THE UNIVERSITY *of* EDINBURGH

The Star-Formation Histories of Massive Quiescent Galaxies

Adam Christopher Carnall

Doctor of Philosophy
The University of Edinburgh
May 2019

For my parents and grandparents

"All you really need to know for the moment is that the Universe is a lot more complicated than you might think, even if you start from a position of thinking it's pretty damn complicated in the first place."

– Douglas Adams

"I mean, look at all the crazy crap surrounding us, I mean, look at that thing right there, what is that thing? ... It defies all logic that thing."

– Rick Sanchez

Lay summary

The Universe contains a large number of widely separated collections of stars, known as galaxies, which are scattered within a huge, mostly empty space. These galaxies can be separated into two types, most commonly referred to as spiral and elliptical galaxies. Spiral galaxies are disk-shaped, and tend to be smaller and bluer than elliptical galaxies, which are roughly spherical, larger and redder. The difference in colour comes from the different kinds of stars spiral and elliptical galaxies contain. Blue stars are very bright, however they are short-lived, meaning that a galaxy that stops forming new stars quickly turns from blue to red.

One of the key goals of modern astronomy is to understand why these red galaxies stopped forming stars. Galaxies that no longer form stars are called quiescent galaxies, and the shutting down of star-formation is referred to as quenching. Because galaxies evolve so slowly, over millions or billions of years, we cannot watch quenching in progress. Instead we have to observe large numbers of galaxies at different stages in their evolution, and try to piece together a coherent timeline of events. The only information we have is the light from galaxies, which we try to interpret in order to understand their properties.

In this thesis, I use observations from a variety of telescopes worldwide and in space to try to understand the past history of star-formation in quiescent galaxies. In order to do this, I have written a piece of software called Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation, or Bagpipes, which can be used to interpret the light from galaxies. I use Bagpipes to analyse large numbers of quiescent galaxies from a variety of different sky surveys to try to understand what processes lead to the quenching of star-formation.

My thesis forms part of a large body of work on this subject, from which a coherent picture of galaxy quenching is beginning to emerge. Within the first five billion years of the Universe, galaxies typically shut down star-formation due to

violent events, such as collisions between galaxies and huge outbursts of energy due to gas falling into supermassive black holes in galaxy cores. Lower levels of supermassive black hole activity continue after these violent events, preventing star-formation from restarting. After this time, approximately nine billion years ago, a transition occurs, with extremely violent events becoming rarer, and star-formation in galaxies dying down more slowly as they become less efficient at drawing in gas from their surroundings.

Many questions remain unanswered, and our various pieces of understanding about the processes by which galaxies form and evolve are only just beginning to be joined up into a coherent whole. Upcoming telescopes and instruments, such as the Multi-Object Optical and Near-infrared Spectrograph (MOONS) on the Very Large Telescope at Paranal Observatory, the James Webb Space Telescope, and the European Extremely Large Telescope will provide us with huge advances in our understanding as we move into the 2020s.

Abstract

This thesis presents several related analyses designed to understand the star-formation histories (SFHs) and quenching mechanisms of massive quiescent galaxies across cosmic time. More generally, it contains research directed at sophisticated modelling and Bayesian fitting of galaxy spectra. I firstly present Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation, or Bagpipes, a new, publicly available Python code that can be used to rapidly generate complex model galaxy spectra and to fit these to arbitrary combinations of spectroscopic and photometric data.

I then perform a detailed analysis of the SFHs of a sample of 9289 quiescent galaxies from UltraVISTA with stellar masses, $M_* > 10^{10}\mathrm{M}_\odot$ and observed redshifts from $0.25 < z < 3.75$. The majority of these galaxies exhibit SFHs that rise gradually then quench relatively rapidly, over 1−2 Gyr. This behaviour is consistent with recent cosmological hydrodynamic simulations, where AGN-driven feedback in the low-accretion (jet) mode is the dominant quenching mechanism. At $z > 1$, I also find a class of objects with SFHs that rise and fall very rapidly, with quenching timescales of < 1 Gyr, consistent with quasar-mode AGN feedback. Finally, at $z < 1$, I find a population with SFHs that quench more slowly than they rise, over > 3 Gyr, consistent with other such analyses in the local Universe. I confirm the trend towards earlier formation with increasing stellar mass (downsizing) at fixed observed redshift, and a trend towards more rapid quenching at higher stellar masses.

