{11} M_{\odot}$, motivating a filtering scale of $0.70 h^{-1} \text{Mpc}$. We again adopt an incremental approach, decreasing the variance by 5, 10 and 20 per cent relative to the value in the reference run. We also increase the variance by 10 per cent to explore if the merger history can be made even rougher. The first initial runs were performed conserving the mean overdensity in the Lagrangian region. While motivated by analytical structure formation, the mean overdensity is a property of the linear density field, and hence not a perfect predictor of the strongly non-linear halo mass. Following these initial modifications, we observed variations in the final halo masses of halo 740's family members of up to 10 per cent. These larger changes compared to halo 839 are due to modifications targeting larger structures, closer in scale to the overall final halo mass. For complete clarity we modified the mean density as stated in Table 3.1 and re-ran the simulations, finding agreement improved to match within 5 per cent.

We emphasize that the strength of the GM method lies in its incremental approach. By making continuous changes to the initial conditions and studying the consequent non-linear response of the halo, we can pinpoint the tipping points of halo and galaxy formation, where small changes in merger histories have large consequences on observed properties. We study next the detailed impact of our modifications on the MAHs and merger tree structure of both haloes (Section 3.3.2), as well as how these changes are reflected in dark matter halo properties (Section 3.4).

3.3.2 Results

Figure 3.2 presents the results of variance modifications on halo 839. The top panel shows the MAHs over cosmic time of the reference halo (thick blue) and three family members with variance decreased by 10 per cent and increased by 10 and 30 per cent (orange solid, red dotted and purple dashed respectively). Variance modifications targeted several merger events around $z \sim 2$, highlighted in time by the grey band. We show in the middle panel the merger trees of the reference halo and the modified halo with increased variance by 30 per cent in this time window. The size of branches scales logarithmically with their mass; the darker bottom branch in each case is the major progenitor. Merging events are matched between simulations (see Section 3.2.2) and highlighted by the linked black boxes. We quantify merger mass ratios by

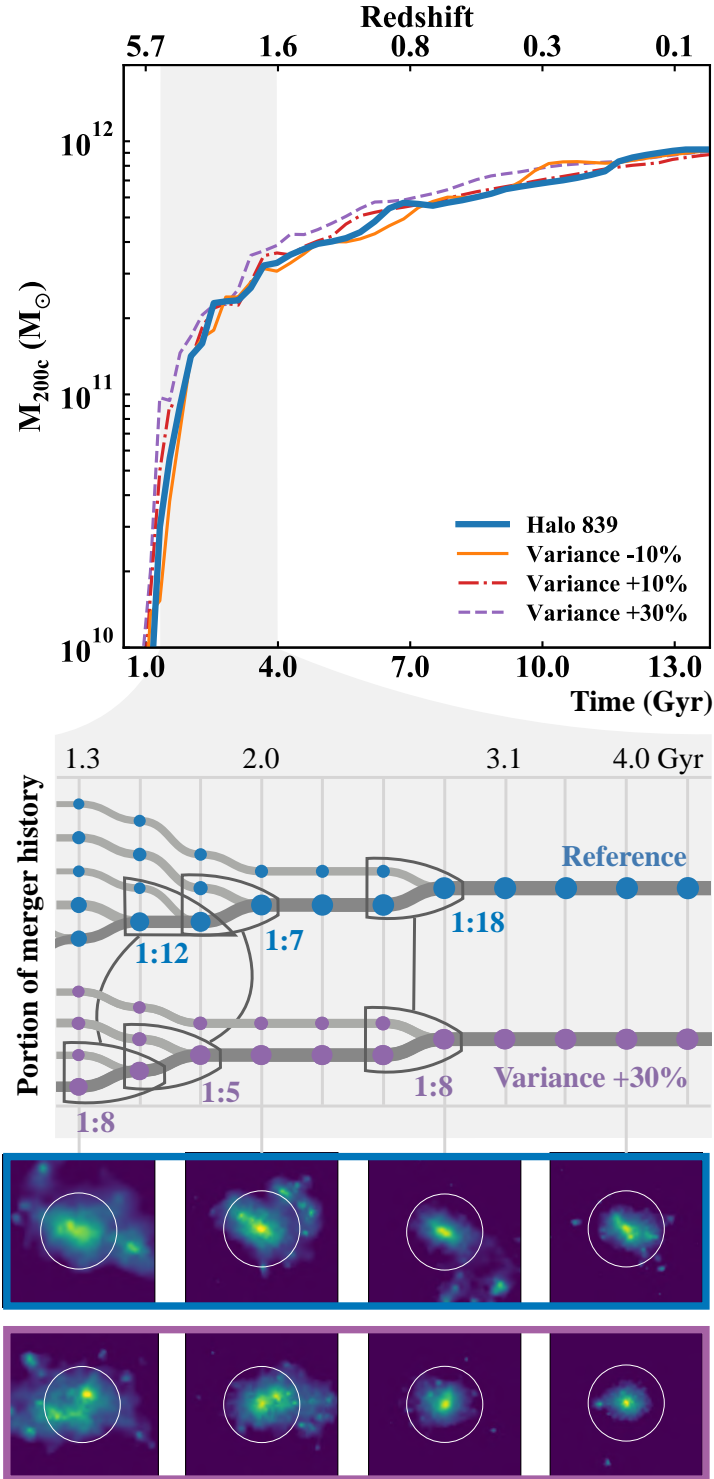


Fig. 3.2.: The growth of virial mass over cosmic time (top panel) for the reference (blue thick curve) halo 839 and its modified counterparts. Halo 839 has a quiet merger history with its most significant activity being three minor mergers around $z \sim 2$. The impact of variance modifications is hence most visible on the merger trees (middle panel). We show the reference halo (blue) and the modified halo with increased variance by 30 per cent (purple), with highlighted mergers matched between the trees. Increasing the variance successfully increases the mass ratios of all three mergers. We visually illustrate how genetic modifications modify mass accretion histories by showing integrated dark matter maps (all with identical color range) at four different times along the two merger trees (bottom panel).

the ratio between the number of particles inside haloes as found by the halo finder. We have verified that changing to merger ratios as defined with M_{200c} or the mass within one scale radius (Hopkins et al. 2010) does not impact our conclusions.

By inspecting the merger trees, we can confirm the ability of a single variance modification to control multiple mergers. The reference halo 839 has three minor events with merger ratios 1:12, 1:7 and 1:18. Increasing the variance by 30 per cent increases the merger ratios of all three mergers, respectively to 1:8, 1:5 and 1:8. Thus, we have successfully made the targeted mergers of halo 839 more significant, confirming the viability of the quadratic modification approach to controlling MAHs.

Figure 3.3 presents the results of variance modifications to halo 740. As before, we show the MAHs of the reference halo and three selected family members (top panel), as well as the merger trees for the reference halo and the modified halo with variance decreased by 10 per cent (middle panels). The reference halo 740 has three major events in the window of interest: two roughly equal-mass mergers (ratios 1:1) that we explicitly targeted with our choice of filtering scale and a less significant 1:6 event. Reducing the variance by 10 per cent reduces the merger ratios of the two targeted mergers (now 1:2 and 1:3). However, the required compensation in mass accretion results in increasing the merger ratio of the smaller event to 1:3. Since we fix the halo mass at $z = 0$, modifications inevitably redistribute the mass between substructures to obtain convergence of MAHs at late time. The minimal nature of GM naturally ensures that MAHs converge before the targeted $z \sim 1$ mergers.

Variance modifications for halo 740 create a reconfiguration of the merger tree topology. Comparing the merger trees for the reference halo and the -10 per cent modified halo (middle panel of Figure 3.3), we find that the same substructures are incorporated into the main progenitor (darker bottom branch of each tree) following two different patterns:

- (A) In the reference merger tree (blue panel) the top two branches merge together. The merger remnant from the top two branches is later incorporated into the main progenitor. Over the course of the time window, the main progenitor experiences *two* mergers.
- (B) In the -10 per cent variance merger tree (orange panel), the same top two branches remain independent. They are incorporated turn by turn into the main progenitor, which now experiences *three* mergers over the course of the time window.

Our incremental approach, with continuous changes to the initial conditions, allows us to explore the merger tree configuration as a function of local variance. For halo 740, the reference halo, as well as $+10$ and -5 per cent modified haloes follow scenario (A), while -10 and -20

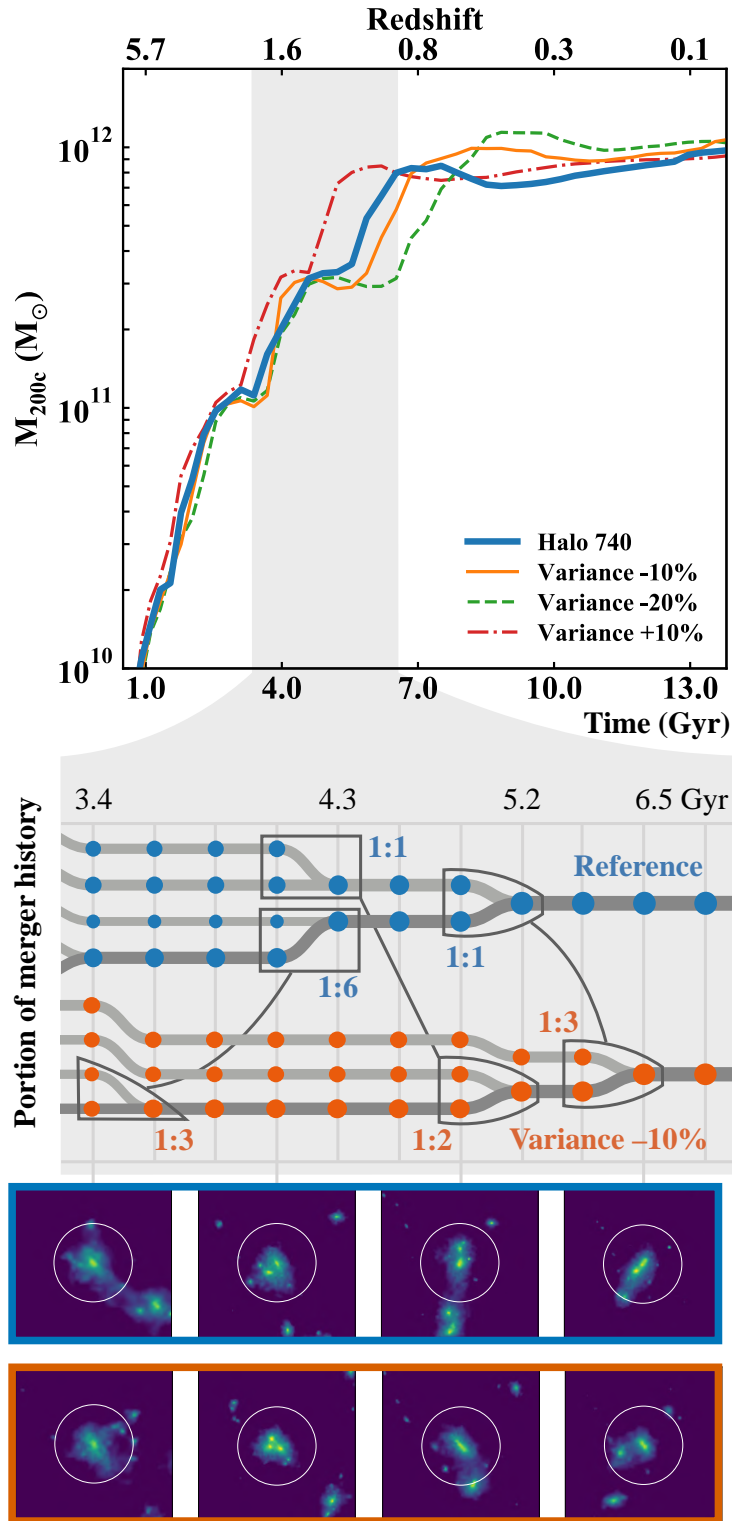


Fig. 3.3.: As Figure 3.2 but for halo 740, which has a more complex merger history that is dominated by equal-mass mergers around $z = 1$. Decreasing the variance by 10 per cent successfully smooths the impact of these two mergers, as measured by their respective merger ratios (orange tree). Because we fix the halo mass at $z = 0$, there must be a compensating increase in mass from other sources: in fact a 1:6 merger in the reference run (blue tree) increases significance to 1:3. Thus, one minor and two major mergers in the reference simulation are mapped onto three medium mergers in the modified case. This also generates a change in the topology of the trees: unlike in the reference, all three structures directly merge into the major progenitor in the modified case.

per cent follow scenario (B). Note that the reconfiguration tipping point is a property of this particular halo; our variance modifications on halo 839 never modify the topology of the merger tree, for example. We will see in Section 3.4 that merger tree topology has consequences for the evolution of dark matter halo properties that cannot be discerned from population studies alone.

In addition to merger ratios and tree structure, variance modifications also impact the timing of mergers. A systematic outcome can be observed for both families: reducing merger ratios pushes mergers to later times; conversely, increasing ratios pushes mergers earlier. This effect was also observed in Pontzen et al. (2017) using a different modification setup. In fact, it is a generic outcome of any GM procedure, since the linear density field is given by the divergence of the velocity field in Λ CDM initial conditions. Reducing merger ratios is achieved by smoothing density gradients, in turn reducing the relative initial velocities of the two substructures. These two structures will then take longer to coalesce, effectively pushing the merger to later time. If fixing the timing of a merger is paramount, one can construct and impose a new velocity modification to conserve the peculiar velocity structure inside the Lagrangian region, at the cost of increasing leakage effects outside the region. We leave the study of simulations with such additional modifications to future work.

In summary, variance modifications create different versions of the reference halo with:

1. a minimally modified large-scale structure (see Figure 3.1), which we will quantify in Section 3.4.3;
2. a predictable but non-trivial effect on MAHs: increased variance increases the mass ratios of mergers and shifts them to slightly earlier times, and may reconfigure the tree topology when critical thresholds are exceeded (see Figure 3.2 and 3.3);
3. a minimal impact on the remaining MAH, which converge before and after the area targeted by variance modifications (e.g. Figure 3.3).

The combination of these features shows that we have achieved the objective of simulating essentially the same halo, up to the selected modifications. Variance modifications create a “dial” through which the MAH of a given halo can be seen as a tuneable parameter. This opens the door for a wide variety of applications, both for galaxy formation studies and for dark matter halo physics on which we will focus next.

3.4 Response of haloes to changes in their merger history

We demonstrated in Section 3.3.2 our ability to generate different versions of the same halo with systematically altered accretion history. We now study the response of $z = 0$ dark matter halo properties to these varying MAHs. Since our families of modified haloes have been generated with fixed final halo mass, we focus on variations in halo secondary properties such as halo concentration, formation time or spin. The statistical correlations between these properties have been extensively studied (e.g. Bullock et al. 2001b; Wechsler et al. 2002; Macciò et al. 2007; Ludlow et al. 2014; Klypin et al. 2016). However, the physical drivers of the emergent correlations and their scatter remains uncertain due to the extremely large number of degrees of freedom in the initial conditions. Genetic modifications alter a small subset of the degrees of freedom, allowing us to gain new physical insight into secondary properties.

We compare our families of modified haloes to a large statistical sample extracted from the Volume-L50N512 simulation, our highest uniform resolution simulation (see Table 3.1) with $\sim 38\,000$ haloes at $z = 0$. Comparing to the overall halo population will allow us to quantify the relative impact of modifications with respect to population-level variations.

3.4.1 Calculating halo secondary properties

The properties that we measure are as follows.

1. *Concentration*: We calculate halo concentrations, c_{NFW} , using the velocity profile method (described in Prada et al. 2012; Klypin et al. 2016). This method is based on computing the ratio of peak circular velocity to circular velocity at r_{200c} (see Section 3.2.2) as a measure of halo concentration. This concentration is a profile-independent quantity so is not impacted by the goodness of fit to a specific analytical form (e.g. Navarro et al. 1997, hereafter NFW, or Einasto 1965). To obtain the velocity ratio, we compute the enclosed mass profile in bins evenly spaced in log radius between 0.7 kpc and r_{200c} . From this, we calculate the circular velocity profile from which we can obtain the desired velocity ratio. We remap this concentration measure into the well-known NFW concentration using equation 20 of Klypin et al. (2016) to facilitate comparison with past studies. We have checked that our results are unchanged if using the velocity ratio as a direct measure of concentration rather than remapping to the NFW definition.
2. *Spin*: We calculate halo spin, λ_{B01} , following the definition of Bullock et al. (2001a), equation 5, inside a sphere of radius r_{200c} .

3. *Formation redshift*: We define the formation redshift, z_F , as the redshift at which the main progenitor has accreted the mass enclosed within the NFW scale radius at $z = 0$ (Ludlow et al. 2014; Correa et al. 2015).
4. *Environment*: We define halo environment, $1 + \delta_{10}$, by calculating the density inside a sphere of radius 10 Mpc centred on the halo. We divide by the mean matter density of the Universe to obtain a dimensionless quantity following Macciò et al. (2007).

Our goal is to study how halo secondary properties vary following variance modifications. There are, however, potential sources of purely numerical variations, which we investigate briefly before turning to the main results. The first is the impact of numerical resolution (see e.g. Power et al. 2003 for halo density profiles). Using the same volume simulation but with degraded mass resolution of $9.6 \times 10^8 M_\odot$, we compare the median and 68 per cent confidence intervals of all secondary properties in the mass range $5 \times 10^{11} M_\odot < M_{200c} < 5 \times 10^{12} M_\odot$ at these two resolutions. We find they are different by less than 4 per cent for halo concentrations and less than 1 per cent for other secondary properties and therefore conclude that our results will not be sensitive to resolution. Second, large variations in halo secondary properties can also be caused by propagation of numerical noise in an intrinsically chaotic system, i.e. the “butterfly effect” (Keller et al. 2019; Genel et al. 2019). To exclude this possibility, we ran all simulations in the family of halo 740 three times. Our simulation code, RAMSES, does not conserve the order of arithmetic operations, hence ensuring that roundoff errors are seeded differently between each re-run. We find that rerunning yields variations in halo secondary properties to up to 2 per cent. We will see next that both variations remain small compared to the overall effect of genetic modifications and we hence conclude that our results are robust to numerical issues.

3.4.2 The evolution of concentration and formation time

Figure 3.4 presents halo properties for the two families of modified haloes (purple squares and red diamonds for halo 740 and halo 839 respectively) compared to the halo statistical sample at $z = 0$ (blue hexagon bins). Each modified halo has been engineered to have systematically varied MAH at fixed halo mass. Arrows show the direction of increasing variance for each family.

The left-hand panel of Figure 3.4 shows the concentration-mass relation at $z = 0$. Our halo population recovers the well-known trend that the median NFW concentration (black points) decreases with increased halo mass, although with considerable scatter (Bullock et al. 2001b; Ludlow et al. 2014; Klypin et al. 2016). Physically, this scatter is thought to be generated in part by the variety of MAHs at given halo mass (Wechsler et al. 2002; Ludlow et al. 2013). We cleanly

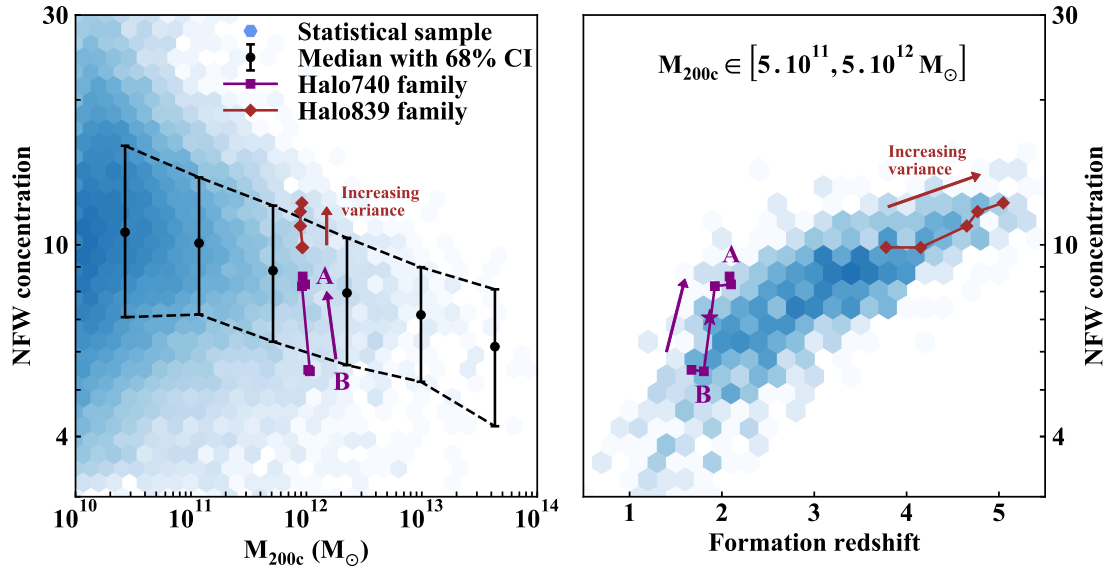


Fig. 3.4.: The concentration-mass relation at $z = 0$ (left) for the two modified families (halo 740 and halo 839 in purple squares and red diamonds respectively). Increasing the variance (arrows) in the initial conditions systematically yields more concentrated haloes. We compare these variations with correlations in the overall halo population (blue hexagon bins with median and 68 per cent confidence intervals in six bins of halo mass). Variations in the mass accretion history of two haloes generate changes in concentration that are comparable to the entire population-level scatter. These variations are consistent with the correlation between halo concentration and formation redshift (right), as both families move along the direction of the correlation. The jump in halo concentration between (A) and (B) is tied to the reconfiguration of the merger tree shown in Figure 3.3. To verify this, we constructed an additional simulation (purple star), intermediate between the ends of the jump. We found that this version undergoes a three-way merger (i.e. it sits precisely on the transition between the two merger topologies) and as a result its concentration is also intermediate.

demonstrate this causal link here: systematically varying MAH yields a systematic change in concentration at fixed halo mass.

The link between MAHs and concentration has been extensively studied through empirical correlations (Wechsler et al. 2002; Zhao et al. 2003; Gao et al. 2005; Gao et al. 2008) and analytical models (Ludlow et al. 2014; Correa et al. 2015; Ludlow et al. 2016). The strongest predictions are obtained by relating the early, fast mass accretion phase (summarised by z_F) to the final halo concentration. At fixed mass, a halo assembling its central mass earlier (when the Universe was denser) will be more centrally concentrated. The population in our volume reproduces this trend (Figure 3.4, right-hand panel) in the mass interval $5 \times 10^{11} M_\odot < M_{200c} < 5 \times 10^{12} M_\odot$. We see from both families that increasing variance also incrementally increases the formation redshift of modified haloes. Thus the population-level correlation is reproduced as a causal connection in the family-level studies.

The evolution of both families is gradual along the direction of the correlation: increasing the variance of a given halo makes it form earlier, in turn making its concentration higher. This is especially visible in the case of halo 839 in which our modifications targeted early

mergers, hence smoothly modifying its formation redshift. By way of contrast, the halo 740 family shows less variation in formation redshift z_F because modifications targeted late-time mergers, leaving the early history unchanged. Nonetheless, we observe a significant variation in halo concentration, highlighted by a “jump” of $\Delta c_{\text{NFW}} = 2.6$ in halo concentration from $c_{\text{NFW}} = 5.5$ to 8.1, untied to any significant variation in formation redshift. This discontinuity is significant at the population level, in the sense that $\Delta c_{\text{NFW}} = 2.2 \sigma$, where σ is defined as half the 68 per cent confidence interval of the population. It occurs at the point that the merger tree is reconfigured, as described in Section 3.3.2. All members of the family with a merger tree structure following (A) (i.e. with one combined merger into the major progenitor) are found on the high end of the discontinuity in concentration, while all members following (B) (with two mergers into the major progenitor) are on the low end. No such reconfigurations are ever generated in halo 839’s family and similarly, we do not observe a discontinuous response in halo concentration.

Our incremental approach with variance modifications therefore allows us to tie the origin of this concentration jump to the merger tree reconfiguration. We can explain *a posteriori* why the two different scenarios generate vastly different concentrations. Scenario (A) is dominated by an equal-mass merger on a mostly radial orbit. The mass being brought by this penetrates deeper in the potential well of the main progenitor, leading to a higher concentration. By contrast, the same mass in scenario (B) is incorporated smoothly through two merger events with smaller mass ratios. The in-falling mass is more evenly distributed through the final halo, leading to an overall lower concentration.

One of the key feature of the GM approach is its ability to refine around tipping points by generating additional intermediate scenarios. To further test our explanation that the merger tree reconfiguration is the source of the concentration jump, we generate a new modification from halo 740 with variance decreased by 8 per cent (in between the two ends of the jump which have variance decreased by respectively 5 per cent and 10 per cent). In this intermediate case, the previously described mergers combine into a three-way event, i.e. the transition scenario in merger tree topology between (A) and (B). The purple star in Figure 3.4 shows the resulting formation time and concentration of this halo which, as expected, bisects the two ends of the jump. This confirms the causal connection between the merger scenario and the concentration.

We conclude that the memory of the different merger configurations between scenario (A) and (B) is retained to $z = 0$ and generates a discontinuity in the resulting halo concentrations. We emphasize that this discontinuity is not stochastic, but rather the result of mapping continu-

ous, incremental changes in the initial conditions to a discrete merger tree topology (i.e. mergers happen in this order or do not). The discrete transition is due to major mergers happening later than the halo’s formation time and hence cannot be captured by models mapping the early mass assembly onto halo concentrations. New summary statistics beyond formation time would be needed to encapsulate knowledge about merger tree topology, for example by counting the number of mergers weighted by their mass ratios. We leave an exploration of such new summary measures as future work and now turn our attention to other secondary halo properties and their relationships with halo concentration.

3.4.3 Halo environment

Another potential source of scatter in the concentration-mass relationship is the diversity of halo local environments at fixed mass (Avila-Reese et al. 2005; Macciò et al. 2007; Maulbetsch et al. 2007; Lee et al. 2017), especially at low halo mass. We now investigate the response of halo environments in families of modified haloes. Visually the large scale structure is unaltered (Figure 3.1), but we should consider a more quantitative measure.

The left panel of Figure 3.5 shows such a measure of environment, the local overdensity at 10 Mpc against halo concentration. The mass bin from which the halo population is extracted and the colour coding are the same as Figure 3.4. We recover no significant correlation between this measure of environment and halo concentration at this mass in the halo population, consistent with past studies (e.g. Macciò et al. 2007; Lee et al. 2017).

We observe a significant and systematic variation with increasing variance in both families, compared to the population scatter, with increasing variance systematically pushes haloes to underdense environment. More specifically, halo 740 and halo 839’s entire families go from $1 + \delta_{10} = 0.42$ to 0.31 and $1 + \delta_{10} = 2.31$ to 1.35 respectively, with increasing variance. This evolution can be compared to the population, ranging 0.66σ and 0.81σ respectively where we have defined σ on the log of the population since cosmological densities are log-normally distributed (e.g. Coles and Jones 1991).

Given the visual similarity of the large scale structures in a family of simulations (Figure 3.1), these differences should be explained. The GM algorithm by construction minimises the differences between reference and modified regions, but also seeks to maintain the correlation function of the field. In the case of variance modifications, we create changes on small scales (i.e. internal to the halo) which then leak to larger scales due to these intrinsic correlations (see the top centre panel of Figure 3.1 for a visual example of leakage). Decreasing the variance on small scales means decreasing the density contrasts in this region, hence creating an *overdensity*

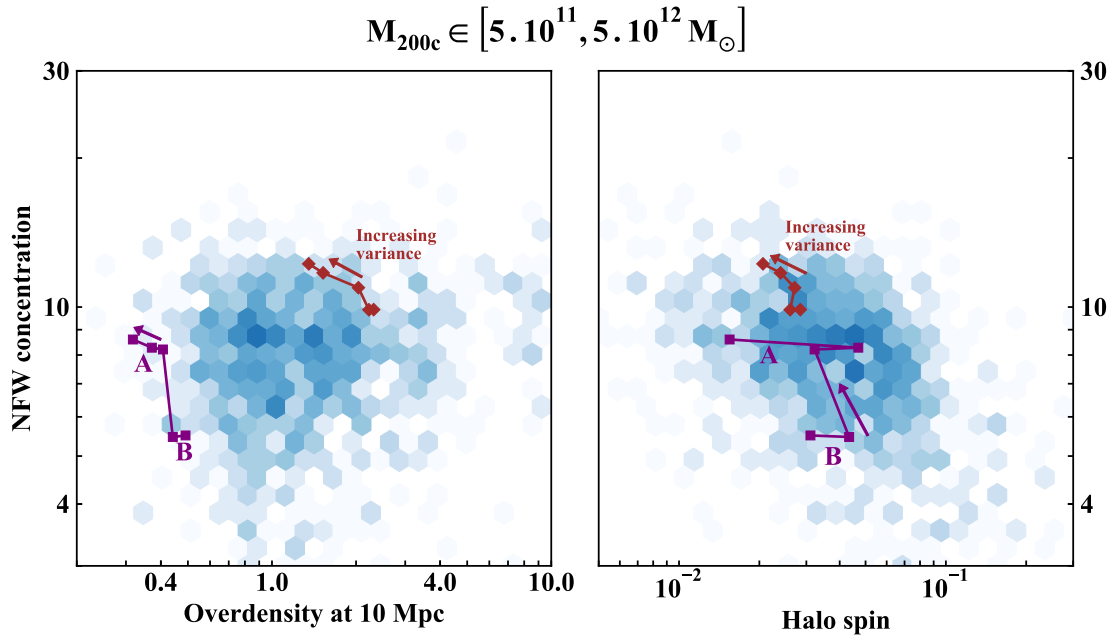


Fig. 3.5.: The impact of variance modifications on the local environment of haloes (left) and halo spin (right). Increasing the variance generates modified haloes in slightly underdense environments compared to their previous version, because internal variance is correlated with a halo’s external density in a Λ CDM universe. We recover a weak correlation between halo spin and halo concentration. Halo 740 shows a jump in spin between its two highest variance scenarios, which can be explained by a late-time minor merger which has been pushed to $z < 0$ (i.e. into the future) in the highest variance, lowest-spin case.

compared to the reference halo as a compensation effect. This directly translates into a higher environment density as observed in Figure 3.5.

In other words, the drift in environment is a result of the algorithm seeking to maintain maximally likely surroundings for the particular field realisation inside the target halo region. Given that environments correlate very weakly with halo concentrations and spin at this mass scale, we do not try to correct for this drift. Nonetheless, if a conserved environmental density is paramount, this can be accomplished with further stipulations on the surrounding field. For example, one may add an additional modification explicitly maintaining constant the mean overdensity of the environment on a chosen scale. In the limit of a perfectly conserved environment, each pixel outside the Lagrangian region could be fixed to its reference value. Though technically feasible, such a drastic approach might result in an unlikely draw from the Λ CDM power spectrum as all degrees of freedom in the field would be constrained. We will investigate the strengths and drawbacks of such an approach in future work.

3.4.4 Halo spin

Finally, we explore the effect of variance modifications on angular momentum. Halo spin and concentration have a weak but significant correlation (Macciò et al. 2007) and spin is often used

as proxy for galactic properties, especially for disc formation (see Somerville et al. 2008; Benson 2012, though see also Jiang et al. 2019).

The right-hand panel of Figure 3.5 shows the evolution of halo spin at $z = 0$ and concentration for both halo families. Tracks are compared to a halo population extracted from the same mass bin as in the left-hand panel. A weak correlation is observed in the halo population, similar to Macciò et al. (2007) but remains tentative due the low number of haloes in our simulation volume.

Halo 839's gradual evolution seems to follow the direction of the correlation, but halo 740 presents a more chaotic behaviour. The previously reported discontinuity in halo concentration between scenarios (A) and (B) is only linked to a change from $\lambda_{B01} = 0.032$ to 0.043 , i.e. $\Delta\lambda_{B01} = 0.48\sigma$ where σ is defined on the log of the population since halo spin are log-normally distributed (Bullock et al. 2001b). However, another discontinuity is visible in halo spin, previously invisible in halo concentration, between the two highest variance points of halo 740 (the reference halo and the +10 per cent halo). This discontinuity ranges from $\lambda_{B01} = 0.047$ to 0.015 with increasing variance, corresponding to a significance of $\Delta\lambda_{B01} = 1.8\sigma$ when compared to the overall spin population.

Unlike for halo concentrations, this jump in halo spin is not tied to a reconfiguration of the merger tree topology. Halo spin has been observed to peak around merger times (Vitvitska et al. 2002), as merging bodies bring in fresh angular momentum during in-fall. Halo spin then decays as smooth halo evolution proceeds and the in-falling body is destroyed. In the reference run, a small merger (mass ratio 1:12) happens at $z = 0.2$, making halo spin peak but without enough time to decay before the end of the simulation. In the +10 per cent halo, the same merger has not yet been incorporated in the main progenitor by the end of the simulation at $z = 0$, hence not generating the same increase in halo spin and yielding an overall much lower $z = 0$ spin value.

The fact that a minor merger can generate a variation in spin comparable to the breadth of the entire population highlights the sensitivity of the halo spin to the detailed merger history. The amplitude and direction of spins, the mass ratio and the orbital in-fall of the two merging bodies all likely play a role; we leave a more detailed analysis of their interplay as future work.

3.5 Conclusion

We have demonstrated our ability to construct alternative versions of a cosmological dark matter halo with fixed final halo mass and varying smoothness for the mass accretion history.

This turns merger histories into a tuneable parameter, opening a new route for galaxy formation and dark matter halo numerical studies.

Control over the smoothness of the mass accretion history is achieved by modifying the local variance on a given scale inside the Lagrangian region of a halo, while maintaining the Gaussianity and the ensemble power spectrum. We implemented the algorithm presented in Chapter 2 for simultaneous linear and quadratic modifications, and applied it successfully in the context of cosmological zoom simulations (see Figure 3.1).

We demonstrated our ability to successfully reduce or increase the mass ratios of multiple mergers (see Figure 3.2) by respectively reducing or enhancing the variance, achieving direct control of the overall smoothness of MAHs at fixed final halo mass. The implementation allows us to target mergers of a given mass scale, making minimal impact on the remaining MAH and large-scale structure around the halo.

The strength of this framework lies in its incremental approach. By gradually varying the local variance, we can causally explore the internal response of dark matter haloes. Specifically, we targeted “tipping points” for which a small change in MAH has large consequences on observed properties. In this way, we established that a reconfiguration in the merger tree topology was responsible for generating a large jump in halo concentrations (see Figure 3.3 and 3.4), comparable to the overall population-level scatter. This kind of dramatic sensitivity to initial conditions may be responsible for much of the scatter in halo and galaxy properties. The sensitivity is not equivalent to stochasticity arising due to chaotic amplification of numerical inaccuracies (Keller et al. 2019; Genel et al. 2019), which we explicitly ruled out as playing any part in our results. Rather, it arises naturally from the mapping of continuous initial conditions to a discrete set of mergers. GMs are therefore highly complementary to population-level studies for understanding how correlations and their scatter emerge through the complexity of halo formation.

In addition to dark matter halo physics, the method will prove invaluable to investigate the role of merger histories in shaping properties of the galaxy population. Mergers can, for example, directly impact star-formation rates by fuelling starbursts (Springel et al. 2005) or alter galaxy morphology (Di Matteo et al. 2005; Naab et al. 2009; Johansson et al. 2009; Pontzen et al. 2017). In particular, the fragility of dwarf galaxies makes them acutely sensitive to their merger and mass growth histories (e.g. Sawala et al. 2016; Fitts et al. 2017). In the next chapter, we will apply a combination of genetic modifications and high-resolution zoom simulations to show

how the variety of possible histories shapes a strong diversity in the properties of ultra-faint dwarf galaxies.

The origin of scatter in ultra-faint dwarf stellar masses and surface brightness

Our ability to control the mass accretion history of an object, developed and demonstrated in Chapters 2 and 3 is now mature and ready to be applied to galaxy formation studies. In this chapter, we focus on the formation and evolution of faint dwarf galaxies. We expect the properties of such low-mass galaxies to be acutely dependent on their precise build-up of mass, thereby providing an appealing playground for genetic modifications.

Results presented in Section 4.3, 4.4 and 4.5 are published in Rey et al. (2019a), accepted in ApJL. The galaxy formation model and numerical setup described in Section 4.2 is the result of collaborative work and is more extensively described in Agertz et al. (2019), of which I am a co-author.

4.1 Introduction

The advent of digital, wide sky photometric surveys is revealing a vast population of low surface brightness galaxies. At the faintest end with absolute V-band magnitudes $M_V \geq -8$, are “ultra-faint” dwarf galaxies, which are amongst the lowest-mass objects able to form stars in a Λ CDM universe (see Simon 2019 for a review).

Analysis of stellar populations within ultra-faint dwarf galaxies reveals that their stars have typical ages approaching that of the Universe (e.g. Brown et al. 2014; Weisz et al. 2014). This implies an early truncation of star formation, thought to arise because the galaxies’ potential wells are too shallow to accrete and cool gas once cosmic reionization has heated the surrounding intergalactic medium at $z \sim 6$ (Efstathiou 1992). Reionization is powered by the entire population of galaxies and quasars, and therefore, to a first approximation, can be modelled without taking account of local conditions (though see Katz et al. 2019). However, galaxies with a given halo mass today have formed at different rates over cosmic time, and therefore had a wide range of masses at the time of reionization. This may lead to a large diversity in the properties of ultra-faint galaxies, depending on the specific coupling between

the galaxy’s history and the timing of cosmic reionization (e.g. Benitez-Llambay et al. 2015; Sawala et al. 2016; Fitts et al. 2017).

Quantifying this expected scatter will be key to interpreting findings from forthcoming surveys (e.g. the Large Synoptic Survey Telescope; LSST). In particular, the low-mass of ultra-faints makes them particularly suited to identifying any fingerprints of alternative dark matter models (see Pontzen and Governato 2014 for a review). Meeting this promise requires us to model the formation of galaxies with a range of cosmological histories, each with sufficient resolution to resolve the interstellar medium and astrophysical processes within such small objects (Munshi et al. 2017; Macciò et al. 2017; Wheeler et al. 2018; Agertz et al. 2019).

In this chapter, we couple our cosmological, high-resolution zoom simulations (Ageretz et al. 2019) to the genetic modification framework (Roth et al. 2016; Chapter 2). This method generates alternative initial conditions for a cosmological galaxy, each new version varying a specific aspect of the galaxy’s mass accretion history (hereafter MAH). Each history is simulated independently, reproducing the same large-scale environment and final dynamical mass. This enables a controlled study, allowing us to construct a causal account of the link between history and observables.

4.2 Genetically modified dwarf galaxies

4.2.1 Initial conditions

We start by creating zoom initial conditions (see Agertz et al. 2019 for a more detailed description of the procedure) for an unmodified, reference galaxy. We first simulate a dark-matter-only cosmological volume, with a size of 50 Mpc at a resolution of 512^3 (i.e. $4 \times 10^7 M_\odot$ dark-matter particle mass). We then select the largest void in this volume, track it to the initial conditions and re-simulate it with a zoomed dark-matter-only simulation with resolution equivalent to 2048^3 (i.e. $6 \times 10^5 M_\odot$ dark-matter particle mass). We then identify haloes using the HOP halo finder (Eisenstein and Hut 1998) and compute their virial mass, where M_{200c} defines the mass enclosed within a sphere of radius r_{200c} encompassing 200 times the critical density of the Universe. We select our reference halo as an isolated central, with no massive neighbors within $5 r_{200c}$, and a present-day virial mass of $M_{200c} = 1.5 \times 10^9 M_\odot$. This halo is tracked again to the initial conditions and re-simulated with our final, zoomed resolution equivalent to 16384^3 (i.e. $960 M_\odot$ dark-matter particle mass), hydrodynamics and galaxy formation physics (see Section 4.2.2).

We then construct, “genetically modified”, alternative initial conditions for this galaxy, modifying its halo mass around reionization, while fixing the halo mass today. The halo mass of an object can be directly controlled from the initial conditions by modifying the height of its associated density peak (Roth et al. 2016). We therefore identify all particles within the major progenitor at redshift $z = 6$ and increase (decrease) the mean density within this region to increase (decrease) the halo mass at $z = 6$. We conserve the final, $z = 0$, halo mass by maintaining a constant mean density in the corresponding Lagrangian region. We emphasize that each modified initial condition makes minimal changes to the surrounding environment, maintaining the same large-scale filamentary structure around the galaxy (e.g. Figure 3.1).

4.2.2 Galaxy formation model

We evolve the modified and reference initial conditions to $z = 0$ using cosmological zoomed galaxy formation simulations. We follow the evolution of dark matter, stars and gas using the adaptative mesh refinement hydrodynamics code RAMSES (Teyssier 2002). The dynamics of collisionless particles (dark matter and stars) are computed using a multi-scale particle-mesh solver, estimating densities through a cloud-in-cell approximation (Guillet and Teyssier 2011). Fluid dynamics are computed using an HLLC Riemann solver (Toro et al. 1994) and the fluid equations are closed by assuming an ideal gas equation of state with adiabatic index $\gamma = 5/3$. We include an extensive galaxy formation model described in detail by Agertz et al. (2019) as “Fiducial”, of which we describe the most important ingredients next.

Gas cooling We track the cooling of a primordial plasma using hydrogen and helium equilibrium thermochemistry (Courty and Alimi 2004) following photoionisation, collisional ionisation and excitation, recombination, bremsstrahlung, Compton cooling and heating and dielectronic recombination. Cooling from metal lines is extracted from tabulated models generated with CLOUDY (Ferland et al. 1998) and we model reionization as a spatially uniform, time-dependent heating source similarly to Chapter 4.

Once gas is allowed to cool, it can collapse under Jeans’ instability and reach densities high enough to self-shield against surrounding radiation. We include an on-the-fly self-shielding prescription (Aubert and Teyssier 2010; Rosdahl and Blaizot 2012) such that gas cells with hydrogen densities $n_{\text{H}} \geq 0.01 \text{ cm}^{-3}$ are associated cooling and heating rates evaluated at the fictional, enhanced density:

$$n_{\text{H, boosted}} = n_{\text{H}} \exp\left(\frac{n_{\text{H}}}{0.01 \text{ cm}^{-3}}\right) \quad (4.1)$$

This self-shielding threshold is fixed at all times and can only be approximate, as it neglects the physical dependency of self-shielding on the column density (rather than number density) of neutral hydrogen, as well as local ionizing conditions. The specific value is therefore calibrated to match full radiative transfer calculations (e.g. Pontzen et al. 2008; Faucher-Giguère et al. 2010; Rosdahl and Blaizot 2012).

Star formation Since gas cooling is increasingly efficient in denser regions, it proceeds as a run-away process and should lead to the formation of stars within a cooling time. Our simulations model the conditions through which gas is turned into stars stochastically, with a recipe following a Schmidt law:

$$\dot{\rho}_* = \epsilon_{\text{ff}} \frac{\rho_g}{t_{\text{ff}}} \text{ for gas cells with } \rho_g > \rho_* \text{ and } T_g > T_* \quad (4.2)$$

where $\dot{\rho}_*$ is the instantaneous star formation rate in a gas cell, ϵ_{ff} is the star formation efficiency per free-fall time, ρ_g and T_g are the gas cell density and temperature, $t_{\text{ff}} = \sqrt{3\pi/32G\rho}$ is the local gas free-fall time and ρ_* and T_* are imposed thresholds that gas must satisfy to qualify for star formation. For every gas cell satisfying these conditions (i.e. densities higher than $\rho_* = 300 m_p \text{ cm}^{-3}$ and temperatures lower 100 K), we sample Equation (4.2) stochastically through a Poisson process, ensuring that the mean number of stellar particles formed is proportional to $\dot{\rho}_*$ (Rasera and Teyssier 2006). Our stellar particles have initial masses of $\sim 300 M_\odot$ to ensure a complete sampling of our chosen initial mass function (IMF) (Kroupa 2001). The final properties of ultra-faint dwarf galaxies are sensitive to both our choices of star formation parameters, and the underlying assumptions of this implementation. We refer to Agertz et al. (2019) for a discussion motivating our chosen values against observational constraints, as well as an exploration of physically-motivated, alternative models tying star formation with the density of molecular hydrogen (see also Munshi et al. 2019).

Stellar Feedback A key aspect of our simulations is the modelling of feedback from massive stars, accounting for Type II and Ia supernovae explosions, as well as stellar winds from massive and asymptotic giant branch (AGB) stars. The amount of energy, momentum, mass and metallicity returned to the interstellar medium by each mechanism is extensively described in Agertz et al. (2013). Our model naturally tracks the stellar evolutionary timescales of stars with different masses within a stellar particle, ensuring that feedback from each process is injected on timescales at which it is known to operate.

At a given simulation timestep, we compute in each stellar particle the number of stars exiting the main sequence and turning into Type II supernovae (Eq. 6 in Agertz et al. 2013). This IMF-averaged number is then randomly sampled through a Poisson process to obtain a discrete number of supernovae within a timestep (Ageretz et al. 2019). The resolution of our simulations greatly reduces uncertainties in modelling the energy injection of supernovae to the surrounding medium. We refine down to a maximum spatial resolution of 3 pc and follow dark matter particles with masses $960 M_{\odot}$. This resolution is sufficient to capture the cooling radius of individual supernovae, allowing us to directly inject a thermal energy of 10^{51} erg in a gas cell and self-consistently follow the build-up of momentum by solving the hydrodynamics equations (Kim and Ostriker 2015; Martizzi et al. 2015). We further include the energy, mass and metallicity returned from Type Ia explosions, assuming they originate from white dwarfs accreting mass from a binary companion. We use a similar approach as previously described for Type II supernovae – we compute the number of stars turning Type Ia supernovae in a stellar particle (Eq. 13 in Agertz et al. 2013) assuming a binary mass function (Raiteri et al. 1996), discretely sample this IMF-averaged number, and inject 10^{51} erg in the parent gas cell for each Type Ia explosion.

Finally, we include the contribution of stellar winds from early, massive stars and late, AGB stars. We model winds from massive stars as a continuous release of energy, momentum, mass and metallicity (Eq. 4 in Agertz et al. 2013) over the first 6.5 Myr of the lifetime of a stellar population. The wind budget in each released component is calibrated to match a stellar evolution model (Leitherer et al. 1999), for stellar populations with varying metallicities. Significant mass-loss also occurs due to winds from AGB stars. We compute the IMF-averaged mass loss due to AGBs over the lifetime of each stellar particle (Eq. 17 in Agertz et al. 2013) and continuously inject the released mass in the parent gas cell.

The major regulating mechanism for our galaxies are explosions from Type II supernovae. Despite our resolution and their accurate modelling, additional feedback channels, e.g. photo-heating, can strengthen or weaken their coupling to the surrounding interstellar medium (e.g. Smith et al. 2019; Agertz et al. 2019). To probe the sensitivity of our results to such residual uncertainties, we evolve all initial conditions with an alternative model, reducing the efficiency of supernova feedback ("Weak feedback" model in Agertz et al. 2019). This model introduces arbitrary temperature and velocity ceilings for supernovae ejecta (10^8 K and 1000 km s^{-1} respectively), thereby limiting their efficiency in driving winds and regulating star formation. We stress that this model should be seen as an explorative test, rather than as an alternate physical prescription.

For each galaxy, we compute the V-band luminosity of each stellar particle using a single stellar population model interpolated over a grid of ages and metallicities (Girardi et al. 2010) and sum them to derive the total magnitude. We choose a random line-of-sight to obtain the projected half-light radius, $r_{1/2}$, and checked that using unprojected 3D half-light radius does not modify our observed trends. Finally, we compute the one-dimensional stellar velocity dispersion as $\sigma_{\star} = \sqrt{\sigma_{\star,x}^2 + \sigma_{\star,y}^2 + \sigma_{\star,z}^2} / \sqrt{3}$ and the total iron metallicity, $[\text{Fe}/\text{H}]$, as the mean of each stellar particle’s iron abundance weighted by its stellar mass (Escala et al. 2018; Agertz et al. 2019).

4.3 Growing the stellar mass of ultra-faints

We show the resulting four genetically modified MAHs in Figure 4.1, top panel. Our modifications generate a range of halo mass growth before cosmic reionization, while these histories have converged by $z \sim 2$ and reach the same dynamical mass today. To illustrate that we probe a cosmologically representative range of histories, we compare these tracks with a statistical sample of $\sim 1\,500$ histories extracted from the parent, lower resolution volume of our zoom simulations. We select central haloes with masses between 0.9 and $4 \times 10^9 M_{\odot}$ at $z = 0$ and compute their fractional mass growth, i.e. their mass growth divided by their total mass. We show the median with 64 and 95 % confidence intervals at each redshift (grey bands), normalised to $1.5 \times 10^9 M_{\odot}$. The four MAHs lie within the 95% contours of the overall population, demonstrating that our objects range across the majority of the population’s scatter in early histories. We stress that this comparison should be seen as qualitative rather than as a rigorous statistical test; we leave more detailed statistical inference to a future work.

In the bottom panel of Figure 4.1, we show the growth of stellar mass of each genetically modified history. As an ultra-faint dwarf galaxy forms earlier (i.e. achieves a higher mass at reionization) its final stellar mass grows. Since the environment surroundings and final mass are all fixed, this trend demonstrates a direct mapping between the halo mass achieved before reionization and the final stellar mass of an ultra-faint.