I then present a general investigation of the use of parametric SFH models in spectral fitting analyses. Parametric models for galaxy SFHs are widely used, though they are known to impose strong priors on physical parameters, with consequences for measurements of the galaxy stellar-mass function, star-formation-rate density (SFRD) and star-forming main sequence (SFMS). I investigate the effects of the exponentially declining, delayed exponentially

declining, lognormal and double power law SFH models. I demonstrate that each of these models imposes strong priors on specific star-formation rates (sSFRs), potentially biasing the SFMS, and also imposes a strong prior preference for young stellar populations. I show that stellar mass, SFR and mass-weighted age inferences from high-quality mock photometry vary with the choice of SFH model by at least 0.1, 0.3 and 0.2 dex respectively. However the biases with respect to the true values depend more on the true SFH shape than the choice of model. I also demonstrate that photometric data cannot discriminate between SFH models, meaning it is important to perform independent tests to find well-motivated priors. In response to this I finally fit a low-redshift, volume-complete sample from the Galaxy and Mass Assembly (GAMA) Survey with each model. I demonstrate that the inferred stellar masses and SFRs at redshift, $z \sim 0.05$ are consistent with other analyses. However, the inferred cosmic SFRDs peak at $z \sim 0.4$, approximately 6 Gyr later than direct observations suggest, meaning that mass-weighted ages are significantly underestimated. This makes the use of parametric SFH models for understanding mass assembly in galaxies challenging.

I finally present a Bayesian full-spectral-fitting analysis of 75 massive ($\boldsymbol{M}_* > 10^{10.3}\mathrm{M}_\odot$) UVJ-selected galaxies at redshifts of $1.0 < z < 1.3$, combining extremely deep rest-frame ultraviolet spectroscopy from VANDELS with multi-wavelength photometry by the use of a sophisticated physical plus systematic uncertainties model. I constrain the stellar mass vs stellar age relationship, finding a strong trend towards earlier formation with increasing stellar mass (downsizing) of $1.48^{+0.34}_{-0.39}$ Gyr per decade in mass. I show that this is consistent with other spectroscopic studies from $0 < z < 2$. This places strong constraints on the AGN-feedback models used in cosmological simulations. I demonstrate that, although the relationships predicted by the Simba and IllustrisTNG simulations agree well with observations at $z = 0.1$, they are too shallow at $z = 1$, predicting an evolution of $\lesssim 0.5$ Gyr per decade in mass. The majority of the lowest-mass galaxies in the sample ($\boldsymbol{M}_* \sim 10^{10.5}\mathrm{M}_\odot$) are consistent with formation in recent ($z < 2$), intense starburst events, with timescales of $\lesssim 500$ Myr. A second class of objects experience extended star-formation epochs before rapidly quenching, passing through both green-valley and post-starburst phases. The most massive galaxies in the sample are extreme systems: already old by $z = 1$, they formed at $z \sim 5$ and quenched by $z = 3$. However, I find evidence for their continued evolution through both AGN and rejuvenated star-formation activity. To understand the detailed SFHs of these objects, similar studies must be extended to the highest redshifts.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

Parts of this work have been published in Carnall et al. (2018, 2019a,b).

(*Adam Christopher Carnall, May 2019*)

Acknowledgements

The work presented in this thesis belongs to a large number of people who've given me the inspiration, opportunities, help and advice I needed along the way.

I'd firstly like to wholeheartedly thank my supervisors, Ross McLure and Jim Dunlop, for their solid, dependable support throughout the past four years. Every single meeting I've had with either Ross or Jim during my PhD has made me feel more confident and more optimistic about my work than I did beforehand.

I'd like to thank all of the coauthors on the published papers that make up this thesis, as well as all of the members of the VANDELS collaboration. I'd also like to thank the European Southern Observatory for giving me the chance to observe on the Very Large Telescope at Paranal Observatory, it was an experience like no other. I'd also like to thank the Scottish Universities Physics Alliance for providing the funding that made the work presented in Chapter 3 possible.

I must also express a huge debt of gratitude to those who helped put me in a position to begin my PhD. Most particularly I thank Tom Shanks for his completely unjustified faith in the scientific abilities of MPhys students. Additional thanks must go a large number of teachers, lecturers, and those who generously invested their time in supervising me for summer projects as an undergraduate student: Diego Altamirano, Andrew Blain and Ryan Houghton.

I'd also like to thank everyone at the Royal Observatory Edinburgh who I've become friends with over the past four years. For fear of omission I won't try to list everyone, but huge numbers of past and present staff, postdocs and PhD students have helped to make my time in Edinburgh enjoyable. I would however particularly like to thank Tom Kemp and Fergus Cullen for getting me back safely from Italy during the broken leg incident.

The greatest thanks must of course go to my family, most particularly my parents, who have always supported me, and always believed in me. I thank my mum for encouraging me to be curious from a young age, and convincing me to do a physics degree. I thank my dad for always being interested in what I'm doing, and for beginning my interest in astronomy. I also thank my grandparents for their inspirational examples.

I must thank Fran Lane for her unconventional yet invaluable support over the past two and a half years. Thank you for making me laugh, thank you for driving me around whilst my leg healed, and thank you for always providing me with change for the vending machine.

I would finally like to express special thanks to everyone involved in the BBC TV series "The Planets" (1999), as well as the BBC Horizon episodes "The Planet Hunters" (1996) and "Supermassive Black Holes" (2000). More than anything else I associate these TV programs with my early interest in astronomy, which I have been hugely privileged to be able to pursue in my adult life.