The link is best explained by the duration of the star-forming phase for each galaxy. Figure 4.1 shows that earlier forming galaxies systematically start building stellar mass earlier in time. We mark by a cross the time of last star formation activity in the main progenitor, showing that reionization quenches star formation at near-identical times ($z \sim 4$) for all mass accretion histories. Star formation continues shortly after the end of reionization from self-shielded cold gas within the halo (Oñorbe et al. 2015). However, heating from the UV-background prevents further gas accretion, quickly leading to starvation and permanent quenching. Earlier forming

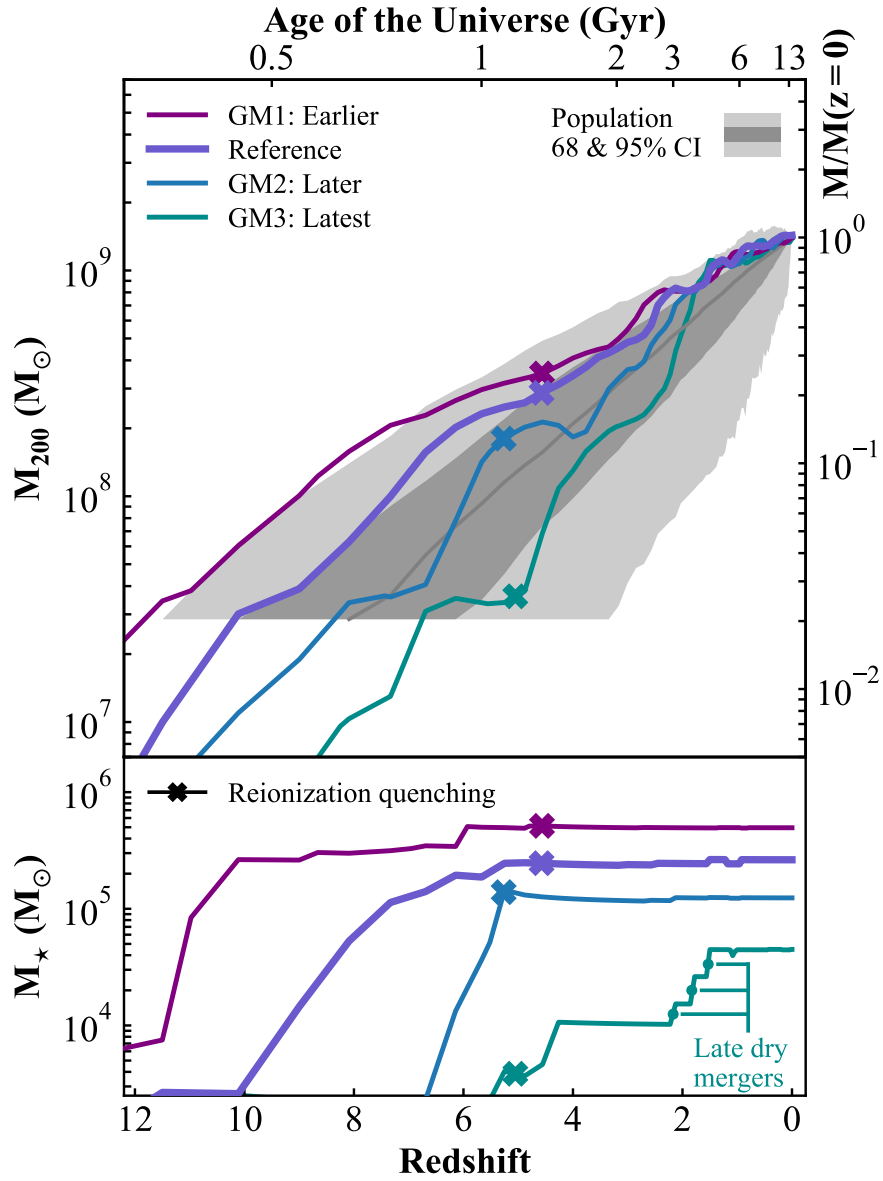


Fig. 4.1: Growth of dynamical mass (top panel) over cosmic time for our reference galaxy and its three genetically modified counterparts. We design each new history to vary the formation time of the reference galaxy, forming earlier (purple) or later (blue, turquoise). Earlier forming ultra-faints start assembling stellar mass (bottom panel) earlier in time, before reionization quenches in-situ star formation by preventing gas inflows (crosses). Earlier forming galaxies therefore have a systematically higher $z = 0$ stellar mass. By construction, all histories converge to the same dynamical mass today while scanning across a representative range of early histories (grey bands). This allows us to quantify scatter in the stellar mass at fixed halo mass (see Figure 4.2).

ultra-faints therefore have a longer period of star formation. In addition, earlier forming ultra-faints have a higher halo mass at a given time and consequently reach higher instantaneous star formation rates before reionization.

More extended and more vigorous star formation thus leads to a higher stellar mass today. We now examine the consequences for the relation between stellar mass and halo mass.

4.4 Generating scatter in the stellar mass-halo mass relation

The mass of a galaxy's dark matter halo is thought to be the primary driver of its properties, as it regulates the depth of the potential well and hence the overall availability of gas. This assumption allows empirical and semi-numerical models of galaxy formation to rely on a parametrized mapping between halo mass and stellar mass (see Wechsler and Tinker 2018 for a review). On the scale of dwarf galaxies however, the functional form of this mapping is highly uncertain (e.g. Behroozi et al. 2013; Brook et al. 2014; Garrison-Kimmel et al. 2017; Read et al. 2017). Our results above further show how the assumption of a one-to-one correspondence between stellar mass and halo mass breaks down at this mass scale.

We show in Figure 4.2 the growth of our modified galaxies in the stellar mass-halo mass plane. The final $z = 0$ points (color diamonds) can be compared with an extrapolated abundance matching prediction (orange-dashed, Read et al. 2017). Our different histories generate a spread in stellar mass over 1 dex, causally demonstrating the existence of extended scatter due to the variety of possible histories for a given ultra-faint.

Since our galaxies have by construction the same environment to isolate the role of histories, the extent of the exposed scatter is a lower bound on the overall population diversity. External factors such as tidal stripping during the infall into a more massive host can provide an additional mechanism to generate scatter on the scale of ultra-faints (Munshi et al. 2017).

To probe the robustness of our results to residual uncertainties in modelling galaxy formation, we show in Figure 4.2 our four modified histories evolved with the alternative, "Weak feedback", model. As expected, all histories have higher stellar masses compared to the "Fiducial" model (solid lines). However, the systematic trend of higher stellar mass with earlier forming galaxies remains, while the scatter in stellar masses at fixed halo mass increases from 1 to 1.3 dex. This approximate conservation of scatter reflects its origin in the relative timing of mass accretion and reionization (Section 3). We conclude that the strength of supernova coupling primarily

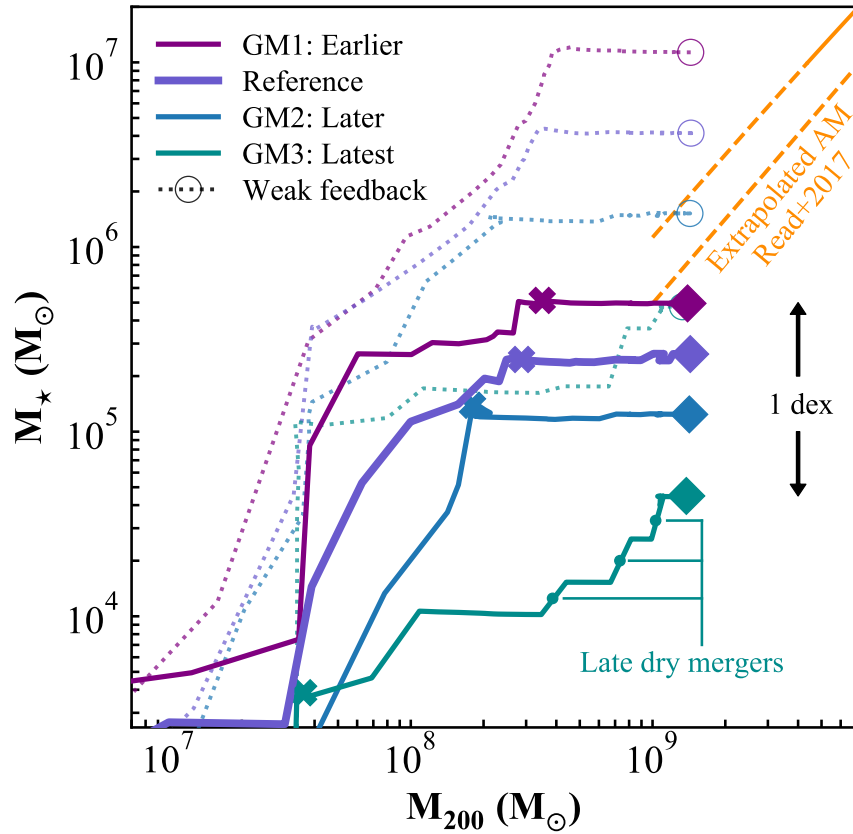


Fig. 4.2.: Stellar mass and halo mass growth of our four modified galaxies compared. The range of histories creates a 1 dex spread in stellar masses at fixed halo mass. Varying the implementation of supernova feedback (faint dotted) modifies the overall normalisation of the stellar masses, but relative differences between histories are conserved. This robustly demonstrates the existence of extended scatter in the relation between stellar mass and halo mass due to the variety of possible formation times. Our latest forming history (turquoise) further sees its mode of stellar mass growth modified, now being dominated by late-time dry stellar accretion. This new formation scenario has a strong impact on observable structural properties (Figure 4.3).

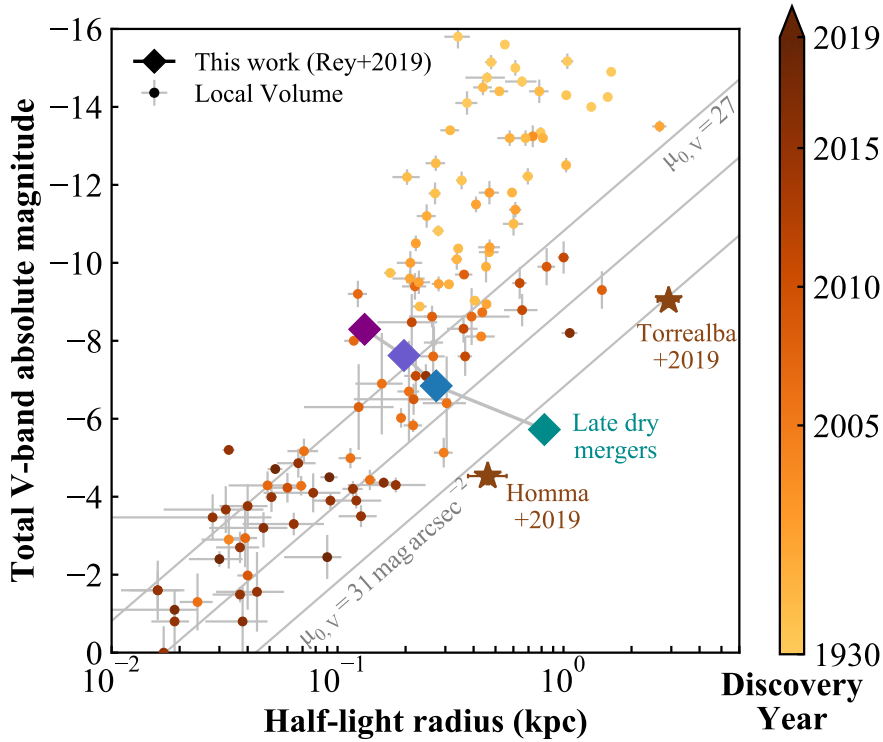


Fig. 4.3.: Impact of modifying the history of an ultra-faint on its V-band absolute magnitude and projected half-light radius. We compare our galaxies to a sample of observed dwarfs from the Local Volume (McConnachie 2012; Kirby et al. 2014; Simon 2019). All our simulated galaxies lie within current observational scatter except the latest forming dwarf that builds its stellar mass through dry mergers. These mergers deposit stars on wide orbits, creating an extremely diffuse ultra-faint. We color individual galaxies by their discovery year, highlighting two trends: towards fainter detections overall and more extended galaxies at a given magnitude. Our dwarf forming through dry mergers has a central surface brightness comparable to the latest discoveries of modern deep imaging surveys (e.g. Homma et al. 2019; Torrealba et al. 2019), highlighting prospects to uncover this diffuse population with e.g. LSST and future GAIA releases.

affects the *absolute* scale of stellar masses, while the diversity in formation times drives the relative scatter around the mean.

In addition to varying the total stellar mass, our modifications change the mode of stellar mass growth for the latest forming MAH (turquoise). Rather than forming stars in-situ, this galaxy accretes most of its stellar mass through late, dry mergers (labeled in Figures 4.1 and 4.2). We now show that this accretion-dominated scenario exposes a new formation pathway for extended, diffuse ultra-faint dwarf galaxies.

4.5 The formation of a diffuse ultra-faint

Figure 4.3 plots the total, absolute V-band magnitude at $z = 0$ against the projected half-light radius of our genetically modified galaxies (color diamonds). A compiled sample of observations from the Local Volume (McConnachie 2012; Kirby et al. 2013; Kirby et al. 2014; Simon 2019) is shown for comparison. We show lines of constant central surface brightness (Martin et al.

2016, Eq 8), assuming a mean ellipticity of 0.5 (Simon 2019). The three earliest forming galaxies (purple, violet, blue) all lie within the observational scatter, while the latest forming (turquoise) has a lower surface brightness than currently detected dwarf galaxies.

The extreme diffuse nature of this object arises due to its assembly from ex-situ accreted stars. Dry stellar mergers after $z = 6$ build 94% of the total stellar mass of this galaxy, depositing stars away from the galaxy centre and growing the galaxy's half-light radius to 820 pc today. The result is an extremely extended ultra-faint with low central surface brightness (~ 32 mag arcsec $^{-2}$). We therefore have exposed a new formation scenario for extended diffuse ultra-faint dwarf galaxies, arising through an early truncation of in-situ star formation by reionization but a later growth via ex-situ dry accretion.

The current rarity of observed analogues for our diffuse ultra-faint reflects observational challenges at this extremely low surface brightness. Each observed dwarf galaxy in Figure 4.3 is colored by the year of their discovery. Two trends are visible as experiments improved over the years: (i) downwards towards overall fainter objects and (ii) rightwards towards more diffuse objects at a given magnitude. Ongoing and next-generation surveys are expecting to pursue these trends, vastly expanding the census of ultra-faint dwarf galaxies. In particular, detections through resolved stars are reaching surface brightnesses similar to our diffuse galaxy, as demonstrated by two recent candidates for Milky-Way satellites (Homma et al. 2019; Torrealba et al. 2019). LSST will further be able to detect individual stars of a $\mathcal{M}_V = -6$ galaxy out to several Mpc, directly probing ultra-faint formation in the field. Future data releases from GAIA will also extend the Milky-Way census, shedding light on the fate of such low surface brightness objects as they fall into a massive host. We therefore expect our prediction of the existence of ultra-faint, diffuse field galaxies to be testable in the near-future.

The mass growth associated to our diffuse ultra-faint is within the 68% contour of the overall population (see Figure 4.1). We therefore conclude that our exposed formation scenario is likely generic for field ultra-faint dwarfs, hence predicting the existence of a population of low-surface brightness ultra-faints waiting to be uncovered by the next-generation of deep imaging cameras. We caution however that the visual comparison with the population's contours is at best incomplete and that a more detailed statistical analysis is required to predict the number density of diffuse ultra-faints. As an example, the likelihood of a galaxy's growth history does not fully account for the likelihood of late major mergers (e.g. Deason et al. 2014), which are key to our formation scenario. Such statistical analysis is outside the scope of this thesis and we leave it as future work (Kim et al. in prep).

Spectroscopic follow-up of stars within such a diffuse galaxy could confirm its formation pathway. Its stellar velocity dispersion $\sigma_* = 7.4 \text{ km s}^{-1}$ is comparable to that of our reference galaxy (7.1 km s^{-1}), highlighting the weak constraining power of stellar dispersions on formation scenarios. However, the assembly from multiple low stellar mass systems produces an extremely metal-poor galaxy with $[\text{Fe}/\text{H}] = -2.9$, compared to -2.4 in our reference case where 88% of the total stellar mass is formed in-situ. A low metallicity at a given stellar mass would therefore complement the extended size as signatures of this new formation scenario.

4.6 Conclusion

We have presented results demonstrating how formation history affects the properties of field ultra-faint dwarf galaxies. We created a series of four “genetically modified” zoom initial conditions (Roth et al. 2016, Chapters 2 and 3) for a single object, systematically varying its accretion history up to the time of cosmic reionization while fixing its $z = 0$ dynamical mass. We evolved these initial conditions with high-resolution zoom cosmological simulations (Agertz et al. 2019) and computed the response of the central galaxy’s observables.

By construction, all our histories converge to the same dynamical mass today and evolve within the same cosmological environment, thereby creating a controlled study. We demonstrate that the halo mass achieved before reionization directly controls the final stellar mass of an ultra-faint (Figure 4.1). Earlier forming galaxies begin forming stars when the Universe was younger and have more vigorous star formation rate at a given time, therefore assembling higher stellar masses before their quenching by reionization. We further show that the variety of possible histories for an ultra-faint leads to an extended scatter in the relation between stellar mass and halo mass (Figure 4.2). This scatter arising from histories is robust to a large variation in our implementation of stellar feedback.

Probing the interaction between merger histories and reionization allows us to expose the potential for highly diffuse, ultra-faint dwarf galaxies. Extremely low surface brightness can be achieved through an early truncation of in-situ star formation and a later growth by stellar accretion, vastly growing stellar size (Figure 4.3). Finding such a population will be within the reach of future facilities such as LSST.

Our study cleanly demonstrates the importance of cosmological histories in explaining the diversity of dwarf galaxy properties. This extends previous results (e.g. Fitts et al. 2017) by (i) targeting smaller dynamical mass haloes more likely to host observed ultra-faints (Jethwa et al. 2018; Read and Erkal 2019) (ii) improved numerical resolution, reaching the critical scale of

the supernova cooling radius and (iii) cleanly isolating the role of histories using the genetic modification technique. Nonetheless, the interaction between histories and reionization is only one factor in determining the full diversity of the ultra-faint dwarf galaxy population (Sawala et al. 2016). We will investigate other sources of diversity such as environment (e.g. Munshi et al. 2017) in future work and turn in the following chapter to the importance of halo masses in shaping the stellar populations of ultra-faint dwarfs.

Rejuvenating star formation in ultra-faint dwarfs

Chapter 4 demonstrated ultra-faint dwarf galaxies with different dynamical masses at the time of reionization had varied stellar masses. Despite this diversity in the total amount of stars formed, every presented galaxy was quenched with a stellar population uniformly old (~ 10 Gyr). In this chapter, we use a larger suite of cosmological zoom simulations to show how ultra-faint dwarf galaxies reaching sufficiently high dynamical masses at late times can accrete and cool gas. This mechanism, leading or not to the rejuvenation of star formation, predicts an additional diversity in the stellar populations and gas content of ultra-faint dwarfs.

Work shown in this chapter is an ongoing analysis. Several open questions are highlighted throughout, which we plan on clarifying before an eventual publication.

5.1 Introduction

The end of reionization and the establishment of a strong background of ultraviolet (hereafter UV) radiation plays a key role in shaping the properties of ultra-faint dwarf galaxies. It provides a near-uniform source of heating across the Universe, raising the Jeans mass of the intergalactic medium (hereafter IGM), and in turn preventing gas accretion in haloes with dynamical masses below a few $10^9 M_{\odot}$ (Efstathiou 1992; Gnedin 2000; Hoeft et al. 2006; Okamoto et al. 2008; Noh and McQuinn 2014; Chapter 4).

The extremely low-mass galaxy Leo T ($M_{\star} \sim 10^5 M_{\odot}$, $\mathcal{M}_V \sim -7$) has challenged this scenario for over a decade. Its resolved stellar population exposed the existence of a young stellar population born in the last 200 Myr (Irwin et al. 2007). In addition to forming stars today, deeper photometric measurements show that Leo T has formed stars at a nearly constant rate over the full age of the Universe (de Jong et al. 2008; Weisz et al. 2012; Clementini et al. 2012), with average star formation rates as low as $10^{-5} M_{\odot} \text{ yr}^{-1}$. Follow-up in the radio wavelengths show an extended HI content dominating in mass the stellar body (Ryan-Weber et al. 2008; Adams and Oosterloo 2018), with a cold and warm phase as expected from an interstellar medium. The survival of Leo T to reionization is puzzling due to its estimated dynamical mass. Matching the observed dynamical properties of the gas with simulations motivates

$M_{200c} \sim 5 \times 10^8 M_{\odot}$ (Read et al. 2016), well below the mass scale at which the UV background should prevent gas accretion and star formation.

The existence of Leo T could be regarded as a statistical peculiarity when modelling the ultra-faint population. However, due to their cold gas content, Leo-T like objects could to be detectable by their HI signatures throughout the Local Volume. Catalogues of HI-selected clouds from GALFA-HI (Peek et al. 2011) and ALFALFA (Haynes et al. 2011) can be matched to overlapping deep optical imaging, respectively uncovering zero (DeFelippis et al. 2019) and five (Janesh et al. 2019) gas-rich ultra-faint candidates. Similar HI selection further exposed a higher mass star forming, gas-rich dwarf (McQuinn et al. 2015) demonstrating that the census of nearby gas-rich faint dwarf galaxies is unlikely to be complete.

The prospect of uncovering a population of gas-rich ultra-faints in the local field by combining HI selected clouds with deep optical imaging from e.g. LSST motivates further exploration of mechanisms through which ultra-faint can replenish their gas content and form stars after reionization. Moreover, such galaxies would provide ideal laboratories to study the nature of dark matter, as late-time star formation could provide sufficient energy to transform dark matter cusps into cores deep in the ultra-faint regime (see Pontzen and Governato 2014 for a review). The extremely low star formation rates also offer key opportunities to study stochastic star formation in low metallicity environments and the corresponding emerging initial mass function (Bastian et al. 2010).

Re-igniting star formation requires gas to be re-accreted onto the central dwarfs (see Ricotti 2009 for analytical insights). Compression of nearby gas by ram pressure when encountering a dense structure, e.g. a filament or an outflow from a nearby massive galaxy, provides such a mechanism (Wright et al. 2019) but would not apply for field isolated systems. In this chapter, we use a suite of cosmological, high-resolution zoom simulations to show that late gas accretion and rejuvenation occurs naturally when the dynamical mass grows sufficiently to allow gas cooling. We further complement these simulations by genetically modified initial conditions (Chapter 2) to expose the strong diversity in gas content and stellar populations of ultra-faints depending on the variety of their late build-up of mass.

5.2 Numerical setup

We first create high-resolution zoom initial conditions for four haloes selected from a parent lower resolution volume using the same procedure as in Chapter 4. The properties of these

Simulation name	$M_{200c} (M_{\odot})$	$M_{\star} (M_{\odot})$	$M_{\text{gas}} (M_{\odot})$	SF at $z=0$
Rejuvenated, major mergers	3.3×10^9	5×10^5	2×10^5	Yes
↔ GM: later mergers	3.2×10^9	4×10^5	1×10^4	Yes
Rejuvenated, steady growth	3.2×10^9	2×10^6	9×10^5	Yes
Quenched, gas-rich	2.5×10^9	6×10^5	1×10^6	No
↔ GM: higher mass	3.7×10^9	1.5×10^6	5×10^5	Yes
Dead	1.4×10^9	5×10^5	2×10^3	No

Tab. 5.1.: Simulations presented in this chapter and their dynamical halo mass, stellar mass and gas mass within the inner 1 kpc at $z = 0$ (second, third and fourth column respectively). The last column indicates whether the galaxy has formed a stellar population in the last 500 Myr. We present four, independent galaxies separated by horizontal lines, as well as two genetically modified variants.

objects are summarised in Table 5.1 and their mass growth history is shown in Figure 5.1. We now briefly describe how we select and evolve these galaxies to $z = 0$.

We choose four dark matter haloes within a narrow window in present-day halo mass, ranging from $M_{200c} = 1.5 \times 10^9 M_{\odot}$ to $3.5 \times 10^9 M_{\odot}$. The goal of this suite is to sample a range of different environments, halo masses and histories thus providing a base for later resimulation with the genetic modification approach (see Section 5.5). We select two high-mass haloes with contrasting histories, one dominated by major mergers at $z \sim 3$ (purple in Figure 5.1) compared to a second with a steady mass growth (blue). We further select a low mass halo (black, described in Chapter 4 as “Earlier”) with a high dynamical mass at the time of reionization but little mass growth thereafter. We close our selection by an intermediate object (brown) in both history and final mass.

We evolve all initial conditions to $z = 0$ using the adaptative mesh refinement code RAMSES (Teyssier 2002), complemented by an extensive galaxy formation model described in Chapter 4 and Agertz et al. (2019). We track the cooling of primordial gas through atomic processes and metal lines (Courty and Alimi 2004), self-shielding above a fixed density threshold (Rosdahl and Blaizot 2012) and heating through a spatially uniform UV background (modified from Haardt and Madau (1996) as in the RAMSES public version).

We model star formation following a Schmidt law, for gas cell with densities above $\rho_{\star} = 300 m_p \text{ cm}^{-3}$ and temperatures lower 100 K (see Chapter 4 and Agertz et al. 2019). As emphasized in Chapter 4, a key aspect of our simulations is the modelling of feedback from massive




Fig. 5.1.: Growth of dynamical mass (top panel) for four ultra-faint dwarf galaxies, spanning a range of dynamical masses at reionization and today. All cumulative star formation histories (bottom panel) plateau by $z \sim 4$, indicative of reionization quenching. The two galaxies achieving the highest dynamical masses today (purple and blue) exhibit late-time star formation after $z = 1$ (stars) and continuously form new stellar populations until the present day. This diversity in star formation histories is best explained by the virial temperature (right axis) of these two galaxies growing over the temperature of the surrounding IGM after reionization (grey dashed), leading to the re-accretion of gas. The intermediate mass galaxy (brown) also grows over this same threshold and is currently undergoing, but has not yet completed, the rejuvenation process (Figure 5.2 and 5.4).

stars. We explicitly track the evolutionary timescales of massive stars (Agertz et al. 2013) within each stellar particle by evolving a single stellar population model (Leitherer et al. 1999). This allows us to independently model the energy injection from exploding Type II and Type Ia supernovae. Furthermore, the resolution of our simulations is sufficient to resolve the cooling radius of supernovae explosions¹, allowing us to model them by directly injecting the required thermal energy in a gas cell and self-consistently follow the build-up of momentum. This approach significantly reduces uncertainties arising from modelling the coupling between supernova feedback and the surrounding interstellar medium.

¹with a maximum spatial resolution of 3 pc and dark matter particles masses of $960 M_{\odot}$, see Chapter 4.

5.3 Late time star formation in ultra-faint dwarfs

Figure 5.1 shows the growth of dynamical mass for each of the four galaxy in the suite (top panel) and the response of their cumulative, archeological star formation history (bottom panel). We compute archeological star formation histories as would be derived from observing a color-magnitude diagram of the galaxy in our Local Group today (e.g. Weisz et al. 2012). We select all stars in the galaxy at $z = 0$ and compute the fraction of final stellar mass formed before a given time. This cumulative history therefore encodes when stars were formed across *all* progenitors, rather than just the most massive object.

All galaxies exhibit a plateau in star formation after $z \sim 4$. This halt of star formation is a signature of cosmic reionization heating the IGM and preventing these galaxies to accrete fresh gas. Self-shielded gas within the halo allows star formation to proceed after the end of reionization ($z \sim 6$) but this reservoir is never replenished, hence quickly leading to quenching (Oñorbe et al. 2015; Agertz et al. 2019; Chapter 4).