Contents

Lay summary v

Abstract vii

Declaration ix

Acknowledgements xi

Contents xiii

List of Figures xix

List of Tables xxiii

1 Introduction 1

1.1 Cosmic history under the Λ-CDM model 2

1.1.1 The expanding Universe and distance measurements 2

1.1.2 Primordial nucleosynthesis 5

1.1.3 Recombination and the cosmic microwave background 5

1.1.4 Dark matter and large-scale structure 7

1.1.5 The first galaxies and reionization 10

1.1.6 Dark energy and accelerating expansion 11

1.2 Astronomical Observations 12
1.2.1 Photometry vs spectroscopy 13
1.2.2 Spatial and spectral resolution 13
1.2.3 Fluxes and magnitudes 14
1.2.4 Observing galaxy populations 15
1.2.5 Data reduction 17
1.3 Modelling galaxy spectra 18
1.3.1 Stellar population synthesis 18
1.3.2 Dust extinction, attenuation and emission 24
1.3.3 Nebular emission 28
1.3.4 Intergalactic medium attenuation 28
1.3.5 The age-metallicity-dust degeneracy 29
1.3.6 Monochromatic star-formation-rate indicators 30
1.4 Bayesian statistics and computational methods 30
1.4.1 Parameter estimation and marginalisation 32
1.4.2 Bayesian model selection 33
1.4.3 Solving Bayes' equation: computational methods 36
1.5 Towards a theory of galaxy formation 40
1.5.1 The galaxy luminosity and stellar-mass functions 40
1.5.2 Cosmic star-formation history 44
1.5.3 The star-forming sequence 45
1.5.4 The galaxy colour bimodality and quenching 47
1.5.5 Metal enrichment 49
1.5.6 Dust attenuation 51

1.6 Thesis outline 52

2 Inferring the star-formation histories of massive quiescent galaxies with BAGPIPES: Evidence for multiple quenching mechanisms 55

2.1 Introduction 55

2.2 The UltraVISTA data 61

2.3 The BAGPIPES code 62

2.3.1 Model generation 62

2.3.2 Model fitting 73

2.4 Testing star-formation-history models with MUFASA 78

2.4.1 Common SFH parameterisations 79

2.4.2 Comparisons between SFH parameterisations 80

2.4.3 Generating mock observations of MUFASA galaxies 85

2.4.4 Recovering SFHs for MUFASA quenched galaxies 86

2.5 UltraVISTA fitting and sample selection 92

2.6 The star-formation histories of UltraVISTA quenched galaxies 95

2.6.1 When did quenched galaxies form their stellar masses? 97

2.6.2 How long did the process of quenching take? 101

2.7 Conclusions 109

3 How to measure star-formation histories with parametric models 113

3.1 Introduction 113

3.2 Parametric star-formation-history models 116

3.2.1 Exponentially declining 117

3.2.2 Delayed exponentially declining 118

3.2.3 Lognormal 118

3.2.4 Double power law 119

3.3 Priors on physical parameters 119

3.3.1 Fiducial models at redshift zero 121

3.3.2 The effects of changing prior probability densities 125

3.3.3 The effects of changing the observed redshift..... 126

3.4 Effects of different priors for the lognormal and DPL models..... 127

3.4.1 The lognormal model..... 128

3.4.2 The double power law model 128

3.5 Testing parametric models with mock observations 131

3.5.1 Generating and fitting mock data 131

3.5.2 Recovery of SFHs and physical parameters 135

3.5.3 Comparisons between SFH models..... 137

3.5.4 Broad-band photometry as a tool for understanding SFHs 138

3.6 Testing parametric models with observational data 140

3.6.1 GAMA data and fitting methodology 141

3.6.2 Inferred stellar masses 143

3.6.3 Inferred star-formation rates 143

3.6.4 Inferred mass-weighted formation times 145

3.7 Non-parametric SFH models..... 147

3.8 Conclusion 149

4 The star-formation histories of VANDELS massive quiescent galaxies at $1.0 < z < 1.3$ 153

4.1 Introduction 153

4.2 VANDELS data and sample selection 157
4.2.1 Photometric catalogues and parent sample 157
4.2.2 VANDELS spectroscopic observations 158
4.2.3 The $1.0 < z < 1.3$ mass-complete sample 159
4.3 Physical model 161
4.3.1 Stellar population model 161
4.3.2 Dust attenuation model 162
4.3.3 Nebular emission model 162
4.4 Combining spectroscopic and photometric data 162
4.4.1 Historical approaches to spectral fitting 163
4.4.2 A simplistic approach to fitting the joint datasets 165
4.4.3 Modelling spectroscopic systematic uncertainties 169
4.4.4 Final fitting of the joint datasets 172
4.5 Results 174
4.5.1 Quiescent and green-valley sub-samples 174
4.5.2 Stellar mass vs formation redshift 175
4.5.3 Distribution in D4000 vs Hδ 177
4.5.4 Trends with rest-frame UVJ colours 179
4.5.5 Oxygen II emission properties 183
4.5.6 Post-starburst and rejuvenated galaxies 184
4.5.7 Evidence of AGN activity 185
4.5.8 Stacking analyses 187
4.6 Discussion 189
4.6.1 The stellar mass vs stellar age relationship 189

4.6.2 Connecting green-valley, post-starburst and quiescent galaxies ... 194

4.6.3 The SFHs of massive quiescent galaxies at $1.0 < z < 1.3$ 201

4.7 Conclusion 204

5 Conclusions and future work 207

5.1 Thesis conclusions 207

5.1.1 The BAGPIPES code 207

5.1.2 Parametric star-formation-history models 208

5.1.3 The star-formation histories of massive quiescent galaxies 209

5.2 Future work 211

5.2.1 Resolved studies of high-redshift quiescent galaxies with JWST .. 211

5.2.2 The stellar mass-metallicity relationship for VANDELS galaxies . 215

5.2.3 Re-analysis of archival data from SDSS, Lega-C and VANDELS . 220

Bibliography 221

List of Figures

(1.1) Cosmic microwave background map from Planck 6

(1.2) Clustering of dark matter in the Millennium Simulation 8

(1.3) 2dFGRS map of galaxy clustering in the local Universe 9

(1.4) A timeline of important events in cosmic history 10

(1.5) Filter response curves for the JWST NIRCam instrument 16

(1.6) Limiting sensitivities for the JWST NIRSpec instrument 17

(1.7) The Hertzsprung-Russell diagram 20

(1.8) Schematic diagram of stellar population synthesis 22

(1.9) Common dust extinction and attenuation laws 25

(1.10) Galaxy spectral model processed through CLOUDY 27

(1.11) Schematic diagram of the nested sampling process 38

(1.12) Evolution of the galaxy stellar-mass function since $z = 3$ 41

(1.13) Stellar mass assembly across cosmic time 42

(1.14) Stellar-mass to halo-mass ratios for central galaxies at $z = 0$ 43

(1.15) The history of the cosmic star-formation-rate density 45

(1.16) The star-forming sequence from $0 < z < 2.5$ 46

(1.17) The galaxy colour bimodality 48

(1.18) The rest-frame UVJ colour selection diagram 49

(1.19) The fundamental metallicity relationship 50

(2.1) An example of a model galaxy built with Bagpipes 63

(2.2) Bagpipes models on the BPT diagram 68

(2.3) Input and Hα-derived SFRs for Bagpipes models 70

(2.4) Example of fitting mock photometry with Bagpipes 74

(2.5) Example of the age-metallicity-dust degeneracy 75

(2.6) Selection criteria for quiescent galaxies 82

(2.7) Diagram of the scheme for describing galaxy SFHs 84

(2.8) SFH recovery for an example object from Mufasa 87

(2.9) Recovery of Mufasa SFHs using the tau model 89

(2.10) Recovery of Mufasa SFHs using the DPL model 90

(2.11) UVJ diagram for UltraVISTA galaxies 93

(2.12) Number of galaxies in mass and redshift bins 94

(2.13) Quenching and formation redshifts for UltraVISTA galaxies 96

(2.14) Median formation redshifts in mass and redshift bins 99

(2.15) Quenching timescales in mass and redshift bins 102

(2.16) Quenching timescales with the prior distribution 106

(3.1) Mass-weighted formation time, sSFR and SFH shape priors 120

(3.2) Effects of alternative τ priors on the tau SFH model 123

(3.3) Effects of moving to higher redshifts on the tau SFH model 124

(3.4) Effects of the $t_{\rm max}$ prior on the lognormal SFH model 127

(3.5) Effects of the τ and β priors on the DPL SFH model 129

(3.6) SEDs and SFHs for the mock catalogue introduced in Section 3.5.1 ... 130

(3.7) Recovery of mock SFHs using the four parametric models 132

(3.8) Physical parameter recovery from mock data 133

(3.9) Posterior predictions for photometry compared to input 134

(3.10) Bayesian evidence plots for each SFH model fitted to mocks 138

(3.11) SFRs derived from SED fitting and from Hα fluxes 142

(3.12) Redshift evolution of cosmic SFRD using different models 145

(3.13) Priors imposed by non-parametric SFH models 147

(4.1) An example object fitted using different methods 163

(4.2) Physical, noise and calibration models fitted to an example object 168

(4.3) Redshifts of formation and quenching for the quiescent sub-sample ... 176

(4.4) Distribution of the VANDELS galaxies in $EW_{H\delta}$ vs D4000 178

(4.5) The VANDELS sample on the UVJ diagram 180

(4.6) Distribution of Oxygen II emission in the VANDELS sample 182

(4.7) Stacked spectra for the green-valley and quiescent sub-samples 188

(4.8) A comparison of quiescent galaxy formation redshifts 189

(4.9) A comparison of formation redshifts from photometric studies 192

(4.10) Dust attenuation vs nSFR for galaxies in the VANDELS sample 195

(4.11) UVJ tracks for four representative galaxy SFHs 197

(4.12) SFHs inferred for the three spectroscopic PSBs 200

(4.13) Posterior median SFHs for the quiescent sub-sample 202

(5.1) Resolved properties of a massive quiescent galaxy from MUFASA 212

(5.2) The stellar mass-metallicity relationship for VANDELS galaxies 216

(5.3) Example KMOS follow-up observation of a VANDELS target 217

(5.4) Distribution of VANDELS quiescent galaxies in CDFS and UDS 218

List of Tables

(2.1) Limiting magnitudes for the Mortlock et al. (2017) catalogue..... 61

(2.2) Input parameters for the BAGPIPES model shown in Fig. 2.1..... 62

(2.3) Model used to fit the UltraVISTA catalogue..... 88

(3.1) Parameters and priors for each SFH model..... 117

(3.2) A comparison of $z \sim 0.05$ SFRD and SMD estimates..... 140

(4.1) Parameters and priors for the model fitted to the VANDELS data.... 160

Chapter 1

Introduction