However, we observe a diversity in the fraction of stars formed prior to reionization. By $z \sim 4$, the “Dead” and “Quenched gas-rich” histories (respectively black and brown) have formed 100% of their stars, thus exhibiting an uniformly old stellar population today. On the other hand, the two highest mass galaxies of our sample (purple and blue) have formed respectively 19% and 2 % of their stars after $z = 1$. We highlight in the top panel the birth of the first generation of late-time stars and stress that both galaxies have formed a stellar population in the last 600 Myr.

This bimodality of star formation is best explained by the interaction between the growth history of each galaxy and a dynamical mass “threshold” to accrete gas set by the temperature of the surrounding IGM (Efstathiou 1992; Gnedin 2000; Hoefl et al. 2006; Okamoto et al. 2008; Noh and McQuinn 2014). To visualize this interaction, we compute the time evolution of each galaxy’s virial temperature defined as $T_{200c} = \frac{\mu m_p G M_{200c}}{2 k_B r_{200c}}$ (right axis), where $\mu = 0.59$ is the mean molecular weight of the gas, m_p is the proton mass and G and k_B are the gravitational and Boltzmann constants. If $T_{200c} \gg T_{\text{IGM}}$, the potential well of the galaxy is sufficiently deep to overcome thermal pressure and accretion is possible.

Reionization heats the IGM to $T_{\text{IGM}} \sim 2 \times 10^4$ K (see McQuinn 2016 for a review), above the corresponding virial temperature of all galaxies at $z \sim 6$. Gas accretion is then prevented, quickly leading to star formation quenching. The three most massive galaxies are however able to grow sufficiently in dynamical mass to fulfill $T_{200c} \gtrsim T_{\text{IGM}}$ at late times. They should then

be able to accrete gas from the IGM, hence explaining how they can replenish their gas content and rejuvenate star formation.

However, only the two most massive galaxies, rather than all three, rejuvenate star formation by $z = 0$ showing that this argument is likely over-simplified. Furthermore, the narrow spread of our galaxies in virial temperature makes our argument particularly sensitive to the exact value chosen for the IGM temperature, when in fact, the state of the IGM is both time and density dependent. We have first assumed that the temperature of the IGM is constant with redshift, thus neglecting heating from the reionization of helium and cooling from adiabatic expansion (McQuinn et al. 2009; Upton Sanderbeck et al. 2016). Secondly, temperature and density in the IGM are correlated through an equation of state (Hui and Gnedin 1997) defined by the balance between adiabatic cooling and photo-heating. The relevant density at which the temperature should be evaluated for halo accretion is however debated (Hoeft et al. 2006; Okamoto et al. 2008; Noh and McQuinn 2014). Both factors can modify the IGM temperature by a factor of a few, hindering a detailed comparison with galaxy virial temperatures. All these factors are self-consistently modelled in our simulation through a time-dependent UV background and we plan in forthcoming work to explore the sensitivity of our findings to different helium reionization models.

In addition, we have so far neglected the physics of gas cooling and heating as densities increase towards the halo centre. These are likely to play a critical role in establishing conditions to re-ignite star formation (Ricotti 2009; Wright et al. 2019). We will now provide a more detailed picture of the evolution of gas content during the rejuvenation transition.

5.4 The fate of gas during the rejuvenation process

Figure 5.2 shows the evolution of inner gas content for three ultra-faints with increasing late-time halo mass. We only present one high-mass galaxy (purple) for visual clarity as the second (“Rejuvenated, steady growth”) exhibit a similar evolution. We compute spherically averaged, enclosed gas mass profiles and evaluate them at 1 kpc. This value is chosen to contain several stellar half-light radii at all times, thus enclosing a volume of gas capable of significantly modify the total stellar mass content of the galaxy. We verified that the trends observed in Figure 5.2 are conserved if choosing smaller and larger radii.

We first observe the strong impact of reionization between $z = 7$ and $z \sim 4$. Gas ejected by the last supernova outflows cannot be replenished from the newly heated IGM, leading all galaxies to an order of magnitude drop in central gas content. After $z \sim 4$ however, evolutions

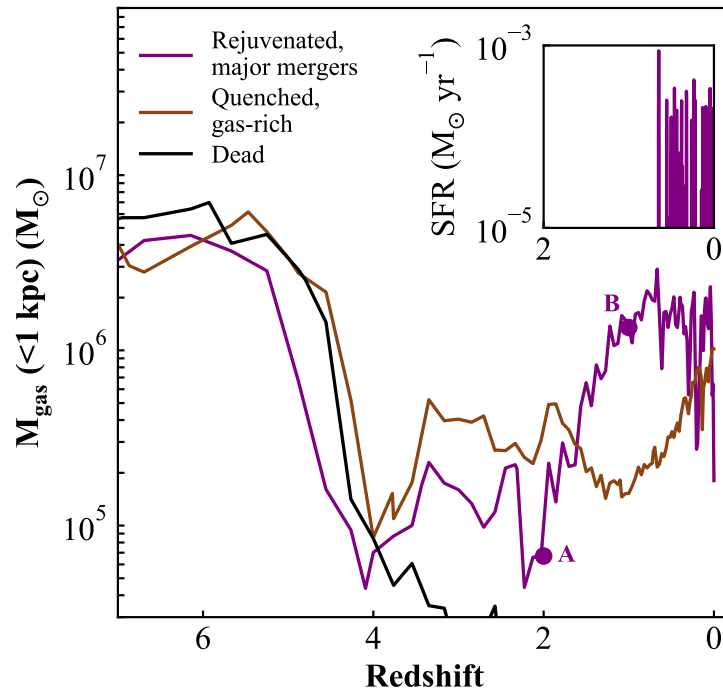


Fig. 5.2.: Evolution of the enclosed gas mass within 1 kpc of three ultra-faint dwarf galaxies with increasing final dynamical mass. Cosmic reionization drastically reduces the gas content of all galaxies by $z \sim 4$, but the two larger objects (purple, brown) stabilise more gas in their centre than their lower mass counterpart (black). By $z \sim 2$, the gas content of the most massive galaxy (purple) steadily increases, eventually leading to the re-ignition of star formation at late-times (inset panel). This increase is best explained by the cooling and collapsing of gas in the halo centre (Figure 5.3). By contrast, the intermediate mass system (brown) shows an increase in gas content but delayed in time. We argue that this system is undergoing, but has not completed, rejuvenation (Figure 5.4).

start diverging. The more massive galaxies² (brown and purple) are able to retain and stabilize significantly more gas in their centre compared to the lowest mass galaxy (black). Furthermore, starting from $z \sim 2$, the most massive example (purple) sees its central gas content steadily increase, eventually leading to the rejuvenation of sustained star formation (inset panel).

The intermediate mass galaxy (brown) presents a similar steady growth in central gas mass, although delayed in time. This galaxy never reaches as high gas masses and further does not exhibit late time star formation. We thus hypothesise that this system is undergoing, but has not completed, rejuvenation. We will see in Section 5.5 that an alteration of its mass growth increasing its dynamical mass at all times successfully reignites star formation, confirming this hypothesis.

To gain insights in the physical mechanisms driving the increase in gas content, we select two snapshots at the beginning and end of the process (respectively A and B in Figure 5.2) and extract the corresponding gas phase diagram (Figure 5.3). We compute two-dimensional histograms of all gas contained in the halo in temperature and density logarithmic bins and colour each bin by the number of contained gas cells.

At the start of the transition ($z = 2$, top panel), most gas is distributed along a diagonal rising from low densities and with temperature between 10^4 and 10^5 K. This equilibrium feature lines along an adiabat ($T \propto \rho^{\gamma-1}$ and $\gamma \sim 1.5$), therefore highlighting gas condensing slowly with long cooling times. Such adiabatic evolution is unlikely to lead to the high densities and low temperatures required to ignite star formation. A second feature is however apparent, horizontally spreading towards higher densities. Once sufficiently high temperatures and densities are reached, radiative cooling becomes efficient. The gas then follows a new equilibrium set by the balance between cooling and photo-heating, leading to near-isothermal evolution (Theuns et al. 1998). Isothermal condensation is of particular interest as it leads to large gas densities and thus potentially to rejuvenation (Ricotti 2009).

When reaching the end of the transition ($z = 1$, bottom panel), most of the gas is distributed along the equilibrium between heating and radiative cooling and has evolved to higher densities. Gas within the halo is nearly isothermal, though regions with densities above $\sim 10^{-2} m_p \text{ cm}^{-3}$ are starting to cool below 10^4 K. This is a natural consequence of our self-shielding prescription (Section 5.2), which reduces the impact of photo-heating above this density threshold. The rejuvenation process is hence linked to a near-isothermal collapse of gas along the equilibrium between photo-heating and radiative cooling, until it reaches self-shielding densities. At this

²We emphasize again that our chosen range of final dynamical masses is narrow, with high-mass galaxies only being twice more massive than their low-mass counterparts.

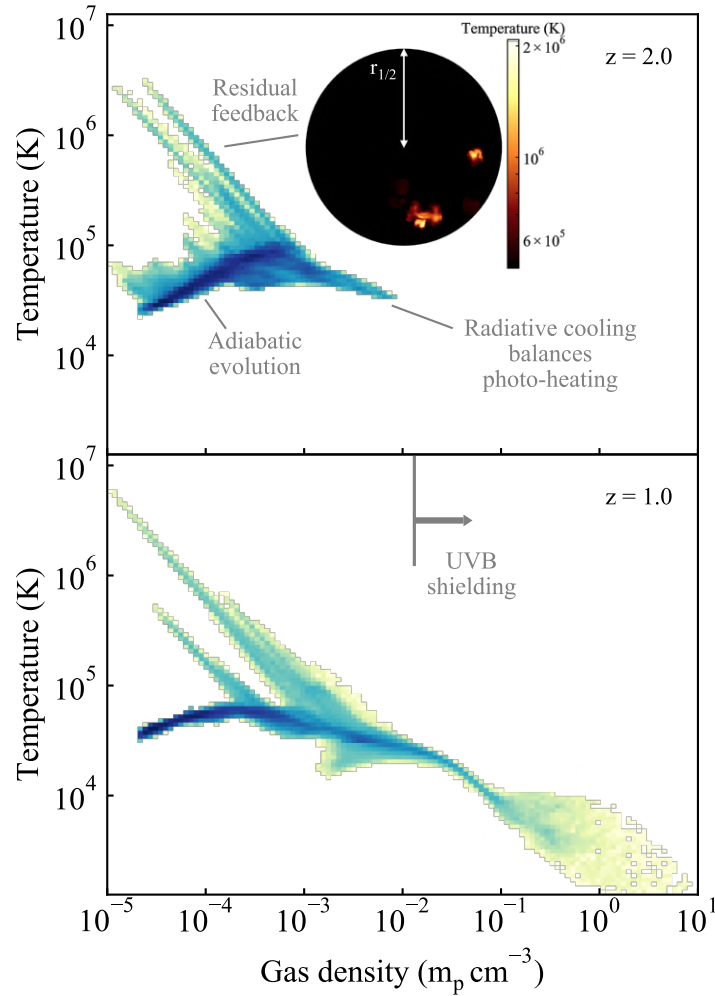


Fig. 5.3: Evolution of gas state during the rejuvenation process. Most gas evolves adiabatically at the start of the transition (top panel), hence unable to condense to high densities. Part of the gas has however already started to cool, evolving along the equilibrium between radiative cooling and photo-heating. This nearly isothermal evolution leads to a strong increase in gas densities at later times (bottom panel), eventually reaching sufficient densities to self-shield. At this point, gas can cool below 10^4 K, which will lead to the re-ignition of star formation shortly after. Plumes of high-temperature gas are linked to individual feedback events within the half-light radius (top, inset panel).

point, the impact of the UV background is drastically reduced hence creating the cold and dense gas required to fuel star formation. There are nonetheless two key open question left unanswered by this scenario:

1. Rejuvenation in our galaxies spreads across several Gyr of cosmic time, while the dynamical and cooling timescales of self-shielded gas are short (~ 100 Myr). However, both panels of Figure 5.3 exhibits discrete plumes of high temperature gas. The inset panel shows the spatial distribution of hot gas (above 5×10^5 K) in a sphere centered on the galaxy and of one half-light radius. Each plume spatially coincides with individual bubbles of hot gas expanding in the interstellar medium, indicating that a residual, stochastic heating source remains active long after quenching by reionization. Being active for billions of years after the last birth of a stellar population (Maoz et al. 2012), Type Ia supernovae could provide such a feedback mechanism and potentially delay rejuvenation. We are currently conducting a detailed exploration of the importance of Type Ia supernovae in the rejuvenation process, for example increasing or reducing the energy input from individual explosions.
2. Condensation, even adiabatic, is only possible if gas is Jeans' unstable, i.e. if the potential well is deep enough to overcome thermal pressure. We hypothesised in Section 5.3 that the rejuvenation transition is triggered by the galaxy's dynamical mass growing sufficiently to meet the previous condition. The existence of such "threshold" in dynamical mass would make rejuvenation highly sensitive to the specific mass growth history of a galaxy, which we now set out to verify.

5.5 Rejuvenation and mass assemblies

We use our ability to generate alternative mass growth histories for a galaxy to explore the response of rejuvenation to modifications in the build-up of mass.

Delaying rejuvenation We create a new initial condition for the "Rejuvenated, major mergers" galaxy, engineered to delay its rapid mass growth due to major mergers around $z \sim 3$. Smoothing the mass accretion can be achieved by reducing the local variance in the Lagrangian region (Chapter 2 and 3). We match our filtering scale to the dynamical mass of the major progenitor around the time of the mergers (i.e. $\sim 1 \times 10^9 M_{\odot}$) and decrease the variance by 20%. By maintaining the same mean density across the Lagrangian region, we hope to conserve the final halo mass. We show the original (solid purple) and modified (dashed purple) mass growth and corresponding star formation histories in Figure 5.4.

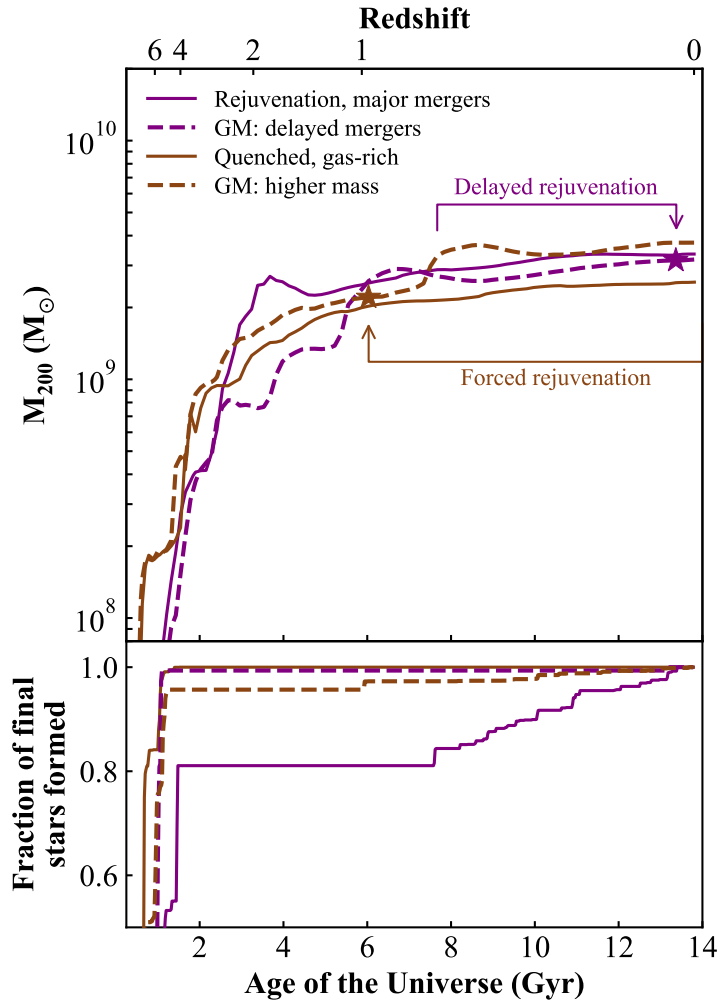


Fig. 5.4.: Growth in dynamical mass (top panel) of reference (solid) ultra-faints and their genetically modified counterparts (dashed). By delaying the build-up of mass in the largest system (purple), we delay the rejuvenation of star formation (bottom) and drastically reduce the fraction of stars formed at late-time from 19% to 0.7%. Alternatively, increasing the dynamical mass at late-times of our intermediate object (brown) triggers rejuvenation at $z \sim 1$ and continuous star formation until today. These two experiments showcase the strong interaction between a galaxy’s mass growth and its ability to rejuvenate, in turn predicting a diversity in stellar populations and gas content at $z = 0$.

Our modifications successfully smooth the late mass accretion history while conserving the early mass growth and converging to the same dynamical mass at $z = 0$ within 5%. The reference galaxy grows from $4 \times 10^8 M_{\odot}$ to $2 \times 10^9 M_{\odot}$ in ~ 1.5 Gyr starting around $z = 3$, while the modified counterpart undergoes the same growth in ~ 4 Gyr. This delay has a radical impact on the cumulative star formation history (bottom panel): 0.7% of the stars are now formed after $z = 4$ and rejuvenation of star formation is delayed from $z = 0.67$ to $z = 0.04$ (arrow) compared to the reference case.

Forcing rejuvenation We further construct a modified initial condition for the “Quenched, gas-rich” ultra-faint. We established in Section 5.4 that this galaxy is likely on the verge of rejuvenation, having started to accrete gas in its centre. We therefore wish to increase its halo mass at late times, hoping to trigger its rejuvenation. The final halo mass is best controlled by the mean density across the Lagrangian region (Roth et al. 2016), which we increase by 20%. The modified mass growth (dashed brown in Figure 5.4) is consistently higher in halo mass at late-times compared to the reference (solid), reaching $3.7 \times 10^9 M_{\odot}$ at $z = 0$ compared to $2.5 \times 10^9 M_{\odot}$. This increased dynamical mass leads to the rejuvenation of star formation (bottom panel), starting at $z = 0.96$ and actively forming 5% of the total stellar mass until $z = 0$.

These two experiments demonstrate the strong interaction between the mass growth of a galaxy and its ability to form stars at late-times. Delaying the build-up of mass of our high-mass galaxy delays the rejuvenation process, nearly preventing it at all. By contrast, increasing the dynamical mass of an otherwise quenched ultra-faint leads to its early rejuvenation and continuous star formation until today. Given the wide variety of possible histories and halo masses at $z = 0$, this mechanism naturally predicts a strong diversity in stellar populations and gas content in the overall population of ultra-faints. This diversity might start being exposed by combining HI observations with deep imaging; we leave quantifying the observable properties of our galaxies as an exciting prospect for future work.

5.6 Discussion and conclusions

In this chapter, we have shown how low-mass ultra-faint dwarf galaxies can accrete and cool gas to fuel star formation at late times. We presented four galaxies with increasing dynamical mass at $z = 0$ and simulated with high-resolution cosmological zoom simulations (Agertz et al. 2019, Chapter 4).

All four ultra-faints have their star formation quenched by reionization but the two most massive rejuvenate star formation after $z = 1$, continuously forming new stars until today at low star formation rates (Figure 5.1). These galaxies are sufficiently massive to retain gas after reionization (Figure 5.2) and later allowing it to cool as they continue to build-up dynamical mass. The gas then condenses near-isothermally in equilibrium between radiative cooling and photo-heating from the surrounding UV radiation (see also Ricotti 2009), until eventually self-shielding to reach the densities and temperatures necessary to ignite star formation (Figure 5.3).

We therefore demonstrate the ability of haloes as small as $M_{200c}(z = 0) \sim 3 \times 10^9 M_{\odot}$ (or $V_{200c} = 20 \text{ km s}^{-1}$) to cool gas and sustain star formation over extended time periods. This mass is significantly smaller than the characteristic mass for UV-suppression ($6.5 \times 10^9 M_{\odot}$ or greater) typically obtained by semi-analytical arguments (e.g. Gnedin 2000; Hoesft et al. 2006; Okamoto et al. 2008; Noh and McQuinn 2014). We argue that this difference boils down to two key reasons: first, our simulations improve the numerical resolution of previously referenced studies by between 2 and 6 orders of magnitude. This increased resolution allows us to, from very high redshift, resolve high densities in the centre of these small haloes and the collapse of surrounding filamentary structures. Both factors would decrease the efficiency of photo-ionization at preventing accretion (Hoesft et al. 2006). Secondly, the characteristic mass is typically computed as the mass-scale at which the baryon fraction of low-mass haloes is reduced by more than half of the cosmic value due to photo-ionization. Despite being “gas-rich” and star-forming in our chosen context of ultra-faint dwarfs, the presented galaxies are still extremely dark matter dominated, with baryon fractions $\leq 1\%$ and would therefore qualify as being UV-suppressed under this definition.

Condensation necessary to the rejuvenation process can only start if gas is Jeans’ unstable, thus introducing a minimum potential well depth that must be reached to trigger gas accretion. We demonstrate, using two genetically modified variants, how this effective “threshold” in dynamical mass generates a strong interaction between a galaxy’s mass growth and its ability to form stars at late-times. We clearly show that rejuvenation can be delayed or forced earlier by respectively slowing or quickening the delivery of dynamical mass (Figure 5.4). Furthermore, we highlight that galaxies overcoming this threshold late enough might have started gas accretion but not yet produced a young generation of stars.

This exposed mechanism therefore naturally predicts (i) field ultra-faint dwarf galaxies with $M_{\star} \sim 10^5 M_{\odot}$ and on-going star formation and (ii) ultra-faint dwarf galaxies with sizeable gas content but an exclusively ancient stellar population. We foresee this scenario to have important observational consequences, providing an explanation for the formation of Leo T

(Irwin et al. 2007) as well as for recent candidates of gas-rich ultra-faints obtained by matching HI observations with optical imaging (Janesh et al. 2019). Due to the vast number of possible histories and final masses, this mechanism predicts a strong diversity in the stellar populations and gas content of ultra-faints. This diversity would be further enlarged by additional rejuvenation mechanisms, such as environmental (Wright et al. 2019). The prospect of matching ever deeper imaging of e.g. LSST with HI detections strongly motivates more detailed predictions of the abundance and observational signatures of rejuvenated ultra-faints. This challenge is beyond the scope of this thesis but one we plan to tackle as future work (see Chapter 6).

Conclusions and future prospects

Building a self-consistent model for the cosmological formation of galaxies is a challenging task: the required dynamic range to model both cosmological environments and non-linear astrophysics places conflicting demands on the best use of human and computing time. Nonetheless, the need to interpret new observations (e.g. JWST, ALMA, MANGA or LSST) to constrain fundamental and astronomical physics requires new methods to study the diversity of the galaxy population while maintaining sufficient resolution to robustly model galaxy interiors.

In this thesis, we have presented, extended and applied such a new method, “genetic modifications”, first proposed in Roth et al. (2016). The key originality of this framework is its ability to create controlled experiments for galaxy simulations, similar in spirit to idealised experiments but retaining a fully cosmological context. By generating alternative initial conditions for a galaxy, we can modify specific aspects of this object’s build-up of mass while recreating a fixed cosmological environment. Because we focus on single objects, our approach naturally couples with concentrating computational power in a “zoom” region, in turn allowing us to afford more detailed modelling of the galaxy’s internal astrophysics than is available in large cosmological volume simulations. These joint abilities are the cornerstones of the “genetic modification” approach: they allow us to construct a causal account of the response of the galaxy’s observables to varying its mass accretion history (e.g. Pontzen et al. 2017; Chapter 4). We now summarize and provide a critical assessment of this thesis’ findings, outlining several improvements and applications that will be tackled in the future.

6.1 Controlling cosmological merger histories

The genetic modification framework achieves controlled experiments by interpreting and engineering the information contained in the cosmological initial conditions of a chosen galaxy. In this context however, perfect control has to be balanced with preserving the statistical properties of the density field. For example, sharp transitions in the density field would create extended power on small scales, in turn extremely unlikely in a Λ CDM Universe.

This trade-off led us to formulating “genetic modifications” as solving a constrained minimization problem in Chapter 2, i.e. finding a new density field satisfying our requested changes while making *minimal* changes in the sense of the statistical distance χ^2 to the original initial

conditions. This formulation ensures that as much as possible of the original galaxy is recreated hence creating controlled experiments. The minimal, but non-zero, nature of changes remains one of the strongest challenges associated with genetic modifications. They are necessary to produce acceptable cosmological initial conditions but naturally introduce correlated “compensation effects” to achieve this goal. An example was highlighted in Chapter 3, in which controlling merger ratios is achieved at the cost of modifying the timings of mergers and the large-scale environmental density. Such uncontrolled changes can blur the causal interpretation of the experiments, as multiple factors are modified at once.

Minimizing these compensations in an acceptable human timescale is a strong motivation to introduce new types of modifications in Chapter 2. We provided the framework for a new class of quadratic controls and implemented modifications targeting the local variance of the density field. We designed these modifications to control multiple mergers at once, hence avoiding manual tracking and tuning of individual structures inherent to linear modifications. We successfully verified this prediction in Chapter 3, enabling comparisons of scenarios in which a galaxy grows through accretion of numerous small bodies or through a smaller number of more significant events.

However, our objectives with variance modifications were only partially met. The imperfect knowledge of the mapping between initial conditions and late-time properties still implies significant exploration to understand the joint consequences of wanted and unwanted changes. Furthermore, Chapter 4 provides an example for which the scientific question addressed was more easily tackled using manual tracking and linear modifications. These limitations in the design of variance modifications are partly unavoidable: if a perfect mapping between initial conditions and late-time properties was available, one would not require numerical simulations. We argue that they in turn warrant exploring new controls and modifications to achieve refined and streamlined control and we now detail several avenues for such extensions.