Whilst undoubtedly observed since ancient times, the first written record of the existence of external galaxies dates to the turn of the first millennium. In al-Sufi (ca. 964), the author notes the existence of a "small cloud" in the constellation of Andromeda, apparently well-known to Persian astronomers of the time, as well as a larger, brighter cloud known to astronomers further South. The nature of these clouds remained unknown until the 18^{th} Century, by which time approximately a thousand had been identified. Wright (1750) and Kant (1755) speculated that a connection might exist between our own Milky Way and certain external clouds, with each being a compact disk of stars, or island Universe. Confirmation of this theory was however delayed until the early 20^{th} Century, when the application of early photometric and spectroscopic techniques to external galaxies began to allow the characterisation of their extreme distances and recession velocities (e.g. Slipher 1913, 1917; Curtis 1917). These observations, along with contemporaneous theoretical advances (e.g. Einstein 1917; Friedmann 1922), eventually gave rise to our modern view of the Universe as composed of a large number of separate galaxies, distributed within an expanding medium.

As well as representing a new cosmological model, this discovery gave rise to an entirely new field of study: the formation and evolution of galaxies. One of the earliest and most fundamental results in this field was the realisation by Hubble (1926) that galaxies could be split broadly into two categories. Elliptical galaxies, now known to be typically massive and red in colour, are distinct from bluer, less-massive spiral systems. The key difference is an absence of ongoing star-formation (quiescence), which causes some galaxies to appear red in colour.

This thesis aims to contribute to our understanding of the processes that drive the formation and evolution of galaxies, as well as the origin of the divide within the galaxy population. In particular, I will present analyses designed to constrain the evolutionary pathways of massive quiescent galaxies by inferring their star-formation histories (SFHs). In this introduction I present the background information necessary to place my work in context. I will begin by giving a brief overview of our current Λ-CDM cosmological model and the cosmic history it implies. I will then give a brief description of observational techniques in astronomy, and details of the important physical ingredients of galaxy spectra. I will then describe the Bayesian statistical and computational techniques that are employed in this thesis to fit models to observational data. I will finally discuss current observational results relating to the evolution of the galaxy population across cosmic time, and the theories that have been developed to explain this evolution.

1.1 Cosmic history under the Λ–CDM model

It was quickly proposed that the expansion of the Universe (e.g. Hubble 1929) implied that it has a finite age and single point of origin (e.g. Lemaître 1927), although this was much contested for several decades. This later became known as the Big Bang theory, and was confirmed in the 1960s with the discovery of the cosmic microwave background (see Section 1.1.3). Modern estimates place the current age of the Universe, otherwise known as the Hubble time, at ~ 13.8 Gyr (Planck Collaboration et al. 2018). Over the last century, Big Bang cosmology has been refined and developed into our current Λ-CDM model, which describes the origins of matter in primordial nucleosynthesis, the role of cold dark matter (CDM) as the backbone of large-scale structure, and the accelerating expansion of the Universe due to dark energy. This Λ-CDM cosmological model is the backdrop against which the formation and evolution of galaxies is set.

1.1.1 The expanding Universe and distance measurements

The discovery that the Universe is expanding, as well as potentially having a non-Euclidean geometry, complicates the definition of the distances between objects. Because of the finite speed of light, c, distant objects are observed as they were

at an earlier time in cosmic history, when the distances between objects were compressed relative to the present time. Additionally, the expansion of space stretches the wavelength of light as it travels, shifting it towards the red end of the spectrum. This gives distant galaxies an apparent recession velocity, v, due to the expansion of the intervening space.

To relate distances at different cosmic epochs, we define the scale factor, $a(t)$, which is equal to 1 at the present time, and < 1 at earlier times. This factor relates the present-day distance between objects, x, often called the comoving distance, to the distance at earlier times, r, known as the proper distance, by

$$r = ax. \tag{1.1}$$

The recession velocity at the present time is related to the comoving distance by Hubble's constant, H_0, such that

$$v = H_0 x. \tag{1.2}$$

We can likewise treat the stretching of light by introducing the redshift, z, which relates the emitted and observed wavelengths, $\lambda_{\rm em}$ and $\lambda_{\rm obs}$, by

$$z = \frac{\lambda_{\rm obs} - \lambda_{\rm em}}{\lambda_{\rm em}}. \tag{1.3}$$

Since the wavelength of light stretches in proportion to the space through which it travels, redshift is also related to the scale factor by

$$a = \frac{1}{(1+z)}. \tag{1.4}$$

Given these relationships, the only missing ingredient is the expansion history of the Universe, which is contained within the time-evolution of the scale factor. This expansion history is described by the Friedmann equation (e.g. Liddle 2003)

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} \tag{1.5}$$

where $\dot{a}$ is the time-derivative of the scale factor, G is the gravitational constant,

ρ is the energy density of the Universe, and k is a constant. The value of k corresponds to the spatial curvature of the Universe, which is determined by its energy density. Current evidence suggests that the total energy density of our Universe corresponds to the critical density, ρ_c, required for the Universe to be spatially flat, with $k = 0$ (e.g. Planck Collaboration et al. 2018).

The energy density of the Universe, which is assumed to be both homogeneously and isotropically distributed on large scales, is made up of several components, namely radiation, ρ_r, matter, ρ_m, and dark energy, ρ_Λ. We define the density parameter, Ω, as the sum of these components, divided by the critical density. The density parameter for each component is the ratio of the energy density in that component to the critical density, for example

$$\Omega_m = \frac{\rho_m}{\rho_c}. \tag{1.6}$$

We can now find the comoving distance to a given redshift by first considering the distance travelled by light in an expanding medium, given by

$$x = \int \frac{c\ dt'}{a(t')} \tag{1.7}$$

and then changing variables to give (e.g. Hogg 1999)

$$x(z) = \int_0^z \frac{c\ dz'}{H(z')}. \tag{1.8}$$

Additionally, it follows from Equation 1.4 that the lookback time to a given redshift is

$$t(z) = \int_0^z \frac{dz'}{(1+z')\ H(z')}. \tag{1.9}$$

Two other important measures of cosmological distance are the angular diameter distance, d_A, and the luminosity distance, d_L. The angular diameter distance is the distance one would obtain by comparing the physical size of an object to its angular diameter on the sky, whereas the luminosity distance is the distance one would obtain by comparing the intrinsic luminosity of a source to the apparent flux observed (k-corrected for redshifting of the light). In a spatially flat Universe,

the angular diameter distance is given by

$$d_A = \frac{x(z)}{(1+z)} \tag{1.10}$$

whereas the luminosity distance is given by

$$d_L = (1+z)\,x(z). \tag{1.11}$$

Throughout this thesis I will assume a spatially flat Universe with $\Omega_M = 0.3$, $\Omega_\Lambda = 0.7$ and $H_0 = 70$ km s^{-1} Mpc^{-1}. All cosmological calculations are carried out using the Astropy Python module (Astropy Collaboration et al. 2013).

1.1.2 Primordial nucleosynthesis

Approximately one second after the Big Bang, the Universe contained a plasma of free protons, neutrons and electrons, all of which were strongly coupled to high-energy thermal radiation. The details of the evolution that took place at earlier times, including baryogenesis and a possible inflationary phase, are still highly speculative and beyond the scope of this thesis. As the Universe expanded it cooled, changing the relative speeds of different reaction pathways between the species of matter present. In particular, when the temperature became low enough that neutrons and electrons no longer spontaneously combined to form protons, the instability of free neutrons (which have a half-life of ~ 611 seconds) caused the neutron-to-proton ratio to begin falling. This fall was quickly arrested by the formation of atomic nuclei, within which neutrons are stable. This process, first postulated by Alpher et al. (1948), means that the elemental abundance pattern of baryonic matter in the Universe is fixed from a very early stage to ~ 75 per cent hydrogen and ~ 25 per cent helium by mass, with only trace quantities of heavier elements (e.g. lithium and beryllium).

1.1.3 Recombination and the cosmic microwave background

As the Universe continued to expand and cool, the thermal radiation that fills it continued to fall in energy. When the average photon energy fell significantly below that required to ionize hydrogen, the reaction rate between the thermal