Derivatives of the density field directly connect to the velocity, gravitational potential and angular momentum content of a region and could be readily tackle as they are linear in the density field. Fixing the timing of mergers could be targeted using velocity modifications to conserve the peculiar velocity structure inside the Lagrangian region. These velocity modifications could be further applied to engineer new scenarios for galaxy formation studies: one could modify the infall time and velocity of a satellite, independently of its mass, to test the triggering of bar instabilities in disk galaxies (e.g. Zana et al. 2018) or the sensitivity of quenching to the orbital trajectory in a galaxy cluster (e.g. Vijayaraghavan and Ricker 2015; Tonnesen 2019). Longer-term extensions could modify the linear tidal torque to try controlling the final spin

of haloes and galaxies (e.g. Hoyle 1949; Efstathiou and Jones 1979; White 1984; Porciani et al. 2002a; Porciani et al. 2002b). A handle on the angular momentum delivery, at fixed mass and merger history, would be key to probe the formation of disk galaxies. In addition, refined control on mergers and infall parameters could be achieved by targeting higher moments in the density field with new quadratic modifications inspired by second-order Lagrangian perturbation theory.

6.2 A new angle on dark matter halo formation

Armed with our implementation to produce genetically modified initial conditions, we turned to studying the formation of dark matter haloes in Chapter 3. We performed a series of variance-modified, dark matter-only simulations showcasing our ability to systematically vary the build-up of mass of a halo while conserving its final mass. Demonstrating this ability proved key for the subsequent applications to the formation of galaxies in Chapter 4 and 5.

Results obtained in Chapter 3 illustrate well the complementarity of genetic modification studies with the more traditional cosmological volume approach. Because computing time is less limiting to study dark matter haloes than galaxies, simulating large halo populations while resolving their structural properties has been possible for decades (see Navarro et al. 1997; Bullock et al. 2001a; Wechsler et al. 2002; Macciò et al. 2007; Ludlow et al. 2013; Diemer and Kravtsov 2015; Klypin et al. 2016; Ragagnin et al. 2019 as examples). Even with ever-increasing sample size, interpreting correlations between the high number of possible halo properties (concentration, mass, formation time, spin, etc) is challenging and we showed how a causal study of just two haloes was able to recover population-level correlations between halo concentration and formation time.

Furthermore, a key aspect of the genetic approach is its continuous nature, particularly effective to refine around unexpected “tipping points”. We showcased an example in which an apparently small change in history (fixed formation time) led to a large change in halo concentration. Generating intermediate scenarios allowed us to tie this evolution to a recombination in the order of major mergers at fixed formation time. In a context of limited human time, such subtle insights can only be gained by focussing on small numbers of closely related objects. This example also emphasizes how non-linear microphysics sometimes turns unplanned modifications into large observable changes. This complicates the causal interpretation but is also precisely where detailed physical insights can be learned.

With these successes in mind, we emphasize that genetic modification are complementary, rather than strictly improving on population studies. All results presented throughout Chapter 3 were guided by our existing knowledge of halo structural relations. In fact, genetic modification does not currently offer a prediction, even approximate, of the relation between the likelihood of a modification in the initial conditions and that of a final non-linear property (see Roth et al. 2016, section 6 for an attempt). We therefore rely on large halo samples to estimate whether our modifications are probing a representative range of the overall population. This is key to defining the magnitude of responses, i.e. what is a significant change, but also to estimating whether single object studies are likely to be general as in Chapter 4.

6.3 The diversity of ultra-faint dwarf galaxies

In Chapter 4 and 5, we finally applied our approach to the formation of ultra-faint dwarf galaxies. Their low mass likely makes these galaxies acutely sensitive to both galaxy formation physics and their cosmological build-up of mass. By combining genetic modifications with high-resolution zoom simulations, we can vary each factor in turn thus providing a more detailed account of how their coupling shapes observable properties.

In Chapter 4, we created four alternative histories for a single ultra-faint galaxy varying its early growth in dynamical mass. We showed that, because cosmic reionization shuts down internal star formation at a fixed time, objects which began forming stars earlier assemble higher final stellar masses (Section 4.3). Since all galaxies have the same final dynamical mass by construction, we therefore cleanly demonstrated that the one-to-one correspondence between halo mass and stellar mass assumed by abundance matching (e.g. Behroozi et al. 2013) breaks down on the scale of ultra-faint dwarfs. This uncovered diversity in stellar mass has important consequences for modelling the galaxy-halo connection at low-masses, and we plan to explore its impact in constraining the nature of dark matter through the abundance of satellites (e.g. Nadler et al. 2019).

Scanning through multiple histories at fixed mass further exposed the potential for a new class of highly diffuse ultra-faints (Section 4.5). Such galaxies arise through an early truncation of internal star formation due to reionization but the later growth of stellar mass through mergers. Mergers deposit stars on extended orbits, thus vastly increasing the galaxy's stellar half-light radius. We illustrated that the latest generation of instruments is just reaching these levels of low surface brightness, in particular detections through resolved stars in the Local Group. This prediction demonstrates that the census of nearby dwarf galaxies is unlikely to be

complete but that such galaxies should be discovered in the near future by deeper photometric surveys such as LSST.

We pursued another source of diversity in Chapter 5 by using a larger suite of objects with increasing dynamical mass today. We showed that the two most massive systems rejuvenate star formation after $z = 1$. We tied this bimodality of star formation with a late growth in dynamical mass overcoming a “threshold”, triggering the near-isothermal condensation of gas in the halo centre. Several questions were left open in Chapter 5 and we plan to tackle them in the near-future, in particular the importance of residual heating from Type Ia supernovae and the sensitivity of our results to variations in the cooling and heating implementations.

We explicitly demonstrated the coupling between rejuvenation and the build-up of mass of ultra-faints. We used genetic modifications to delay and force rejuvenation in two distinct objects by respectively slowing and quickening their growth of mass (Section 5.5). Because rejuvenation spreads across billions of years, galaxies might begin gas accretion but have insufficient time to rejuvenate before $z = 0$. Such a generic mechanism has long-reaching consequences, as it naturally predicts field ultra-faints with (i) ongoing star formation and (ii) reionization-old stars but a sizeable gas content. This prediction could provide an explanation of the long-standing challenge posed by the star-forming ultra-faint galaxy, Leo T (Irwin et al. 2007), as well as recent detections of gas-rich ultra-faints (Janesh et al. 2019) obtained by matching HI detections with optical imaging. Quantifying the observable properties and abundances of these objects is therefore of paramount importance, and we plan to soon link our results with ongoing and future HI and optical surveys.

In a broader context, Chapters 4 and 5 are part of the collaborative EDGE¹ project introduced in Agertz et al. (2019). In this work, we presented an exploration of the theoretical uncertainties associated with modeling galaxy formation physics at the scale of ultra-faints and the constraining power of stellar metallicities to differentiate models (Ageretz et al. 2019). In addition to questions linked to Chapter 5 which I will be exploring personally, we will tackle additional science cases in the near future, e.g. whether low-mass dwarf galaxies are capable of transforming dark matter cusps into cores during their early star-forming phase (Pontzen and Governato 2012); whether these cores can survive today following mergers with dark cuspy haloes (Laporte and Penarrubia 2015); and whether they can be observed by the dynamical properties of globular clusters (Amorisco 2017).

¹Engineering Dwarfs at Galaxy Formation’s Edge

6.4 Future work

We have seen throughout this thesis how studying a small number of high-resolution, genetically modified galaxies provides a powerful new angle on the complexity of galaxy formation. The ability to modify a galaxy's history, while maintaining its final dynamical mass and environment, provides a nearly endless list of possible applications. In addition to previously highlighted future work arising directly from this manuscript, we now focus on novel ideas for the longer term future.

A unique window on the galaxy assembly process is provided by the diffuse stellar halo surrounding galaxies, originating from the accretion and stripping of stars previously formed in smaller galaxies. Because such stars in the outskirts of galaxies are collisionless and have long orbital timescales, memory of past mergers is likely retained in their kinematics. Stellar haloes are faint, extended, low surface brightness objects, making them challenging to tackle observationally, but this picture is now rapidly evolving. Data from the GAIA sky survey provides an unprecedented view of the Milky Way stellar halo, including resolved kinematics and photometry for a large fraction of halo stars. Furthermore, stellar haloes are now targeted in extragalactic surveys, with dedicated time programs (GHOSTS; Monachesi et al. 2016), new instruments (Dragonfly; Merritt et al. 2016) or are complementary science programs for cosmological cameras (DES, HSC and LSST in the future). These instruments discovered an unexpected diversity in the stellar halo population, with Milky-Way-like galaxies showing both non-detection (e.g. M101) or extended stellar haloes (e.g. M31). This diversity seems to be generated by the variety of possible merger histories for a given galaxy (Elias et al. 2018; Monachesi et al. 2019), as well as through the response of the central galaxy to these mergers (Zolotov et al. 2009; Cooper et al. 2015). Simulating a handful of modified, high-resolution galaxies would help disentangle how merger histories build a stellar halo, as well as clarifying the observational imprints of these mergers in the hope of reconstructing a galaxy's past (e.g. Belokurov et al. 2018; Helmi et al. 2018).

A broader effort to apply genetic modifications from dwarf galaxies to clusters is also underway. Results and expertise derived in this thesis will form a base for these collaborative studies. In addition to focusing on different scientific questions, these projects also differ in the numerical methods and overall project scale, being part of a larger effort in both computing and human time. In particular, we briefly highlight applications to (i) the formation of recently discovered ultra-diffuse dwarf galaxies (e.g. van Dokkum et al. 2015), in which we will engineer initial conditions to test if the extended size of these galaxies could be the result of specific

merger scenarios (e.g. Di Cintio et al. 2019) and (ii) the co-evolution of super-massive black holes and galaxies, in which we will create alternative merger histories to investigate how this co-evolution drives quenching of star formation, galaxy morphological transformations (e.g. Pontzen et al. 2017) and the emission of gravitational waves.

Much remains to be established in constructing an *ab initio* model of galaxy formation. Our incomplete understanding of this process hinders our ability to answer some of the most fundamental questions of modern physics: *what is our Universe made of* and *where do we come from*. We hope the new approach presented in this thesis and its numerous applications across the galaxy formation field will prove influential in this quest.

Appendix: Supplementary interpretations of quadratically modified fields

A.1 Constructing constrained ensembles for quadratic constraints

In Section 2.2, we contrasted the notion of a linearly-constrained ensemble (Section 2.2.1) against that of genetic modifications (Section 2.2.2). While conceptually different, constrained ensembles can be sampled using the HR91 procedure which can in turn be seen as applying suitable modifications to realisations from the unconstrained probability distribution.

In this section we show that there is no such similarity between quadratically-constrained ensembles and quadratically modified fields. To put it another way, there is no HR91-like method for generating samples from a quadratically-constrained ensemble.

Following the same Bayesian argument as in Section 2.2.1, we define the quadratically-constrained ensemble for a fixed \mathbf{Q} and q by

$$P(\boldsymbol{\delta}|q) \propto \exp\left(-\frac{1}{2}\boldsymbol{\delta}^\dagger \mathbf{C}_0^{-1}\boldsymbol{\delta}\right) \delta_D(\boldsymbol{\delta}^\dagger \mathbf{Q} \boldsymbol{\delta} - q). \quad (\text{A.1})$$

We will show that the modification procedure does not generate samples from the ensemble (A.1), even when \mathbf{Q} and q are known and fixed in advance.

We start by defining the alternative ensemble $P(\boldsymbol{\delta}_1|q)$ to be that sampled by drawing an unconstrained field from $P(\boldsymbol{\delta})$ and using the GM procedure to enforce the constraint $\boldsymbol{\delta}^\dagger \mathbf{Q} \boldsymbol{\delta} = q$. In Section 2.3.1, the mapping $\boldsymbol{\delta}_0 \rightarrow \boldsymbol{\delta}_1$ was given by

$$\boldsymbol{\delta}_1 = \exp(\alpha(\boldsymbol{\delta}_0, q) \mathbf{C}_0 \mathbf{Q}) \boldsymbol{\delta}_0, \quad (\text{A.2})$$

where \mathbf{C}_0 is the covariance matrix of the Gaussian distribution $P(\boldsymbol{\delta})$ and the value of $\alpha(\boldsymbol{\delta}_0, q)$ is implicitly defined by the need to satisfy the quadratic constraint $\boldsymbol{\delta}_1^\dagger \mathbf{Q} \boldsymbol{\delta}_1 = q$.

To incorporate this implicit requirement to choose the correct value of α into an expression for the ensemble, we make use of Bayes' theorem:

$$\begin{aligned}
P(\boldsymbol{\delta}_1|q) &= \int P(\boldsymbol{\delta}_1|\alpha)P(\alpha|q) d\alpha \\
&= \iint P(\boldsymbol{\delta}_1|\alpha, \boldsymbol{\delta}_0)P(\alpha|q, \boldsymbol{\delta}_0)P(\boldsymbol{\delta}_0) d\alpha d\boldsymbol{\delta}_0 \\
&= \iint P(\boldsymbol{\delta}_1|\alpha, \boldsymbol{\delta}_0)P(q|\alpha, \boldsymbol{\delta}_0)\frac{P(\alpha|\boldsymbol{\delta}_0)}{P(q)}P(\boldsymbol{\delta}_0) d\alpha d\boldsymbol{\delta}_0.
\end{aligned} \tag{A.3}$$

Note that the constraint demands $P(q|\alpha, \boldsymbol{\delta}_0) = \delta_D(\boldsymbol{\delta}_1^\dagger \mathbf{Q} \boldsymbol{\delta}_1 - q)$, where $\boldsymbol{\delta}_1$ and $\boldsymbol{\delta}_0$ are related by the condition (A.2). Writing out the normalisation condition for $P(\alpha|q, \boldsymbol{\delta})$ then gives

$$1 = \int P(\alpha|q, \boldsymbol{\delta}) d\alpha = \int \delta_D(\boldsymbol{\delta}_1^\dagger \mathbf{Q} \boldsymbol{\delta}_1 - q) \frac{P(\alpha|\boldsymbol{\delta}_0)}{P(q)} d\alpha. \tag{A.4}$$

Because \mathbf{Q} and $\mathbf{C}_0 \mathbf{Q}$ are positive semi-definite, q is a monotonically increasing function of α ; there is only one value of α which satisfies the Dirac delta function on the right-hand-side. Consequently, we can perform the integration by a change of variables to yield

$$\left. \frac{P(\alpha|\boldsymbol{\delta}_0)}{P(q)} \right|_{\boldsymbol{\delta}_1^\dagger \mathbf{Q} \boldsymbol{\delta}_1 = q} = \left. \frac{\partial}{\partial \alpha} \right|_{\boldsymbol{\delta}_0} \left(\boldsymbol{\delta}_1^\dagger \mathbf{Q} \boldsymbol{\delta}_1 \right) = 2\boldsymbol{\delta}_1^\dagger \mathbf{Q} \mathbf{C}_0 \mathbf{Q} \boldsymbol{\delta}_1. \tag{A.5}$$

Substituting this result back into Equation (A.3) and performing the integral over $\boldsymbol{\delta}_0$ using $P(\boldsymbol{\delta}_1|\alpha, \boldsymbol{\delta}_0) = \delta_D(\boldsymbol{\delta}_1 - \exp(\alpha \mathbf{C}_0 \mathbf{Q})\boldsymbol{\delta}_0)$, one obtains

$$\begin{aligned}
P(\boldsymbol{\delta}_1|q) &\propto \delta_D(q - \boldsymbol{\delta}_1^\dagger \mathbf{Q} \boldsymbol{\delta}_1) \boldsymbol{\delta}_1^\dagger \mathbf{Q} \mathbf{C}_0 \mathbf{Q} \boldsymbol{\delta}_1 \\
&\quad \times \int d\alpha |e^{-\alpha \mathbf{C}_0 \mathbf{Q}}| \exp\left(-\frac{1}{2} \boldsymbol{\delta}_1^\dagger e^{-\alpha \mathbf{Q} \mathbf{C}_0} \mathbf{C}_0^{-1} e^{-\alpha \mathbf{C}_0 \mathbf{Q}} \boldsymbol{\delta}_1\right), \tag{A.6}
\end{aligned}$$

where normalisation factors depending only on \mathbf{C}_0 have been dropped. This expression no longer has any explicit reference to $\boldsymbol{\delta}_0$, which was our primary aim. It can now be compared with the distribution for a true constrained ensemble, Equation (A.1). The two distributions appear different (as expected given our claim of inequivalence), but the comparison is complicated by the unsolved integral over α which obscures the content of the expression.

We can see that this integral will never regenerate the true quadratic constrained ensemble by taking an illustrative example. Consider a three-dimensional field $\delta_1 = (x, y, z)$ with unit power spectrum ($C_0 = I$). Let us further choose an explicit form for Q such that

$$Q = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \Rightarrow e^{\alpha C_0 Q} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{-\alpha} & 0 \\ 0 & 0 & e^{\alpha} \end{pmatrix}. \quad (\text{A.7})$$

Inserting these results into Equation (A.6) gives

$$P(\delta_1|q) \propto \delta_D(q + y^2 - z^2)(y^2 + z^2) \times \int_0^\infty \frac{d\beta}{\beta} \exp\left(-\frac{1}{2}(x^2 + \beta^{-2}y^2 + \beta^2z^2)\right), \quad (\text{A.8})$$

where we have made the substitution $\beta = e^{-\alpha}$. The integral over β has an analytical solution using the further substitution $t = (\beta z)^2/2$ and introducing the modified Bessel function of the second kind

$$K_0(x) = \frac{1}{2} \int_0^\infty \frac{dt}{t} e^{-t - \frac{x^2}{4t}}. \quad (\text{A.9})$$

Equation (A.6) can then be evaluated explicitly to obtain

$$P(\delta_1|q) \propto e^{-\frac{x^2}{2}} \delta_D(q + y^2 - z^2) (y^2 + z^2) K_0(|yz|). \quad (\text{A.10})$$

For comparison, the quadratic constrained ensemble in this example is given by

$$P(\delta|q) \propto e^{-\frac{x^2}{2}} \delta_D(q + y^2 - z^2) e^{-\frac{y^2+z^2}{2}}. \quad (\text{A.11})$$

The distributions defined by (A.10) and (A.11) have identical x -dependence. This is a general property: degrees of freedom for which Q has a null direction are unconstrained and, similarly, left unchanged by our GM transformation. The distribution generated by these degrees of freedom will therefore coincide at all times with the constrained ensemble. However, the y and z dependences differ between Equations (A.10) and (A.11). In general, non-null directions in field space will behave differently between the GM and constrained ensemble cases.

The result establishes that our formulation of quadratic GM as a matrix exponential mapping does not reproduce a quadratically-constrained ensemble when used analogously to the HR91 algorithm. A similar argument allows one to verify that applying the alternative non-linear modification specified by Equation (2.14) also fails to regenerate the constrained result. In fact,

one can go even further and write a general power series expansion for the mapping between δ_0 and δ_1 , writing

$$\delta_1 = \sum_{i=0}^{\infty} A_i (\mu \mathbf{C}_0 \mathbf{Q})^i \delta_0, \quad (\text{A.12})$$

without further specifying the power series coefficients A_i . Even in this case, which generalises away from a specific mapping, it is not possible to generate a constrained ensemble from the modification procedure. This underlines that modifications and constraints need to be regarded as entirely separate procedures. Only in the linear case do they appear to be cosmetically related.

A.2 Geometrical interpretation

Throughout the main text, we used fields sampled at a finite number of points n ; the resulting algorithms can therefore be interpreted geometrically as acting on vectors in an n -dimensional space. For instance, Roth et al. (2016) noted that the linear GM procedure is equivalent to an orthonormal projection of the unmodified field onto a subspace defining the modification objective (see their appendix A). In this Section we provide the geometric interpretation for our extended formulation of GM.

For the purposes of visualising the connection, we use fields with only two samples, $\delta = (x, y)$. The arguments of this section generalise to higher dimensions but are easiest to visualise with $n = 2$. Figure A.1 shows the resulting two-dimensional geometry in terms of the displacements Δx and Δy from the unmodified field. By construction, the unmodified field is at the origin.

The left panel shows the elliptical geometry generated by the covariance matrix in the $\Delta x - \Delta y$ plane; specifically, the ellipses are of constant distance Δs^2 from the origin, where

$$\Delta s^2 \equiv \begin{pmatrix} \Delta x & \Delta y \end{pmatrix} \mathbf{C}_0^{-1} \begin{pmatrix} \Delta x \\ \Delta y \end{pmatrix}. \quad (\text{A.13})$$

The linear objective $\mathbf{A}\delta = \mathbf{b}$ defines a line in two dimensions. The modification consists of finding the value of $(\Delta x, \Delta y)$ lying on the line, while minimising Δs^2 . Since Δs^2 is measured in terms of \mathbf{C}_0^{-1} , the solution does not correspond to the closest point on the page but to the point at which a covariance ellipse is tangent to the modification line.

Similarly, the quadratic modifications (right panel of Figure A.1) are associated with ellipses of constant $q = (x, y) \mathbf{Q} (x, y)^\top$. These targets are shown as dotted lines; note that they are centred on $(x, y) = (0, 0)$ and therefore appear offset from the origin in the $\Delta x - \Delta y$ plane.

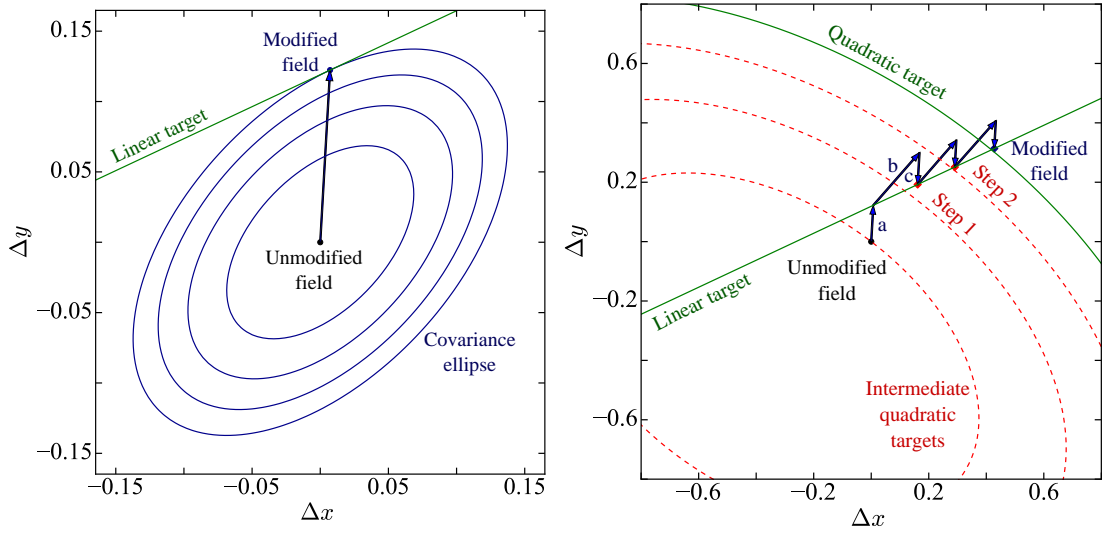


Fig. A.1.: *Left panel:* Geometry of linear GM for a field with two components $\delta = (x, y)^\top$. The axes represent displacements $\Delta\delta = (\Delta x, \Delta y)^\top$ from the unmodified realisation. The distance measure, Equation (A.13), gives rise to elliptical surfaces of constant distance (blue). The linear target corresponds to a line (green). The GM algorithm (arrow) takes the unmodified realisation (black dot) to the first intersection between this line and ellipses of increasing distance, defining the modified field. *Right panel:* Geometry of making simultaneous quadratic and linear modifications using the algorithm from Section 2.3.1. Two target modifications are shown, a linear (green line) and a quadratic (green ellipse). The algorithm defines intermediate quadratic modifications (red dotted ellipses) to step towards the final result. The first operation is the projection of the unmodified field onto the linear modification (a); each iterative step then displaces the field along the normal of the ellipse (b), and projects it again onto the linear modification (c).

The right panel of Figure A.1 also illustrates the algorithm for finding the modified field with a simultaneous quadratic and linear objective. For visual clarity, an unrealistically small ($N = 3$) number of steps are taken. We start by defining three intermediate ellipses (red-dotted) between the value of the modification at the unmodified field and the target. As explained in Section 2.3.2, we first apply the global linear modifications from Equation (2.8)

$$\delta \rightarrow \delta - \underbrace{C_0 A^\dagger (A C_0 A^\dagger)^{-1} (A \delta - b)}_a. \quad (\text{A.14})$$

The algorithm then iterates the step ϵ defined by Equation (2.24)

$$\epsilon = - \underbrace{\mu C_0 Q \delta}_b + \underbrace{\mu C_0 A^\dagger (A C_0 A^\dagger)^{-1} A C_0 Q \delta}_c. \quad (\text{A.15})$$

These operations can be understood geometrically as:

- (a) A projection of the current field on the linear modification. This term is similar to the case with linear modifications only.
- (b) A displacement along the normal of the ellipse at the current field value. This term is towards the next intermediate ellipse.

(c) The projection of the previous term back onto the linear modification to ensure that both are always satisfied.

Term (c) ensures that the current field at the end of each step always lies on the linear constraint. Term (a) therefore vanishes after the first step; it is an overall offset that needs to be applied only once. Together, (b) and (c) are locally orthogonalizing the quadratic modification with respect to the global linear modification. The orthogonalisation must be repeated at each step since the local linearisation changes as we progress towards the final value of q .