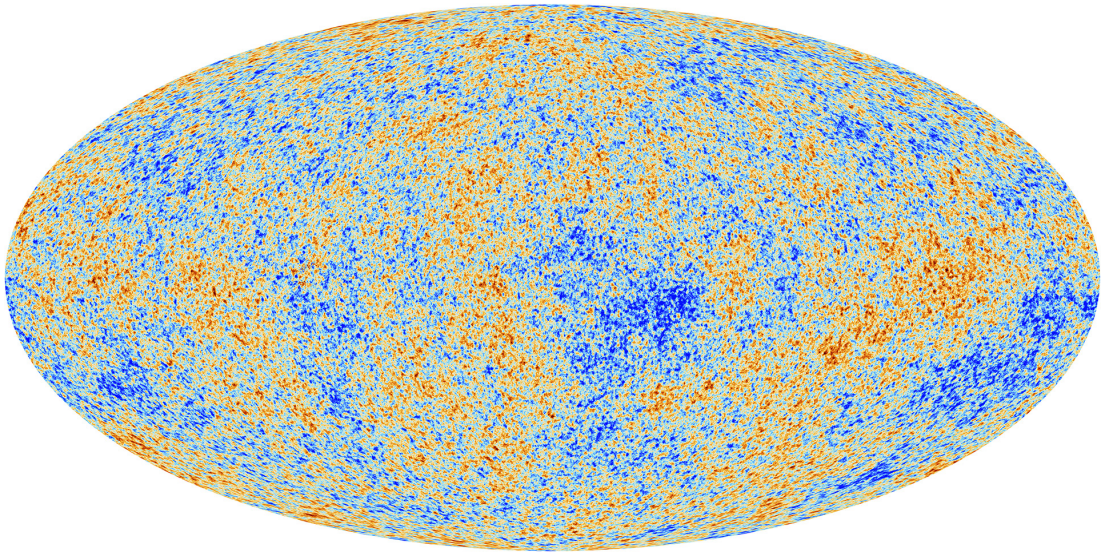

Figure 1.1 *A map of temperature fluctuations in the cosmic microwave background from the Planck satellite (e.g. Planck Collaboration et al. 2018). Redder colours indicate hotter temperatures and bluer colours cooler temperatures. The variations have a root-mean-squared amplitude of 18μK, demonstrating that the Universe was highly uniform ~380,000 years after the Big Bang at $z \sim 1000$.*

photons and matter rapidly fell to near zero. This allowed the thermal radiation to propagate freely, gradually being redshifted to lower energies by the expansion of the Universe. This is often referred to as the time of last scattering. Predicted by Alpher & Herman (1948), this background radiation, now redshifted into the microwave regime, was first detected by Penzias & Wilson (1965).

Today we can make extremely precise measurements of the cosmic microwave background (CMB). An example CMB map from the Planck satellite (e.g. Planck Collaboration et al. 2018) is shown in Fig. 1.1. The CMB is the most perfect black-body spectrum ever observed, with flux per unit frequency, ν, given by

$$B_\nu(\nu, T) = \frac{2h\nu^3}{c^2\left(\exp\left(\frac{h\nu}{kT}\right) - 1\right)} \tag{1.12}$$

where h is Planck's constant and T is the temperature of the radiation field. The CMB has an average temperature of ~ 2.723 K, which is observed to be extremely uniform over the whole sky, with the variations shown in Fig. 1.1 having a root-mean-squared amplitude of 18μK.

The time of last scattering for the CMB radiation, at $z \sim 1000$, was approximately

380,000 years after the Big Bang. The extreme uniformity of the CMB radiation demonstrates that the Universe contained no significant structure at this epoch. Our ability to directly measure the CMB gives us strong constraints on the properties of the Universe at this very early time, upon which attempts to model the later formation of structure can be founded.

The time of last scattering between the CMB photons and the baryonic matter in the Universe is often known as recombination. This is because, once the thermal radiation in the Universe can no longer ionize hydrogen, the atomic nuclei formed during primordial nucleosynthesis combine with electrons to form a neutral gas. This primordial hydrogen-helium gas is the material from which the first stars and galaxies were formed, with heavier elements, commonly known in astronomy as metals, formed later by stellar nucleosynthesis processes (see Section 1.3.1).

1.1.4 Dark matter and large-scale structure

The total matter density of the Universe implied by CMB measurements is significantly higher than the total density of baryonic matter observed in galaxies and the intergalactic medium (IGM). This implies that the majority of the matter in the Universe is of a different kind, which, since it cannot be seen, does not interact electromagnetically. Even before precision measurements of the CMB, evidence for the existence of this "dark matter" had been observed as early as the 1930s in the rotation curves of galaxies and velocity dispersions of galaxy clusters (e.g. Zwicky 1937, Babcock 1939, Oort 1940). By the 1970s strong evidence was available from spiral galaxy rotation curves (e.g. Rubin & Ford 1970), with the nature of dark matter firmly established as a major unsolved problem in physics.

The main importance of dark matter in the formation and evolution of galaxies is its role in the growth of cosmic structure. As dark matter does not interact electromagnetically, it readily collapses under gravity and clusters together to form dark matter halos. It was quickly realised that, since dark matter dominates the matter density of the Universe, gravitational interactions dominate the formation of cosmic structure. Early attempts to model these interactions with N-body simulations (e.g. Press & Schechter 1974; White & Rees 1978) suggested that structure forms in a hierarchical fashion, with smaller structures collapsing earliest and gradually being incorporated into larger and larger halos. Fig. 1.2 shows an example of the hierarchical clustering of dark matter halos in the Millennium Simulation (Springel et al. 2005) on a variety of scales.