Bibliography

- Adamek, Julian et al. (July 2016). ‘gevolution: a cosmological N-body code based on General Relativity’. In: *J. Cosmology Astropart. Phys.* 2016.7, 053, p. 053. arXiv: 1604.06065 [astro-ph.CO] (cit. on p. 27).
- Adams, Elizabeth A. K. et al. (Apr. 2018). ‘Deep neutral hydrogen observations of Leo T with the Westerbork Synthesis Radio Telescope’. In: *A&A* 612, A26, A26. arXiv: 1712.06636 [astro-ph.GA] (cit. on p. 95).
- Agertz, O. et al. (Sept. 2007). ‘Fundamental differences between SPH and grid methods’. In: *MNRAS* 380.3, pp. 963–978. arXiv: astro-ph/0610051 [astro-ph] (cit. on p. 28).
- Agertz, O. et al. (June 2013). ‘Toward a Complete Accounting of Energy and Momentum from Stellar Feedback in Galaxy Formation Simulations’. In: *ApJ* 770.1, 25, p. 25. arXiv: 1210.4957 [astro-ph.CO] (cit. on pp. 84, 85, 98).
- Agertz, O. et al. (May 2015). ‘On the Interplay between Star Formation and Feedback in Galaxy Formation Simulations’. In: *ApJ* 804.1, 18, p. 18. arXiv: 1404.2613 [astro-ph.GA] (cit. on p. 30).
- Agertz, O. et al. (Apr. 2019). ‘EDGE I: the mass-metallicity relation as a critical test of galaxy formation physics’. In: *arXiv e-prints*, arXiv:1904.02723, arXiv:1904.02723. arXiv: 1904.02723 [astro-ph.GA] (cit. on pp. 3, 30, 37, 81–86, 92, 97, 99, 106, 113).
- Alam, Shadab et al. (Sept. 2017). ‘The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample’. In: *Monthly Notices of the Royal Astronomical Society* 470.3, pp. 2617–2652. arXiv: 1607.03155 [astro-ph.CO] (cit. on p. 22).
- Albrecht, Andreas et al. (Sept. 2006). ‘Report of the Dark Energy Task Force’. In: *arXiv e-prints*, astro-ph/0609591, astro-ph/0609591. arXiv: astro-ph/0609591 [astro-ph] (cit. on p. 22).
- Amorisco, Nicola C. (July 2017). ‘Deadly Dark Matter Cusps versus Faint and Extended Star Clusters: Eridanus II and Andromeda XXV’. In: *ApJ* 844.1, 64, p. 64. arXiv: 1704.06262 [astro-ph.GA] (cit. on p. 113).
- Arzoumanian, Zaven et al. (Apr. 2018). ‘The NANOGrav 11-year Data Set: High-precision Timing of 45 Millisecond Pulsars’. In: *The Astrophysical Journal Supplement Series* 235.2, 37, p. 37. arXiv: 1801.01837 [astro-ph.HE] (cit. on p. 26).
- Aubert, Dominique et al. (Nov. 2010). ‘Reionization Simulations Powered by Graphics Processing Units. I. On the Structure of the Ultraviolet Radiation Field’. In: *ApJ* 724.1, pp. 244–266. arXiv: 1004.2503 [astro-ph.CO] (cit. on p. 83).
- Avila-Reese, V. et al. (Nov. 2005). ‘The Dependence on Environment of Cold Dark Matter Halo Properties’. In: *ApJ* 634, pp. 51–69. eprint: astro-ph/0508053 (cit. on pp. 60, 75).
- Bacon, R. et al. (Mar. 2015). ‘The MUSE 3D view of the Hubble Deep Field South’. In: *A&A* 575, A75, A75. arXiv: 1411.7667 (cit. on p. 25).

- Bagnasco, Giorgio et al. (Sept. 2007). ‘Overview of the near-infrared spectrograph (NIRSpec) instrument on-board the James Webb Space Telescope (JWST)’. In: *Proc. SPIE*. Vol. 6692. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 66920M (cit. on p. 25).
- Balogh, Michael L. et al. (Oct. 2001). ‘Revisiting the cosmic cooling crisis’. In: *MNRAS* 326.4, pp. 1228–1234. arXiv: astro-ph/0104041 [astro-ph] (cit. on p. 29).
- Bañados, Eduardo et al. (Jan. 2018). ‘An 800-million-solar-mass black hole in a significantly neutral Universe at a redshift of 7.5’. In: *Nature* 553.7689, pp. 473–476. arXiv: 1712.01860 [astro-ph.GA] (cit. on p. 26).
- Barcons, X. et al. (Mar. 2017). ‘Athena: ESA’s X-ray observatory for the late 2020s’. In: *Astronomische Nachrichten* 338.153, pp. 153–158 (cit. on p. 26).
- Bardeen, J. M. et al. (May 1986). ‘The statistics of peaks of Gaussian random fields’. In: *ApJ* 304, pp. 15–61 (cit. on pp. 32, 40, 61).
- Barnes, Josh et al. (Dec. 1986). ‘A hierarchical $O(N \log N)$ force-calculation algorithm’. In: *Nature* 324.6096, pp. 446–449 (cit. on p. 27).
- Bastian, Nate et al. (Sept. 2010). ‘A Universal Stellar Initial Mass Function? A Critical Look at Variations’. In: *ARA&A* 48, pp. 339–389. arXiv: 1001.2965 [astro-ph.GA] (cit. on p. 96).
- Behroozi, Peter S. et al. (June 2013). ‘The Average Star Formation Histories of Galaxies in Dark Matter Halos from $z = 0-8$ ’. In: *ApJ* 770.1, 57, p. 57. arXiv: 1207.6105 [astro-ph.CO] (cit. on pp. 65, 88, 112).
- Belokurov, V. et al. (July 2018). ‘Co-formation of the disc and the stellar halo’. In: *MNRAS* 478.1, pp. 611–619. arXiv: 1802.03414 [astro-ph.GA] (cit. on p. 114).
- Benitez-Llambay, A. et al. (July 2015). ‘The imprint of reionization on the star formation histories of dwarf galaxies’. In: *MNRAS* 450.4, pp. 4207–4220. arXiv: 1405.5540 [astro-ph.GA] (cit. on p. 82).
- Benson, A. J. (Feb. 2012). ‘G ALACTICUS: A semi-analytic model of galaxy formation’. In: *New Astron.* 17, pp. 175–197. arXiv: 1008.1786 (cit. on p. 77).
- Bernal, José Luis et al. (Oct. 2016). ‘The trouble with H_0 ’. In: *J. Cosmology Astropart. Phys.* 2016.10, 019, p. 019. arXiv: 1607.05617 [astro-ph.CO] (cit. on p. 22).
- Bernardeau, F. et al. (Sept. 2002). ‘Large-scale structure of the Universe and cosmological perturbation theory’. In: *Phys. Rep.* 367.1-3, pp. 1–248. arXiv: astro-ph/0112551 [astro-ph] (cit. on pp. 19, 32).
- Bertone, Gianfranco et al. (Oct. 2018). ‘A new era in the search for dark matter’. In: *Nature* 562.7725, pp. 51–56. arXiv: 1810.01668 [astro-ph.CO] (cit. on p. 22).
- Bertschinger, E. (Dec. 1987). ‘Path integral methods for primordial density perturbations - Sampling of constrained Gaussian random fields’. In: *ApJ* 323, pp. L103–L106 (cit. on pp. 33, 40, 43).
- (Nov. 2001). ‘Multiscale Gaussian Random Fields and Their Application to Cosmological Simulations’. In: *ApJS* 137, pp. 1–20. eprint: astro-ph/0103301 (cit. on pp. 34, 35, 40).
- Bertschinger, Edmund (Jan. 1998). ‘Simulations of Structure Formation in the Universe’. In: *ARA&A* 36, pp. 599–654 (cit. on p. 27).
- Betoule, M. et al. (Aug. 2014). ‘Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples’. In: *Astronomy and Astrophysics* 568, A22, A22. arXiv: 1401.4064 (cit. on p. 22).

- Binney, J. (July 1977). ‘The physics of dissipational galaxy formation.’ In: *Apj* 215, pp. 483–491 (cit. on p. 20).
- Bode, Paul et al. (July 2001). ‘Halo Formation in Warm Dark Matter Models’. In: *The Astrophysical Journal* 556.1, pp. 93–107. arXiv: astro-ph/0010389 [astro-ph] (cit. on p. 24).
- Bond, J. R. et al. (Oct. 1991). ‘Excursion set mass functions for hierarchical Gaussian fluctuations’. In: *Apj* 379, pp. 440–460 (cit. on pp. 32, 46, 55, 57, 61).
- Brook, C. B. et al. (Mar. 2014). ‘The Stellar-to-halo Mass Relation for Local Group Galaxies’. In: *Apj* 784.1, L14, p. L14. arXiv: 1311.5492 [astro-ph.CO] (cit. on p. 88).
- Brown, Thomas M. et al. (Dec. 2014). ‘The Quenching of the Ultra-faint Dwarf Galaxies in the Reionization Era’. In: *Apj* 796.2, 91, p. 91. arXiv: 1410.0681 [astro-ph.GA] (cit. on p. 81).
- Bryan, Greg L. et al. (Apr. 2014). ‘ENZO: An Adaptive Mesh Refinement Code for Astrophysics’. In: *ApJS* 211.2, 19, p. 19. arXiv: 1307.2265 [astro-ph.IM] (cit. on p. 27).
- Bullock, J. S. et al. (July 2001a). ‘A Universal Angular Momentum Profile for Galactic Halos’. In: *Apj* 555, pp. 240–257. eprint: astro-ph/0011001 (cit. on pp. 71, 111).
- Bullock, J. S. et al. (Mar. 2001b). ‘Profiles of dark haloes: evolution, scatter and environment’. In: *MNRAS* 321, pp. 559–575. eprint: astro-ph/9908159 (cit. on pp. 60, 71, 72, 77).
- Bullock, James S. et al. (Aug. 2017). ‘Small-Scale Challenges to the Λ CDM Paradigm’. In: *ARA&A* 55.1, pp. 343–387. arXiv: 1707.04256 [astro-ph.CO] (cit. on p. 23).
- Bundy, Kevin et al. (Jan. 2015). ‘Overview of the SDSS-IV MaNGA Survey: Mapping nearby Galaxies at Apache Point Observatory’. In: *The Astrophysical Journal* 798.1, 7, p. 7. arXiv: 1412.1482 [astro-ph.GA] (cit. on p. 25).
- Carroll, Sean M. et al. (Aug. 2004). ‘Is cosmic speed-up due to new gravitational physics?’ In: *Phys. Rev. D* 70.4, 043528, p. 043528. arXiv: astro-ph/0306438 [astro-ph] (cit. on p. 22).
- Chisari, N. E. et al. (Nov. 2018). ‘The impact of baryons on the matter power spectrum from the Horizon-AGN cosmological hydrodynamical simulation’. In: *Monthly Notices of the Royal Astronomical Society* 480.3, pp. 3962–3977. arXiv: 1801.08559 [astro-ph.CO] (cit. on p. 24).
- Christensen, Charlotte R. et al. (June 2016). ‘In-N-Out: The Gas Cycle from Dwarfs to Spiral Galaxies’. In: *Apj* 824.1, 57, p. 57. arXiv: 1508.00007 [astro-ph.GA] (cit. on p. 30).
- Clementini, Gisella et al. (Sept. 2012). ‘Variability and Star Formation in Leo T, the Lowest Luminosity Star-forming Galaxy Known Today’. In: *Apj* 756.2, 108, p. 108. arXiv: 1207.2764 [astro-ph.GA] (cit. on p. 95).
- Coles, Peter et al. (Jan. 1991). ‘A lognormal model for the cosmological mass distribution.’ In: *MNRAS* 248, pp. 1–13 (cit. on p. 75).
- Cooper, A. P. et al. (Dec. 2015). ‘Formation of in situ stellar haloes in Milky Way-mass galaxies’. In: *MNRAS* 454, pp. 3185–3199. arXiv: 1501.04630 (cit. on p. 114).
- Cooray, A. et al. (Dec. 2002). ‘Halo models of large scale structure’. In: *Phys. Rep.* 372, pp. 1–129. eprint: astro-ph/0206508 (cit. on p. 59).
- Correa, Camila A. et al. (Sept. 2015). ‘The accretion history of dark matter haloes - III. A physical model for the concentration-mass relation’. In: *MNRAS* 452, pp. 1217–1232. arXiv: 1502.00391 [astro-ph.CO] (cit. on pp. 72, 73).

- Courty, S. et al. (Mar. 2004). ‘Thermodynamic evolution of cosmological baryonic gas. I. Influence of non-equipartition processes’. In: *A&A* 416, pp. 875–888. arXiv: astro-ph/0312532 [astro-ph] (cit. on pp. 83, 97).
- Croom, Scott M. et al. (Mar. 2012). ‘The Sydney-AAO Multi-object Integral field spectrograph’. In: *Monthly Notices of the Royal Astronomical Society* 421.1, pp. 872–893. arXiv: 1112.3367 [astro-ph.CO] (cit. on p. 25).
- de Haan, T. et al. (Nov. 2016). ‘Cosmological Constraints from Galaxy Clusters in the 2500 Square-degree SPT-SZ Survey’. In: *The Astrophysical Journal* 832.1, 95, p. 95. arXiv: 1603.06522 [astro-ph.CO] (cit. on p. 22).
- de Jong, J. T. A. et al. (June 2008). ‘The Structural Properties and Star Formation History of Leo T from Deep LBT Photometry’. In: *ApJ* 680.2, pp. 1112–1119. arXiv: 0801.4027 [astro-ph] (cit. on p. 95).
- Deason, Alis et al. (Oct. 2014). ‘Satellite Dwarf Galaxies in a Hierarchical Universe: The Prevalence of Dwarf-Dwarf Major Mergers’. In: *ApJ* 794.2, 115, p. 115. arXiv: 1406.3344 [astro-ph.GA] (cit. on p. 91).
- DeFelippis, Daniel et al. (July 2019). ‘A Correlated Search for Local Dwarf Galaxies in GALFA-H I and Pan-STARRS’. In: *ApJ* 879.1, 22, p. 22. arXiv: 1906.05287 [astro-ph.GA] (cit. on p. 96).
- Dehnen, W. et al. (May 2011). ‘N-body simulations of gravitational dynamics’. In: *European Physical Journal Plus* 126, 55, p. 55. arXiv: 1105.1082 [astro-ph.IM] (cit. on p. 27).
- Dekel, A. et al. (Apr. 1986). ‘The Origin of Dwarf Galaxies, Cold Dark Matter, and Biased Galaxy Formation’. In: *ApJ* 303, p. 39 (cit. on p. 21).
- DES collaboration Abbott, T. M. C. et al. (Aug. 2018). ‘Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing’. In: *Physical Review D* 98.4, 043526, p. 043526. arXiv: 1708.01530 [astro-ph.CO] (cit. on p. 22).
- DES collaboration Abbott, T. M. C. et al. (Feb. 2019). ‘First Cosmology Results using Type Ia Supernovae from the Dark Energy Survey: Constraints on Cosmological Parameters’. In: *The Astrophysical Journal* 872.2, L30, p. L30. arXiv: 1811.02374 [astro-ph.CO] (cit. on p. 23).
- Di Cintio, Arianna et al. (Jan. 2014). ‘The dependence of dark matter profiles on the stellar-to-halo mass ratio: a prediction for cusps versus cores’. In: *MNRAS* 437, pp. 415–423. arXiv: 1306.0898 [astro-ph.CO] (cit. on p. 24).
- Di Cintio, Arianna et al. (June 2019). ‘NIHAO XXI: the emergence of low surface brightness galaxies’. In: *MNRAS* 486.2, pp. 2535–2548. arXiv: 1901.08559 [astro-ph.GA] (cit. on p. 115).
- Di Matteo, Tiziana et al. (Feb. 2005). ‘Energy input from quasars regulates the growth and activity of black holes and their host galaxies’. In: *Nature* 433, pp. 604–607. eprint: astro-ph/0502199 (cit. on pp. 31, 40, 78).
- Diemer, Benedikt et al. (Jan. 2015). ‘A Universal Model for Halo Concentrations’. In: *ApJ* 799, 108, p. 108. arXiv: 1407.4730 [astro-ph.CO] (cit. on pp. 60, 111).
- Donnert, J. et al. (Jan. 2009). ‘Cluster magnetic fields from galactic outflows’. In: *MNRAS* 392.3, pp. 1008–1021. arXiv: 0808.0919 [astro-ph] (cit. on p. 33).
- Drlica-Wagner, Alex et al. (Feb. 2019). ‘Probing the Fundamental Nature of Dark Matter with the Large Synoptic Survey Telescope’. In: *arXiv e-prints*, arXiv:1902.01055, arXiv:1902.01055. arXiv: 1902.01055 [astro-ph.CO] (cit. on p. 24).

- Dubois, Yohan et al. (Dec. 2016). ‘The HORIZON-AGN simulation: morphological diversity of galaxies promoted by AGN feedback’. In: *MNRAS* 463.4, pp. 3948–3964. arXiv: 1606.03086 [astro-ph.GA] (cit. on pp. 29, 31).
- Efstathiou, G. (May 1992). ‘Suppressing the formation of dwarf galaxies via photoionization’. In: *MNRAS* 256.2, 43P–47P (cit. on pp. 24, 81, 95, 99).
- Efstathiou, G. et al. (Jan. 1979). ‘The rotation of galaxies: numerical investigations of the tidal torque theory.’ In: *MNRAS* 186, pp. 133–144 (cit. on p. 111).
- Efstathiou, George (June 1995). ‘An anthropic argument for a cosmological constant’. In: *MNRAS* 274.4, pp. L73–L76 (cit. on p. 22).
- Einasto, J. (1965). ‘On the Construction of a Composite Model for the Galaxy and on the Determination of the System of Galactic Parameters’. In: *Trudy Astrofizicheskogo Instituta Alma-Ata* 5, pp. 87–100 (cit. on p. 71).
- Eisenstein, D. J. et al. (Mar. 1998). ‘Baryonic Features in the Matter Transfer Function’. In: *ApJ* 496.2, pp. 605–614. arXiv: astro-ph/9709112 [astro-ph] (cit. on p. 34).
- Eisenstein, D. J. et al. (May 1998). ‘HOP: A New Group-Finding Algorithm for N-Body Simulations’. In: *ApJ* 498, pp. 137–142. eprint: astro-ph/9712200 (cit. on pp. 64, 82).
- Elias, L. M. et al. (Sept. 2018). ‘Stellar halos in Illustris: probing the histories of Milky Way-mass galaxies’. In: *MNRAS* 479, pp. 4004–4016. arXiv: 1801.07273 (cit. on p. 114).
- Escala, Ivanna et al. (Feb. 2018). ‘Modelling chemical abundance distributions for dwarf galaxies in the Local Group: the impact of turbulent metal diffusion’. In: *MNRAS* 474.2, pp. 2194–2211. arXiv: 1710.06533 [astro-ph.GA] (cit. on p. 86).
- Fabian, A. C. (Jan. 1994). ‘Cooling Flows in Clusters of Galaxies’. In: *ARA&A* 32, pp. 277–318 (cit. on p. 20).
- (Sept. 2012). ‘Observational Evidence of Active Galactic Nuclei Feedback’. In: *ARA&A* 50, pp. 455–489. arXiv: 1204.4114 [astro-ph.CO] (cit. on p. 29).
- Falck, B. et al. (Mar. 2017). ‘The Effect of Corner Modes in the Initial Conditions of Cosmological Simulations’. In: *ApJ* 837.2, 181, p. 181. arXiv: 1610.04862 [astro-ph.CO] (cit. on p. 34).
- Fan, Xiaohui et al. (July 2006). ‘Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with $z \sim 6$ Quasars. II. A Sample of 19 Quasars’. In: *The Astronomical Journal* 132.1, pp. 117–136. arXiv: astro-ph/0512082 [astro-ph] (cit. on p. 26).
- Faucher-Giguère, Claude-André et al. (Dec. 2010). ‘Ly α Cooling Emission from Galaxy Formation’. In: *ApJ* 725.1, pp. 633–657. arXiv: 1005.3041 [astro-ph.CO] (cit. on p. 84).
- Ferland, G. J. et al. (July 1998). ‘CLOUDY 90: Numerical Simulation of Plasmas and Their Spectra’. In: *PASP* 110.749, pp. 761–778 (cit. on p. 83).
- Fitts, Alex et al. (Nov. 2017). ‘fire in the field: simulating the threshold of galaxy formation’. In: *MNRAS* 471.3, pp. 3547–3562. arXiv: 1611.02281 [astro-ph.GA] (cit. on pp. 78, 82, 92).
- Flores, Ricardo A. et al. (May 1994). ‘Observational and Theoretical Constraints on Singular Dark Matter Halos’. In: *The Astrophysical Journal* 427, p. L1. arXiv: astro-ph/9402004 [astro-ph] (cit. on p. 23).
- Frenk, C. S. et al. (Nov. 1999). ‘The Santa Barbara Cluster Comparison Project: A Comparison of Cosmological Hydrodynamics Solutions’. In: *ApJ* 525.2, pp. 554–582. arXiv: astro-ph/9906160 [astro-ph] (cit. on p. 33).

- Fukugita, M. et al. (Aug. 1998). ‘The Cosmic Baryon Budget’. In: *Apj* 503.2, pp. 518–530. arXiv: astro-ph/9712020 [astro-ph] (cit. on p. 21).
- Gao, Liang et al. (Oct. 2005). ‘The age dependence of halo clustering’. In: *MNRAS* 363, pp. L66–L70. eprint: astro-ph/0506510 (cit. on pp. 60, 73).
- Gao, Liang et al. (June 2008). ‘The redshift dependence of the structure of massive Λ cold dark matter haloes’. In: *MNRAS* 387, pp. 536–544. arXiv: 0711.0746 [astro-ph] (cit. on p. 73).
- Garrison-Kimmel, Shea et al. (Jan. 2017). ‘Organized chaos: scatter in the relation between stellar mass and halo mass in small galaxies’. In: *MNRAS* 464.3, pp. 3108–3120. arXiv: 1603.04855 [astro-ph.GA] (cit. on p. 88).
- Garrison, Lehman H. et al. (Oct. 2016). ‘Improving initial conditions for cosmological N-body simulations’. In: *MNRAS* 461.4, pp. 4125–4145. arXiv: 1605.02333 [astro-ph.CO] (cit. on p. 34).
- Genel, Shy et al. (Jan. 2019). ‘A Quantification of the Butterfly Effect in Cosmological Simulations and Implications for Galaxy Scaling Relations’. In: *Apj* 871.1, 21, p. 21. arXiv: 1807.07084 [astro-ph.GA] (cit. on pp. 72, 78).
- Girardi, L. et al. (Dec. 2010). ‘The ACS Nearby Galaxy Survey Treasury. IX. Constraining Asymptotic Giant Branch Evolution with Old Metal-poor Galaxies’. In: *Apj* 724, pp. 1030–1043. arXiv: 1009.4618 [astro-ph.SR] (cit. on p. 86).
- Gnedin, Nickolay Y. (Oct. 2000). ‘Effect of Reionization on Structure Formation in the Universe’. In: *Apj* 542.2, pp. 535–541. arXiv: astro-ph/0002151 [astro-ph] (cit. on pp. 95, 99, 107).
- Gnedin, Oleg Y. (Jan. 2003). ‘Tidal Effects in Clusters of Galaxies’. In: *Apj* 582.1, pp. 141–161. arXiv: astro-ph/0302497 [astro-ph] (cit. on p. 33).
- Governato, F. et al. (Aug. 1999). ‘Properties of galaxy clusters: mass and correlation functions’. In: *MNRAS* 307.4, pp. 949–966. arXiv: astro-ph/9810189 [astro-ph] (cit. on p. 32).
- Guillet, Thomas et al. (June 2011). ‘A simple multigrid scheme for solving the Poisson equation with arbitrary domain boundaries’. In: *Journal of Computational Physics* 230.12, pp. 4756–4771. arXiv: 1104.1703 [physics.comp-ph] (cit. on p. 83).
- Guth, Alan H. (Jan. 1981). ‘Inflationary universe: A possible solution to the horizon and flatness problems’. In: *Phys. Rev. D* 23.2, pp. 347–356 (cit. on p. 18).
- Guth, Alan H. et al. (Oct. 1982). ‘Fluctuations in the New Inflationary Universe’. In: *Phys. Rev. Lett.* 49.15, pp. 1110–1113 (cit. on p. 18).
- Haardt, Francesco et al. (Apr. 1996). ‘Radiative Transfer in a Clumpy Universe. II. The Ultraviolet Extragalactic Background’. In: *Apj* 461, p. 20. arXiv: astro-ph/9509093 [astro-ph] (cit. on p. 97).
- Hahn, Oliver et al. (Aug. 2011). ‘Multi-scale initial conditions for cosmological simulations’. In: *MNRAS* 415.3, pp. 2101–2121. arXiv: 1103.6031 [astro-ph.CO] (cit. on pp. 34, 35).
- Harrison, E. R. (May 1970). ‘Fluctuations at the Threshold of Classical Cosmology’. In: *Phys. Rev. D* 1.10, pp. 2726–2730 (cit. on p. 19).
- Hawking, S. W. (Sept. 1982). ‘The development of irregularities in a single bubble inflationary universe’. In: *Physics Letters B* 115.4, pp. 295–297 (cit. on p. 18).
- Hayes, Matthew et al. (Sept. 2016). ‘O VI Emission Imaging of a Galaxy with the Hubble Space Telescope: a Warm Gas Halo Surrounding the Intense Starburst SDSS J115630.63+500822.1’. In: *The Astrophysical Journal* 828.1, 49, p. 49. arXiv: 1606.04536 [astro-ph.GA] (cit. on p. 26).

- Haynes, Martha P. et al. (Nov. 2011). ‘The Arecibo Legacy Fast ALFA Survey: The α .40 H I Source Catalog, Its Characteristics and Their Impact on the Derivation of the H I Mass Function’. In: *AJ* 142.5, 170, p. 170. arXiv: 1109.0027 [astro-ph.CO] (cit. on p. 96).
- Helmi, Amina et al. (Nov. 2018). ‘The merger that led to the formation of the Milky Way’s inner stellar halo and thick disk’. In: *Nature* 563.7729, pp. 85–88. arXiv: 1806.06038 [astro-ph.GA] (cit. on p. 114).
- Hernquist, Lars (June 1993). ‘N-body realizations of compound galaxies’. In: *ApJS* 86, pp. 389–400 (cit. on pp. 31, 40).
- Heß, S. et al. (Nov. 2013). ‘Simulating structure formation of the Local Universe’. In: *MNRAS* 435, pp. 2065–2076. arXiv: 1304.6565 (cit. on pp. 33, 40).
- Hikage, Chiaki et al. (Apr. 2019). ‘Cosmology from cosmic shear power spectra with Subaru Hyper Suprime-Cam first-year data’. In: *Publications of the Astronomical Society of Japan* 71.2, 43, p. 43. arXiv: 1809.09148 [astro-ph.CO] (cit. on p. 22).
- Hinshaw, G. et al. (Oct. 2013). ‘Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results’. In: *The Astrophysical Journal Supplement Series* 208.2, 19, p. 19. arXiv: 1212.5226 [astro-ph.CO] (cit. on p. 19).
- Hockney, R. W. et al. (1988). *Computer simulation using particles* (cit. on p. 27).
- Hoft, Matthias et al. (Sept. 2006). ‘Dwarf galaxies in voids: suppressing star formation with photoheating’. In: *MNRAS* 371.1, pp. 401–414. arXiv: astro-ph/0501304 [astro-ph] (cit. on pp. 95, 99, 100, 107).
- Hoffman, Y. et al. (Oct. 1991). ‘Constrained realizations of Gaussian fields - A simple algorithm’. In: *ApJ* 380, pp. L5–L8 (cit. on pp. 33, 35, 39–41, 43–45, 57, 117, 119).
- Hoffman, Y. et al. (Jan. 2017). ‘The dipole repeller’. In: *Nature Astronomy* 1, 0036, p. 0036. arXiv: 1702.02483 (cit. on pp. 33, 40).
- Homma, Daisuke et al. (July 2019). ‘Bootes. IV. A new Milky Way satellite discovered in the Subaru Hyper Suprime-Cam Survey and implications for the missing satellite problem’. In: *PASJ*, p. 91. arXiv: 1906.07332 [astro-ph.GA] (cit. on pp. 90, 91).
- Hopkins, P. F. et al. (Dec. 2010). ‘Mergers in Λ CDM: Uncertainties in Theoretical Predictions and Interpretations of the Merger Rate’. In: *ApJ* 724, pp. 915–945. arXiv: 1004.2708 [astro-ph.CO] (cit. on p. 68).
- Hopkins, P. F. et al. (Dec. 2012). ‘Stellar feedback and bulge formation in clumpy discs’. In: *MNRAS* 427, pp. 968–978. arXiv: 1111.6591 [astro-ph.CO] (cit. on p. 40).
- Hopkins, Philip F. (Feb. 2013). ‘A general class of Lagrangian smoothed particle hydrodynamics methods and implications for fluid mixing problems’. In: *MNRAS* 428.4, pp. 2840–2856. arXiv: 1206.5006 [astro-ph.IM] (cit. on p. 28).
- Hopkins, Philip F. et al. (Oct. 2018). ‘FIRE-2 simulations: physics versus numerics in galaxy formation’. In: *MNRAS* 480.1, pp. 800–863. arXiv: 1702.06148 [astro-ph.GA] (cit. on p. 30).
- Hoyle, F. (1949). *Problems of cosmical aerodynamics: proceedings of the Symposium on the Motion of Gaseous Masses of Cosmical Dimensions held at Paris, August 16-19, 1949*. English. Conference Proceedings. Organized by the International Union of Theoretical and Applied Mechanics and the International Astronomical Union with the cooperation of the United Nations Educational, Scientific and Cultural Organization (cit. on p. 111).

- Huang, Kuan-Wei et al. (June 2019). ‘The early growth of supermassive black holes in cosmological hydrodynamic simulations with constrained Gaussian realizations’. In: *arXiv e-prints*, arXiv:1906.00242, arXiv:1906.00242. arXiv: 1906 . 00242 [astro-ph.GA] (cit. on p. 33).
- Hui, Lam et al. (Nov. 1997). ‘Equation of state of the photoionized intergalactic medium’. In: *MNRAS* 292.1, pp. 27–42. arXiv: astro-ph/9612232 [astro-ph] (cit. on p. 100).
- Iršič, Vid et al. (Apr. 2017). ‘The Lyman α forest power spectrum from the XQ-100 Legacy Survey’. In: *Monthly Notices of the Royal Astronomical Society* 466.4, pp. 4332–4345. arXiv: 1702 . 01761 [astro-ph.CO] (cit. on p. 22).
- Irwin, M. J. et al. (Feb. 2007). ‘Discovery of an Unusual Dwarf Galaxy in the Outskirts of the Milky Way’. In: *ApJ* 656.1, pp. L13–L16. arXiv: astro-ph/0701154 [astro-ph] (cit. on pp. 95, 108, 113).
- Janesh, William et al. (May 2019). ‘Five Gas-rich Ultrafaint Dwarf Galaxy Candidates Discovered in WIYN Imaging of ALFALFA Sources’. In: *AJ* 157.5, 183, p. 183. arXiv: 1908 . 00438 [astro-ph.GA] (cit. on pp. 96, 108, 113).
- Jenkins, Adrian (Apr. 2010). ‘Second-order Lagrangian perturbation theory initial conditions for resimulations’. In: *MNRAS* 403.4, pp. 1859–1872. arXiv: 0910 . 0258 [astro-ph.CO] (cit. on p. 34).
- Jethwa, P. et al. (Jan. 2018). ‘The upper bound on the lowest mass halo’. In: *MNRAS* 473.2, pp. 2060–2083. arXiv: 1612 . 07834 [astro-ph.GA] (cit. on p. 92).
- Jiang, Fangzhou et al. (Oct. 2019). ‘Is the dark-matter halo spin a predictor of galaxy spin and size?’ In: *MNRAS* 488.4, pp. 4801–4815. arXiv: 1804 . 07306 [astro-ph.GA] (cit. on p. 77).
- Joachimi, Benjamin et al. (Nov. 2015). ‘Galaxy Alignments: An Overview’. In: *Space Science Reviews* 193.1-4, pp. 1–65. arXiv: 1504 . 05456 [astro-ph.GA] (cit. on p. 25).
- Johansson, P. H. et al. (Jan. 2009). ‘Equal- and Unequal-Mass Mergers of Disk and Elliptical Galaxies with Black Holes’. In: *ApJ* 690, pp. 802–821. arXiv: 0802 . 0210 (cit. on p. 78).
- Katz, Harley et al. (May 2019). ‘How to Quench a Dwarf Galaxy: The Impact of Inhomogeneous Reionization on Dwarf Galaxies and Cosmic Filaments’. In: *arXiv e-prints*, arXiv:1905.11414, arXiv:1905.11414. arXiv: 1905 . 11414 [astro-ph.GA] (cit. on p. 81).
- Katz, Neal et al. (Nov. 1992). ‘Galaxies and Gas in a Cold Dark Matter Universe’. In: *ApJ* 399, p. L109 (cit. on p. 29).
- Katz, Neal et al. (Aug. 1993). ‘Hierarchical Galaxy Formation: Overmerging and the Formation of an X-Ray Cluster’. In: *ApJ* 412, p. 455 (cit. on pp. 34, 61).
- Keller, B. W. et al. (Jan. 2019). ‘Chaos and variance in galaxy formation’. In: *MNRAS* 482.2, pp. 2244–2261. arXiv: 1803 . 05445 [astro-ph.GA] (cit. on pp. 72, 78).
- Kim, Chang-Goo et al. (Apr. 2015). ‘Momentum Injection by Supernovae in the Interstellar Medium’. In: *ApJ* 802.2, 99, p. 99. arXiv: 1410 . 1537 [astro-ph.GA] (cit. on p. 85).
- Kim, Ji-hoon et al. (Jan. 2014). ‘The AGORA High-resolution Galaxy Simulations Comparison Project’. In: *ApJS* 210.1, 14, p. 14. arXiv: 1308 . 2669 [astro-ph.GA] (cit. on p. 28).
- Kim, Ji-hoon et al. (Dec. 2016). ‘The AGORA High-resolution Galaxy Simulations Comparison Project. II. Isolated Disk Test’. In: *ApJ* 833.2, 202, p. 202. arXiv: 1610 . 03066 [astro-ph.GA] (cit. on p. 28).
- Kirby, Evan N. et al. (Dec. 2013). ‘The Universal Stellar Mass-Stellar Metallicity Relation for Dwarf Galaxies’. In: *ApJ* 779.2, 102, p. 102. arXiv: 1310 . 0814 [astro-ph.GA] (cit. on p. 90).

- Kirby, Evan N. et al. (Mar. 2014). ‘The dynamics of isolated Local Group galaxies’. In: *MNRAS* 439.1, pp. 1015–1027. arXiv: 1401.1208 [astro-ph.GA] (cit. on p. 90).
- Klein, Antoine et al. (Jan. 2016). ‘Science with the space-based interferometer eLISA: Supermassive black hole binaries’. In: *Physical Review D* 93.2, 024003, p. 024003. arXiv: 1511.05581 [gr-qc] (cit. on p. 26).
- Klypin, A. et al. (Apr. 2016). ‘MultiDark simulations: the story of dark matter halo concentrations and density profiles’. In: *MNRAS* 457, pp. 4340–4359. arXiv: 1411.4001 (cit. on pp. 60, 71, 72, 111).
- Klypin, Anatoly A. et al. (Oct. 2011). ‘Dark Matter Halos in the Standard Cosmological Model: Results from the Bolshoi Simulation’. In: *ApJ* 740.2, 102, p. 102. arXiv: 1002.3660 [astro-ph.CO] (cit. on p. 27).
- Klypin, Anatoly et al. (Sept. 1999). ‘Where Are the Missing Galactic Satellites?’ In: *The Astrophysical Journal* 522.1, pp. 82–92. arXiv: astro-ph/9901240 [astro-ph] (cit. on p. 23).
- Kroupa, Pavel (Apr. 2001). ‘On the variation of the initial mass function’. In: *MNRAS* 322.2, pp. 231–246. arXiv: astro-ph/0009005 (cit. on p. 84).
- Krumholz, Mark R. (June 2014). ‘The big problems in star formation: The star formation rate, stellar clustering, and the initial mass function’. In: *Phys. Rep.* 539, pp. 49–134. arXiv: 1402.0867 [astro-ph.GA] (cit. on p. 25).
- Laporte, C. F. P. et al. (Apr. 2015). ‘Under the sword of Damocles: plausible regeneration of dark matter cusps at the smallest galactic scales.’ In: *MNRAS* 449, pp. L90–L94. arXiv: 1409.3848 [astro-ph.GA] (cit. on p. 113).
- Laporte, N. et al. (Mar. 2017). ‘Dust in the Reionization Era: ALMA Observations of a $z = 8.38$ Gravitationally Lensed Galaxy’. In: *The Astrophysical Journal* 837.2, L21, p. L21. arXiv: 1703.02039 [astro-ph.GA] (cit. on p. 25).
- Lee, C. T. et al. (Apr. 2017). ‘Properties of dark matter haloes as a function of local environment density’. In: *MNRAS* 466, pp. 3834–3858. arXiv: 1610.02108 (cit. on pp. 60, 75).
- Leitherer, Claus et al. (July 1999). ‘Starburst99: Synthesis Models for Galaxies with Active Star Formation’. In: *ApJS* 123.1, pp. 3–40. arXiv: astro-ph/9902334 [astro-ph] (cit. on pp. 85, 98).
- Lesgourgues, Julien (Apr. 2011). ‘The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview’. In: *arXiv e-prints*, arXiv:1104.2932, arXiv:1104.2932. arXiv: 1104.2932 [astro-ph.IM] (cit. on p. 34).
- Lewis, Antony et al. (Aug. 2000). ‘Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models’. In: *ApJ* 538.2, pp. 473–476. arXiv: astro-ph/9911177 [astro-ph] (cit. on p. 34).
- Libeskind, Noam I. et al. (Jan. 2010). ‘Constrained simulations of the Local Group: on the radial distribution of substructures’. In: *MNRAS* 401.3, pp. 1889–1897. arXiv: 0909.4423 [astro-ph.CO] (cit. on p. 33).
- LIGO Scientific Collaboration et al. (Oct. 2017). ‘GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral’. In: *Phys. Rev. Lett.* 119 (16), p. 161101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101> (cit. on p. 26).
- Linde, A. D. (Feb. 1982). ‘A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems’. In: *Physics Letters B* 108.6, pp. 389–393 (cit. on p. 18).

- Lokhorst, Deborah et al. (May 2019). ‘On the Detectability of Visible-wavelength Line Emission from the Local Circumgalactic and Intergalactic Medium’. In: *The Astrophysical Journal* 877.1, 4, p. 4. arXiv: 1904 . 07874 [astro-ph.GA] (cit. on p. 26).
- Lucie-Smith, Luisa et al. (Sept. 2018). ‘Machine learning cosmological structure formation’. In: *MNRAS* 479.3, pp. 3405–3414. arXiv: 1802 . 04271 [astro-ph.CO] (cit. on p. 32).
- Ludlow, Aaron D. et al. (June 2013). ‘The mass profile and accretion history of cold dark matter haloes’. In: *MNRAS* 432, pp. 1103–1113. arXiv: 1302 . 0288 (cit. on pp. 36, 60, 72, 111).
- Ludlow, Aaron D. et al. (June 2014). ‘The mass-concentration-redshift relation of cold dark matter haloes’. In: *MNRAS* 441, pp. 378–388. arXiv: 1312 . 0945 (cit. on pp. 60, 71–73).
- Ludlow, Aaron D. et al. (Aug. 2016). ‘The mass-concentration-redshift relation of cold and warm dark matter haloes’. In: *MNRAS* 460, pp. 1214–1232. arXiv: 1601 . 02624 [astro-ph.CO] (cit. on pp. 36, 73).
- LUX collaboration et al. (Jan. 2017). ‘Results from a Search for Dark Matter in the Complete LUX Exposure’. In: *Phys. Rev. Lett.* 118.2, 021303, p. 021303. arXiv: 1608 . 07648 [astro-ph.CO] (cit. on p. 22).
- Lynden-Bell, D. (Jan. 1967). ‘Statistical mechanics of violent relaxation in stellar systems’. In: *MNRAS* 136, p. 101 (cit. on p. 20).
- Macciò, A. V. et al. (June 2007). ‘Concentration, spin and shape of dark matter haloes: scatter and the dependence on mass and environment’. In: *MNRAS* 378, pp. 55–71. eprint: astro-ph/0608157 (cit. on pp. 36, 60, 71, 72, 75–77, 111).
- Macciò, Andrea V. et al. (Dec. 2017). ‘The edge of galaxy formation - I. Formation and evolution of MW-satellite analogues before accretion’. In: *MNRAS* 472.2, pp. 2356–2366. arXiv: 1707 . 01106 [astro-ph.GA] (cit. on p. 82).
- Maggiore, Michele et al. (July 2010). ‘The Halo Mass Function from Excursion Set Theory. III. Non-Gaussian Fluctuations’. In: *ApJ* 717.1, pp. 526–541. arXiv: 0903 . 1251 [astro-ph.CO] (cit. on p. 32).
- Maoz, Dan et al. (Nov. 2012). ‘The delay-time distribution of Type Ia supernovae from Sloan II’. In: *MNRAS* 426.4, pp. 3282–3294. arXiv: 1206 . 0465 [astro-ph.CO] (cit. on p. 104).
- Martin, Nicolas F. et al. (Dec. 2016). ‘The PAndAS View of the Andromeda Satellite System. II. Detailed Properties of 23 M31 Dwarf Spheroidal Galaxies’. In: *ApJ* 833.2, 167, p. 167. arXiv: 1610 . 01158 [astro-ph.GA] (cit. on p. 90).
- Martizzi, Davide et al. (June 2015). ‘Supernova feedback in an inhomogeneous interstellar medium’. In: *MNRAS* 450.1, pp. 504–522. arXiv: 1409 . 4425 [astro-ph.GA] (cit. on p. 85).
- Mathis, H. et al. (July 2002). ‘Simulating the formation of the local galaxy population’. In: *MNRAS* 333.4, pp. 739–762. arXiv: astro-ph/0111099 [astro-ph] (cit. on p. 33).
- Maulbetsch, C. et al. (Jan. 2007). ‘The Dependence of the Mass Assembly History of Cold Dark Matter Halos on Environment’. In: *ApJ* 654, pp. 53–65. eprint: astro-ph/0606360 (cit. on pp. 60, 75).
- McConnachie, Alan W. (July 2012). ‘The Observed Properties of Dwarf Galaxies in and around the Local Group’. In: *AJ* 144.1, 4, p. 4. arXiv: 1204 . 1562 [astro-ph.CO] (cit. on p. 90).
- McQuinn, Kristen et al. (Oct. 2015). ‘Leo P: An Unquenched Very Low-mass Galaxy’. In: *ApJ* 812.2, 158, p. 158. arXiv: 1506 . 05495 [astro-ph.GA] (cit. on p. 96).
- McQuinn, Matthew (Sept. 2016). ‘The Evolution of the Intergalactic Medium’. In: *ARA&A* 54, pp. 313–362. arXiv: 1512 . 00086 [astro-ph.CO] (cit. on p. 99).

- McQuinn, Matthew et al. (Apr. 2009). ‘He II Reionization and its Effect on the Intergalactic Medium’. In: *ApJ* 694.2, pp. 842–866. arXiv: 0807 . 2799 [astro-ph] (cit. on p. 100).
- Merritt, A. et al. (Oct. 2016). ‘The Dragonfly nearby Galaxies Survey. I. Substantial Variation in the Diffuse Stellar Halos around Spiral Galaxies’. In: *ApJ* 830, 62, p. 62. arXiv: 1606 . 08847 (cit. on p. 114).
- Mo, H. J. et al. (Sept. 1996). ‘An analytic model for the spatial clustering of dark matter haloes’. In: *MNRAS* 282.2, pp. 347–361. arXiv: astro-ph/9512127 [astro-ph] (cit. on p. 32).
- Mo, H. et al. (May 2010). *Galaxy Formation and Evolution* (cit. on pp. 18, 32).
- Monachesi, Antonela et al. (Apr. 2016). ‘The GHOSTS survey - II. The diversity of halo colour and metallicity profiles of massive disc galaxies’. In: *MNRAS* 457.2, pp. 1419–1446. arXiv: 1507 . 06657 [astro-ph.GA] (cit. on p. 114).
- Monachesi, Antonela et al. (May 2019). ‘The Auriga stellar haloes: connecting stellar population properties with accretion and merging history’. In: *MNRAS* 485.2, pp. 2589–2616. arXiv: 1804 . 07798 [astro-ph.GA] (cit. on p. 114).
- Monaghan, J. J. (Jan. 1992). ‘Smoothed particle hydrodynamics.’ In: *ARA&A* 30, pp. 543–574 (cit. on p. 27).
- Moore, Ben (Aug. 1994). ‘Evidence against dissipation-less dark matter from observations of galaxy haloes’. In: *Nature* 370.6491, pp. 629–631 (cit. on p. 23).
- Moore, Ben et al. (Oct. 1999). ‘Dark Matter Substructure within Galactic Halos’. In: *The Astrophysical Journal* 524.1, pp. L19–L22. arXiv: astro-ph/9907411 [astro-ph] (cit. on p. 23).
- Moore, Gordon E. (Apr. 1965). ‘Cramming more components onto integrated circuits’. In: *Electronics* 38.8 (cit. on p. 30).
- Morrissey, Patrick et al. (Sept. 2018). ‘The Keck Cosmic Web Imager Integral Field Spectrograph’. In: *ApJ* 864.1, 93, p. 93. arXiv: 1807 . 10356 [astro-ph.IM] (cit. on p. 25).
- Mortlock, Daniel J. et al. (June 2011). ‘A luminous quasar at a redshift of $z = 7.085$ ’. In: *Nature* 474.7353, pp. 616–619. arXiv: 1106 . 6088 [astro-ph.CO] (cit. on p. 26).
- Moster, Benjamin P. et al. (Feb. 2013). ‘Galactic star formation and accretion histories from matching galaxies to dark matter haloes’. In: *MNRAS* 428.4, pp. 3121–3138. arXiv: 1205 . 5807 [astro-ph.CO] (cit. on p. 60).
- Mukhanov, V. F. et al. (June 1992). ‘Theory of cosmological perturbations’. In: *Phys. Rep.* 215.5-6, pp. 203–333 (cit. on pp. 19, 32).
- Munshi, Ferah et al. (May 2017). ‘Going, going, gone dark: Quantifying the scatter in the faintest dwarf galaxies’. In: *arXiv e-prints*, arXiv:1705.06286, arXiv:1705.06286. arXiv: 1705 . 06286 [astro-ph.GA] (cit. on pp. 37, 82, 88, 93).
- Munshi, Ferah et al. (Mar. 2019). ‘Dancing in the Dark: Uncertainty in Ultrafaint Dwarf Galaxy Predictions from Cosmological Simulations’. In: *ApJ* 874.1, 40, p. 40. arXiv: 1810 . 12417 [astro-ph.GA] (cit. on p. 84).
- Naab, Thorsten et al. (July 2009). ‘Minor Mergers and the Size Evolution of Elliptical Galaxies’. In: *ApJ* 699, pp. L178–L182. arXiv: 0903 . 1636 [astro-ph.CO] (cit. on p. 78).
- Naab, Thorsten et al. (Aug. 2017). ‘Theoretical Challenges in Galaxy Formation’. In: *ARA&A* 55.1, pp. 59–109. arXiv: 1612 . 06891 [astro-ph.GA] (cit. on pp. 21, 25, 29).

- Nadler, Ethan O. et al. (June 2019). ‘Constraints on Dark Matter Microphysics from the Milky Way Satellite Population’. In: *ApJ* 878.2, L32, p. L32. arXiv: 1904 . 10000 [astro-ph.CO] (cit. on p. 112).
- Natarajan, Priyamvada et al. (Apr. 2017). ‘Unveiling the First Black Holes With JWST:Multi-wavelength Spectral Predictions’. In: *ApJ* 838.2, 117, p. 117. arXiv: 1610 . 05312 [astro-ph.GA] (cit. on p. 26).
- Navarro, J. F. et al. (Nov. 1993). ‘Simulations of Dissipative Galaxy Formation in Hierarchically Clustering Universes - Part One - Tests of the Code’. In: *MNRAS* 265, p. 271 (cit. on p. 29).
- Navarro, Julio F. et al. (Dec. 1997). ‘A Universal Density Profile from Hierarchical Clustering’. In: *ApJ* 490, pp. 493–508. eprint: astro-ph/9611107 (cit. on pp. 23, 71, 111).
- Nelson, Dylan et al. (Feb. 2019a). ‘First Results from the TNG50 Simulation: Galactic outflows driven by supernovae and black hole feedback’. In: *arXiv e-prints*, arXiv:1902.05554, arXiv:1902.05554. arXiv: 1902 . 05554 [astro-ph.GA] (cit. on p. 30).
- Nelson, Dylan et al. (May 2019b). ‘The IllustrisTNG simulations: public data release’. In: *Computational Astrophysics and Cosmology* 6.1, 2, p. 2. arXiv: 1812 . 05609 [astro-ph.GA] (cit. on pp. 29, 31).
- Nocedal, J. et al. (2006). *Numerical Optimization*. 2nd. Springer (cit. on p. 50).
- Noh, Yookyung et al. (Oct. 2014). ‘A physical understanding of how reionization suppresses accretion on to dwarf haloes’. In: *MNRAS* 444.1, pp. 503–514. arXiv: 1401 . 0737 [astro-ph.CO] (cit. on pp. 95, 99, 100, 107).
- Nori, Matteo et al. (Aug. 2018). ‘AX-GADGET: a new code for cosmological simulations of Fuzzy Dark Matter and Axion models’. In: *MNRAS* 478.3, pp. 3935–3951. arXiv: 1801 . 08144 [astro-ph.CO] (cit. on p. 27).
- Okamoto, Takashi et al. (Nov. 2008). ‘Mass loss of galaxies due to an ultraviolet background’. In: *MNRAS* 390.3, pp. 920–928. arXiv: 0806 . 0378 [astro-ph] (cit. on pp. 95, 99, 100, 107).
- Oñorbe, Jose et al. (Jan. 2014). ‘How to zoom: bias, contamination and Lagrange volumes in multimass cosmological simulations’. In: *MNRAS* 437.2, pp. 1894–1908. arXiv: 1305 . 6923 [astro-ph.CO] (cit. on p. 34).
- Oñorbe, Jose et al. (Dec. 2015). ‘Forged in FIRE: cusps, cores and baryons in low-mass dwarf galaxies’. In: *MNRAS* 454.2, pp. 2092–2106. arXiv: 1502 . 02036 [astro-ph.GA] (cit. on pp. 86, 99).
- Palanque-Delabrouille, Nathalie et al. (Nov. 2013). ‘The one-dimensional Ly α forest power spectrum from BOSS’. In: *Astronomy and Astrophysics* 559, A85, A85. arXiv: 1306 . 5896 [astro-ph.CO] (cit. on p. 22).
- Paranjape, Aseem et al. (Nov. 2012). ‘Peaks theory and the excursion set approach’. In: *MNRAS* 426.4, pp. 2789–2796. arXiv: 1206 . 3506 [astro-ph.CO] (cit. on p. 32).
- Peek, J. E. G. et al. (June 2011). ‘The GALFA-HI Survey: Data Release 1’. In: *ApJS* 194.2, 20, p. 20. arXiv: 1101 . 1879 [astro-ph.GA] (cit. on p. 96).
- Planck Collaboration et al. (Sept. 2016a). ‘Planck 2015 results. XIII. Cosmological parameters’. In: *A&A* 594, A13, A13. arXiv: 1502 . 01589 (cit. on p. 63).
- Planck Collaboration et al. (Sept. 2016b). ‘Planck 2015 results. XXIV. Cosmology from Sunyaev-Zeldovich cluster counts’. In: *Astronomy and Astrophysics* 594, A24, A24. arXiv: 1502 . 01597 [astro-ph.CO] (cit. on p. 22).

- Planck Collaboration et al. (July 2018). ‘Planck 2018 results. VI. Cosmological parameters’. In: *arXiv e-prints*, arXiv:1807.06209, arXiv:1807.06209. arXiv: 1807 . 06209 [astro-ph.CO] (cit. on pp. 19, 23).
- Pontzen, Andrew et al. (Nov. 2008). ‘Damped Lyman α systems in galaxy formation simulations’. In: *MNRAS* 390.4, pp. 1349–1371. arXiv: 0804 . 4474 [astro-ph] (cit. on p. 84).
- Pontzen, Andrew et al. (Apr. 2012). ‘How supernova feedback turns dark matter cusps into cores’. In: *MNRAS* 421.4, pp. 3464–3471. arXiv: 1106 . 0499 [astro-ph.CO] (cit. on pp. 24, 113).
- Pontzen, Andrew et al. (May 2013). *pynbody: N-Body/SPH analysis for python*. Astrophysics Source Code Library. ascl: 1305 . 002 (cit. on p. 64).
- Pontzen, Andrew et al. (Feb. 2014). ‘Cold dark matter heats up’. In: *Nature* 506.7487, pp. 171–178. arXiv: 1402 . 1764 [astro-ph.CO] (cit. on pp. 82, 96).
- Pontzen, Andrew et al. (Feb. 2017). ‘How to quench a galaxy’. In: *MNRAS* 465.1, pp. 547–558. arXiv: 1607 . 02507 [astro-ph.GA] (cit. on pp. 33, 36, 39, 41, 53, 60, 61, 70, 78, 109, 115).
- Pontzen, Andrew et al. (Aug. 2018). ‘TANGOS: The Agile Numerical Galaxy Organization System’. In: *The Astrophysical Journal Supplement Series* 237, 23, p. 23 (cit. on p. 64).
- Porciani, C. (Dec. 2016). ‘Constrained simulations and excursion sets: understanding the risks and benefits of ‘genetically modified’ haloes’. In: *MNRAS* 463, pp. 4068–4082. arXiv: 1609 . 00730 (cit. on pp. 39, 45, 46).
- Porciani, Cristiano et al. (May 2002a). ‘Testing tidal-torque theory - I. Spin amplitude and direction’. In: *MNRAS* 332.2, pp. 325–338. arXiv: astro-ph/0105123 [astro-ph] (cit. on p. 111).
- (May 2002b). ‘Testing tidal-torque theory - II. Alignment of inertia and shear and the characteristics of protohaloes’. In: *MNRAS* 332.2, pp. 339–351. arXiv: astro-ph/0105165 [astro-ph] (cit. on p. 111).
- Power, C. et al. (Jan. 2003). ‘The inner structure of Λ CDM haloes - I. A numerical convergence study’. In: *MNRAS* 338.1, pp. 14–34. arXiv: astro-ph/0201544 [astro-ph] (cit. on pp. 64, 72).
- Prada, F. et al. (July 2012). ‘Halo concentrations in the standard Λ cold dark matter cosmology’. In: *MNRAS* 423, pp. 3018–3030. arXiv: 1104 . 5130 (cit. on p. 71).
- Press, W. H. et al. (Feb. 1974). ‘Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation’. In: *ApJ* 187, pp. 425–438 (cit. on pp. 32, 46, 57, 61).
- Prunet, S. et al. (Oct. 2008). ‘Initial Conditions For Large Cosmological Simulations’. In: *ApJS* 178.2, pp. 179–188. arXiv: 0804 . 3536 [astro-ph] (cit. on p. 34).
- Ragagnin, Antonio et al. (July 2019). ‘Dependency of halo concentration on mass, redshift and fossilness in Magneticum hydrodynamic simulations’. In: *MNRAS* 486.3, pp. 4001–4012. arXiv: 1810 . 08212 [astro-ph.CO] (cit. on p. 111).
- Raiteri, C. M. et al. (Nov. 1996). ‘Simulations of Galactic chemical evolution. I. O and Fe abundances in a simple collapse model.’ In: *A&A* 315, pp. 105–115 (cit. on p. 85).
- Rasera, Y. et al. (Jan. 2006). ‘The history of the baryon budget. Cosmic logistics in a hierarchical universe’. In: *A&A* 445.1, pp. 1–27. arXiv: astro-ph/0505473 [astro-ph] (cit. on p. 84).
- Ratra, Bharat et al. (June 1988). ‘Cosmological consequences of a rolling homogeneous scalar field’. In: *Phys. Rev. D* 37.12, pp. 3406–3427 (cit. on p. 22).
- Read, J. I. et al. (Feb. 2006). ‘The tidal stripping of satellites’. In: *Monthly Notices of the Royal Astronomical Society* 366.2, pp. 429–437. arXiv: astro-ph/0506687 [astro-ph] (cit. on p. 24).

- Read, J. I. et al. (July 2016). ‘Dark matter cores all the way down’. In: *MNRAS* 459.3, pp. 2573–2590. arXiv: 1508.04143 [astro-ph.GA] (cit. on p. 96).
- Read, J. I. et al. (May 2017). ‘The stellar mass-halo mass relation of isolated field dwarfs: a critical test of Λ CDM at the edge of galaxy formation’. In: *MNRAS* 467.2, pp. 2019–2038. arXiv: 1607.03127 [astro-ph.GA] (cit. on p. 88).
- Read, J. I. et al. (June 2019). ‘Abundance matching with the mean star formation rate: there is no missing satellites problem in the Milky Way’. In: *MNRAS*, p. 1615. arXiv: 1807.07093 [astro-ph.GA] (cit. on p. 92).
- Reed, Darren et al. (Dec. 2003). ‘Evolution of the mass function of dark matter haloes’. In: *MNRAS* 346.2, pp. 565–572. arXiv: astro-ph/0301270 [astro-ph] (cit. on p. 32).
- Rees, M. J. et al. (June 1977). ‘Cooling, dynamics and fragmentation of massive gas clouds: clues to the masses and radii of galaxies and clusters.’ In: *MNRAS* 179, pp. 541–559 (cit. on p. 20).
- Rey, Martin P. et al. (Feb. 2018). ‘Quadratic genetic modifications: a streamlined route to cosmological simulations with controlled merger history’. In: *MNRAS* 474.1, pp. 45–54. arXiv: 1706.04615 [astro-ph.GA] (cit. on pp. 3, 40).
- Rey, Martin P. et al. (Sept. 2019a). ‘EDGE: The origin of scatter in ultra-faint dwarf stellar masses and surface brightnesses’. In: *arXiv e-prints*, arXiv:1909.04664, arXiv:1909.04664. arXiv: 1909.04664 [astro-ph.GA] (cit. on pp. 3, 81).
- Rey, Martin P. et al. (May 2019b). ‘Sensitivity of dark matter haloes to their accretion histories’. In: *MNRAS* 485.2, pp. 1906–1915. arXiv: 1810.09473 [astro-ph.CO] (cit. on pp. 3, 59).
- Ricotti, Massimo (Jan. 2009). ‘Late gas accretion on to primordial minihaloes: a model for Leo T, dark galaxies and extragalactic high-velocity clouds’. In: *MNRAS* 392.1, pp. L45–L49. arXiv: 0806.2402 [astro-ph] (cit. on pp. 96, 100, 102, 107).
- Robertson, Brant et al. (May 2004). ‘Disk Galaxy Formation in a Λ Cold Dark Matter Universe’. In: *ApJ* 606.1, pp. 32–45. arXiv: astro-ph/0401252 [astro-ph] (cit. on p. 29).
- Rodriguez-Puebla, A. et al. (Jan. 2016). ‘Is main-sequence galaxy star formation controlled by halo mass accretion?’ In: *MNRAS* 455, pp. 2592–2606. arXiv: 1508.04842 (cit. on p. 60).
- Romano-Diaz, Emilio et al. (Feb. 2006). ‘Constrained Cosmological Simulations of Dark Matter Halos’. In: *ApJ* 637.2, pp. L93–L96. arXiv: astro-ph/0508272 [astro-ph] (cit. on p. 33).
- Romano-Diaz, Emilio et al. (Sept. 2011). ‘Galaxy Formation in Heavily Overdense Regions at $z \sim 10$: The Prevalence of Disks in Massive Halos’. In: *ApJ* 738.2, L19, p. L19. arXiv: 1106.1214 [astro-ph.CO] (cit. on p. 33).
- Romano-Diaz, Emilio et al. (Aug. 2014). ‘The Gentle Growth of Galaxies at High Redshifts in Overdense Environments’. In: *ApJ* 790.2, L32, p. L32. arXiv: 1407.2614 [astro-ph.GA] (cit. on p. 33).
- Rosdahl, J. et al. (June 2012). ‘Extended Ly α emission from cold accretion streams’. In: *MNRAS* 423.1, pp. 344–366. arXiv: 1112.4408 [astro-ph.CO] (cit. on pp. 83, 84, 97).
- Roth, Nina et al. (Jan. 2016). ‘Genetically modified haloes: towards controlled experiments in Λ CDM galaxy formation’. In: *MNRAS* 455.1, pp. 974–986. arXiv: 1504.07250 [astro-ph.GA] (cit. on pp. 33, 35, 36, 39, 40, 57, 60, 61, 63, 82, 83, 92, 106, 109, 112, 120).
- Ryan-Weber, Emma V. et al. (Feb. 2008). ‘The Local Group dwarf Leo T: HI on the brink of star formation’. In: *MNRAS* 384.2, pp. 535–540. arXiv: 0711.2979 [astro-ph] (cit. on p. 95).
- Sawala, Till et al. (Feb. 2016). ‘The chosen few: the low-mass haloes that host faint galaxies’. In: *MNRAS* 456.1, pp. 85–97. arXiv: 1406.6362 [astro-ph.CO] (cit. on pp. 78, 82, 93).

- Schaye, Joop et al. (Jan. 2015). 'The EAGLE project: simulating the evolution and assembly of galaxies and their environments'. In: *MNRAS* 446.1, pp. 521–554. arXiv: 1407.7040 [astro-ph.GA] (cit. on pp. 28, 29, 31).
- Scolnic, D. M. et al. (June 2018). 'The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample'. In: *The Astrophysical Journal* 859.2, 101, p. 101. arXiv: 1710.00845 [astro-ph.CO] (cit. on p. 22).
- Seljak, Uros et al. (Oct. 1996). 'A Line of Sight Integration Approach to Cosmic Microwave Background Anisotropies'. In: *ApJ* 469, p. 437. arXiv: astro-ph/9603033 [astro-ph] (cit. on p. 34).
- Sheth, Ravi K. et al. (May 2001). 'Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes'. In: *MNRAS* 323.1, pp. 1–12. arXiv: astro-ph/9907024 [astro-ph] (cit. on p. 32).
- Sheth, Ravi K. et al. (Jan. 2002). 'An excursion set model of hierarchical clustering: ellipsoidal collapse and the moving barrier'. In: *MNRAS* 329.1, pp. 61–75. arXiv: astro-ph/0105113 [astro-ph] (cit. on p. 32).
- Silk, Joseph et al. (Mar. 1998). 'Quasars and galaxy formation'. In: *A&A* 331, pp. L1–L4. arXiv: astro-ph/9801013 [astro-ph] (cit. on p. 21).
- Simon, Joshua D. (2019). 'The Faintest Dwarf Galaxies'. In: *Annual Review of Astronomy and Astrophysics* 57.1, pp. 375–415. arXiv: 1901.05465 [astro-ph.GA] (cit. on pp. 37, 81, 90, 91).
- Smith, Matthew C. et al. (May 2019). 'Cosmological simulations of dwarfs: the need for ISM physics beyond SN feedback alone'. In: *MNRAS* 485.3, pp. 3317–3333. arXiv: 1807.04288 [astro-ph.GA] (cit. on p. 85).
- Smoot, G. F. et al. (Sept. 1992). 'Structure in the COBE Differential Microwave Radiometer First-Year Maps'. In: *ApJ* 396, p. L1 (cit. on p. 19).
- Somerville, R. S. et al. (Dec. 2008). 'A semi-analytic model for the co-evolution of galaxies, black holes and active galactic nuclei'. In: *MNRAS* 391, pp. 481–506. arXiv: 0808.1227 (cit. on p. 77).
- Somerville, Rachel S. (June 2002). 'Can Photoionization Squelching Resolve the Substructure Crisis?' In: *ApJ* 572.1, pp. L23–L26. arXiv: astro-ph/0107507 [astro-ph] (cit. on p. 24).
- Somerville, Rachel S. et al. (May 1999). 'How to plant a merger tree'. In: *MNRAS* 305.1, pp. 1–14. arXiv: astro-ph/9711080 [astro-ph] (cit. on p. 32).
- Somerville, Rachel S. et al. (Aug. 2015). 'Physical Models of Galaxy Formation in a Cosmological Framework'. In: *ARA&A* 53, pp. 51–113. arXiv: 1412.2712 [astro-ph.GA] (cit. on pp. 18, 25, 29).
- Sommer-Larsen, Jesper et al. (July 1999). 'Formation of Disk Galaxies: Feedback and the Angular Momentum Problem'. In: *ApJ* 519.2, pp. 501–512. arXiv: astro-ph/9801094 [astro-ph] (cit. on p. 29).
- Sorce, J. G. et al. (Jan. 2016). 'Cosmicflows Constrained Local UniversE Simulations'. In: *MNRAS* 455, pp. 2078–2090. arXiv: 1510.04900 (cit. on pp. 33, 40).
- Spergel, David N. et al. (Apr. 2000). 'Observational Evidence for Self-Interacting Cold Dark Matter'. In: *Physical Review Letters* 84.17, pp. 3760–3763. arXiv: astro-ph/9909386 [astro-ph] (cit. on p. 24).
- Springel, Volker (Dec. 2005). 'The cosmological simulation code GADGET-2'. In: *MNRAS* 364.4, pp. 1105–1134. arXiv: astro-ph/0505010 [astro-ph] (cit. on p. 27).

- Springel, Volker (Jan. 2010a). ‘E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh’. In: *MNRAS* 401.2, pp. 791–851. arXiv: 0901.4107 [astro-ph.CO] (cit. on p. 28).
- (Sept. 2010b). ‘Smoothed Particle Hydrodynamics in Astrophysics’. In: *ARA&A* 48, pp. 391–430. arXiv: 1109.2219 [astro-ph.CO] (cit. on p. 27).
- Springel, Volker et al. (Mar. 2005). ‘Formation of a Spiral Galaxy in a Major Merger’. In: *ApJ* 622.1, pp. L9–L12. arXiv: astro-ph/0411379 [astro-ph] (cit. on p. 31).
- Springel, Volker et al. (Aug. 2005). ‘Modelling feedback from stars and black holes in galaxy mergers’. In: *MNRAS* 361, pp. 776–794. arXiv: astro-ph/0411108 [astro-ph] (cit. on p. 78).
- Sutherland, Ralph S. et al. (Sept. 1993). ‘Cooling Functions for Low-Density Astrophysical Plasmas’. In: *ApJS* 88, p. 253 (cit. on p. 20).
- Teyssier, R. (Apr. 2002). ‘Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES’. In: *A&A* 385, pp. 337–364. arXiv: astro-ph/0111367 [astro-ph] (cit. on pp. 27, 64, 83, 97).
- Teyssier, Romain (Aug. 2015). ‘Grid-Based Hydrodynamics in Astrophysical Fluid Flows’. In: *ARA&A* 53, pp. 325–364 (cit. on p. 27).
- Thacker, R. J. et al. (July 2001). ‘Star Formation, Supernova Feedback, and the Angular Momentum Problem in Numerical Cold Dark Matter Cosmogony: Halfway There?’ In: *ApJ* 555.1, pp. L17–L20. arXiv: astro-ph/0106060 [astro-ph] (cit. on p. 29).
- Theuns, Tom et al. (Dec. 1998). ‘P³M-SPH simulations of the Ly α forest’. In: *MNRAS* 301.2, pp. 478–502. arXiv: astro-ph/9805119 [astro-ph] (cit. on p. 102).
- Tinker, Jeremy et al. (Dec. 2008). ‘Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality’. In: *ApJ* 688.2, pp. 709–728. arXiv: 0803.2706 [astro-ph] (cit. on p. 32).
- Tonnesen, Stephanie (Apr. 2019). ‘The Journey Counts: The Importance of Including Orbits when Simulating Ram Pressure Stripping’. In: *ApJ* 874.2, 161, p. 161. arXiv: 1903.08178 [astro-ph.GA] (cit. on p. 110).
- Toro, E. F. et al. (July 1994). ‘Restoration of the contact surface in the HLL-Riemann solver’. In: *Shock Waves* 4.1, pp. 25–34 (cit. on p. 83).
- Torrealba, G. et al. (Sept. 2019). ‘The hidden giant: discovery of an enormous Galactic dwarf satellite in Gaia DR2’. In: *MNRAS* 488.2, pp. 2743–2766. arXiv: 1811.04082 [astro-ph.GA] (cit. on pp. 90, 91).
- Tumlinson, Jason et al. (Nov. 2013). ‘The COS-Halos Survey: Rationale, Design, and a Census of Circumgalactic Neutral Hydrogen’. In: *The Astrophysical Journal* 777.1, 59, p. 59. arXiv: 1309.6317 [astro-ph.CO] (cit. on p. 26).
- Upton Sanderbeck, Phoebe R. et al. (Aug. 2016). ‘Models of the thermal evolution of the intergalactic medium after reionization’. In: *MNRAS* 460.2, pp. 1885–1897. arXiv: 1511.05992 [astro-ph.CO] (cit. on p. 100).
- Valkenburg, Wessel et al. (June 2017). ‘Accurate initial conditions in mixed dark matter-baryon simulations’. In: *MNRAS* 467.4, pp. 4401–4409. arXiv: 1610.08501 [astro-ph.CO] (cit. on p. 34).
- van Dokkum, Pieter G. et al. (Jan. 2015). ‘Forty-seven Milky Way-sized, Extremely Diffuse Galaxies in the Coma Cluster’. In: *ApJ* 798.2, L45, p. L45. arXiv: 1410.8141 [astro-ph.GA] (cit. on p. 114).

- van Uitert, Edo et al. (June 2018). ‘KiDS+GAMA: cosmology constraints from a joint analysis of cosmic shear, galaxy-galaxy lensing, and angular clustering’. In: *Monthly Notices of the Royal Astronomical Society* 476.4, pp. 4662–4689. arXiv: 1706.05004 [astro-ph.CO] (cit. on p. 22).
- Veilleux, Sylvain et al. (Sept. 2005). ‘Galactic Winds’. In: *ARA&A* 43.1, pp. 769–826. arXiv: astro-ph/0504435 [astro-ph] (cit. on p. 29).
- Vijayaraghavan, Rukmani et al. (May 2015). ‘Ram pressure stripping of hot coronal gas from group and cluster galaxies and the detectability of surviving X-ray coronae’. In: *MNRAS* 449.3, pp. 2312–2335. arXiv: 1503.01121 [astro-ph.GA] (cit. on p. 110).
- Vitvitska, Maya et al. (Dec. 2002). ‘The Origin of Angular Momentum in Dark Matter Halos’. In: *ApJ* 581, pp. 799–809. arXiv: astro-ph/0105349 [astro-ph] (cit. on p. 77).
- Vogelsberger, Mark et al. (Oct. 2014). ‘Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe’. In: *MNRAS* 444.2, pp. 1518–1547. arXiv: 1405.2921 [astro-ph.CO] (cit. on pp. 29, 31).
- Volonteri, Marta (July 2010). ‘Formation of supermassive black holes’. In: *A&ARv* 18.3, pp. 279–315. arXiv: 1003.4404 [astro-ph.CO] (cit. on p. 25).
- Wadsley, J. W. et al. (Feb. 2004). ‘Gasoline: a flexible, parallel implementation of TreeSPH’. In: *New Astron.* 9.2, pp. 137–158. arXiv: astro-ph/0303521 [astro-ph] (cit. on p. 27).
- Wadsley, James W. et al. (Oct. 2017). ‘Gasoline2: a modern smoothed particle hydrodynamics code’. In: *MNRAS* 471.2, pp. 2357–2369. arXiv: 1707.03824 [astro-ph.IM] (cit. on p. 28).
- Walter, Fabian et al. (Dec. 2016). ‘ALMA Spectroscopic Survey in the Hubble Ultra Deep Field: Survey Description’. In: *The Astrophysical Journal* 833.1, 67, p. 67. arXiv: 1607.06768 [astro-ph.GA] (cit. on p. 25).
- Wang, H. et al. (Nov. 2016). ‘ELUCID - Exploring the Local Universe with ReConstructed Initial Density Field III: Constrained Simulation in the SDSS Volume’. In: *ApJ* 831, 164, p. 164. arXiv: 1608.01763 (cit. on pp. 33, 40).
- Wechsler, Risa H. et al. (Mar. 2002). ‘Concentrations of Dark Halos from Their Assembly Histories’. In: *ApJ* 568, pp. 52–70. eprint: astro-ph/0108151 (cit. on pp. 36, 60, 71–73, 111).
- Wechsler, Risa H. et al. (Nov. 2006). ‘The Dependence of Halo Clustering on Halo Formation History, Concentration, and Occupation’. In: *ApJ* 652, pp. 71–84. eprint: astro-ph/0512416 (cit. on p. 60).
- Wechsler, Risa H. et al. (Sept. 2018). ‘The Connection Between Galaxies and Their Dark Matter Halos’. In: *ARA&A* 56, pp. 435–487. arXiv: 1804.03097 [astro-ph.GA] (cit. on pp. 20, 25, 59, 88).
- Weinberg, David H. et al. (Sept. 2013). ‘Observational probes of cosmic acceleration’. In: *Physics Reports* 530.2, pp. 87–255. arXiv: 1201.2434 [astro-ph.CO] (cit. on p. 22).
- Weinberg, Steven (Jan. 1989). ‘The cosmological constant problem’. In: *Reviews of Modern Physics* 61.1, pp. 1–23 (cit. on p. 22).
- Weinberger, Rainer et al. (Sept. 2019). ‘The Arepo public code release’. In: *arXiv e-prints*, arXiv:1909.04667, arXiv:1909.04667. arXiv: 1909.04667 [astro-ph.IM] (cit. on p. 28).
- Weisz, Daniel R. et al. (Apr. 2012). ‘The Star Formation History of Leo T from Hubble Space Telescope Imaging’. In: *ApJ* 748.2, 88, p. 88. arXiv: 1201.4859 [astro-ph.CO] (cit. on pp. 95, 99).
- Weisz, Daniel R. et al. (July 2014). ‘The Star Formation Histories of Local Group Dwarf Galaxies. I. Hubble Space Telescope/Wide Field Planetary Camera 2 Observations’. In: *ApJ* 789.2, 147, p. 147. arXiv: 1404.7144 [astro-ph.GA] (cit. on p. 81).

- Wetzel, Andrew R. et al. (Feb. 2007). ‘The Clustering of Massive Halos’. In: *Apj* 656, pp. 139–147. arXiv: astro-ph/0606699 [astro-ph] (cit. on p. 60).
- Wheeler, C. et al. (Dec. 2018). ‘Be it therefore resolved: Cosmological Simulations of Dwarf Galaxies with Extreme Resolution’. In: *arXiv e-prints*, arXiv:1812.02749, arXiv:1812.02749. arXiv: 1812.02749 [astro-ph.GA] (cit. on pp. 30, 37, 82).
- White, S. D. M. (Nov. 1984). ‘Angular momentum growth in protogalaxies’. In: *Apj* 286, pp. 38–41 (cit. on p. 111).
- White, S. D. M. et al. (May 1978). ‘Core condensation in heavy halos: a two-stage theory for galaxy formation and clustering.’ In: *MNRAS* 183, pp. 341–358 (cit. on p. 21).
- Wright, Anna C. et al. (Jan. 2019). ‘Reignition of star formation in dwarf galaxies’. In: *MNRAS* 482.1, pp. 1176–1189. arXiv: 1802.03019 [astro-ph.GA] (cit. on pp. 96, 100, 108).
- Xenon collaboration et al. (Sept. 2018). ‘Dark Matter Search Results from a One Ton-Year Exposure of XENON1T’. In: *Phys. Rev. Lett.* 121.11, 111302, p. 111302. arXiv: 1805.12562 [astro-ph.CO] (cit. on p. 22).
- Zana, Tommaso et al. (Jan. 2018). ‘External versus internal triggers of bar formation in cosmological zoom-in simulations’. In: *MNRAS* 473.2, pp. 2608–2621. arXiv: 1705.02348 [astro-ph.GA] (cit. on p. 110).
- Zel’dovich, Y. B. (Mar. 1970). ‘Gravitational instability: An approximate theory for large density perturbations.’ In: *A&A* 5, pp. 84–89 (cit. on pp. 34, 63).
- Zeldovich, Yaa B. (Jan. 1972). ‘A hypothesis, unifying the structure and the entropy of the Universe’. In: *MNRAS* 160, 1P (cit. on p. 19).
- Zhao, D. H. et al. (Feb. 2003). ‘The growth and structure of dark matter haloes’. In: *MNRAS* 339, pp. 12–24. arXiv: astro-ph/0204108 [astro-ph] (cit. on pp. 36, 60, 73).
- Zolotov, A. et al. (Sept. 2009). ‘The Dual Origin of Stellar Halos’. In: *Apj* 702, pp. 1058–1067. arXiv: 0904.3333 [astro-ph.GA] (cit. on p. 114).