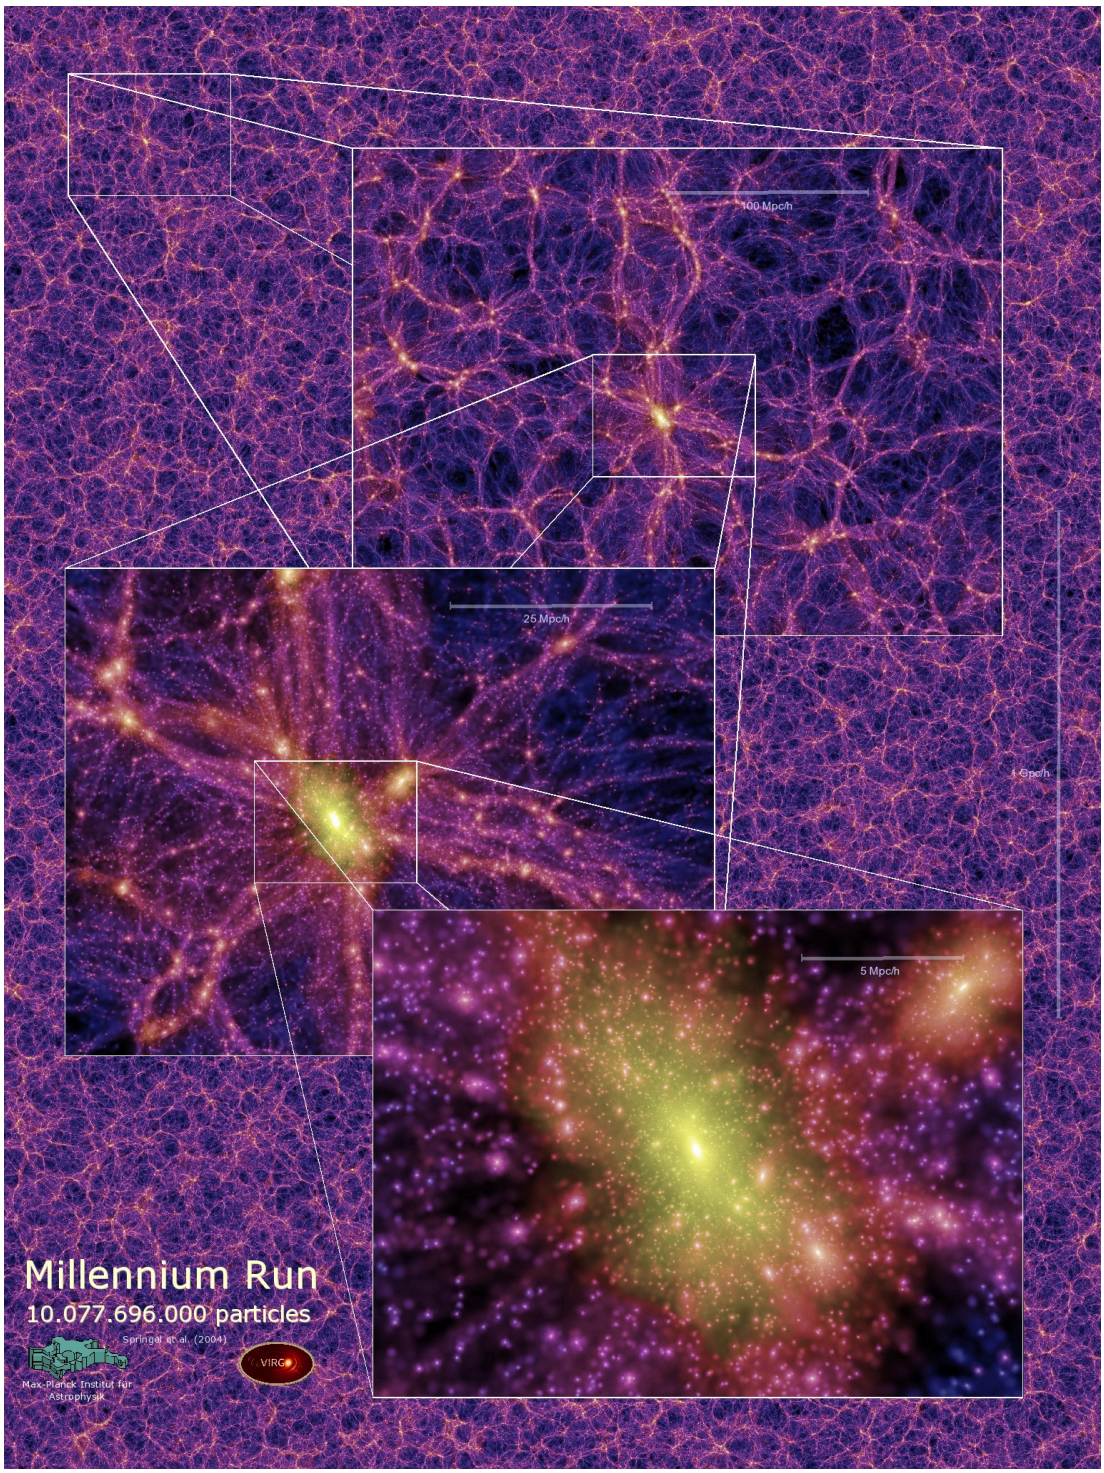

Figure 1.2 *The clustering of dark matter in the Millennium Simulation (Springel et al. 2005). This pattern is often referred to as the cosmic web. The clustering is hierarchical, beginning on smaller scales and progressing towards larger scales at later times. Gravitationally bound dark matter halos form at nodes in the cosmic web. On large scales, the distribution of dark matter halos in simulations agrees well with the distribution of galaxies in the Universe, indicating that galaxies form within the gravitational potential wells of dark matter halos.*

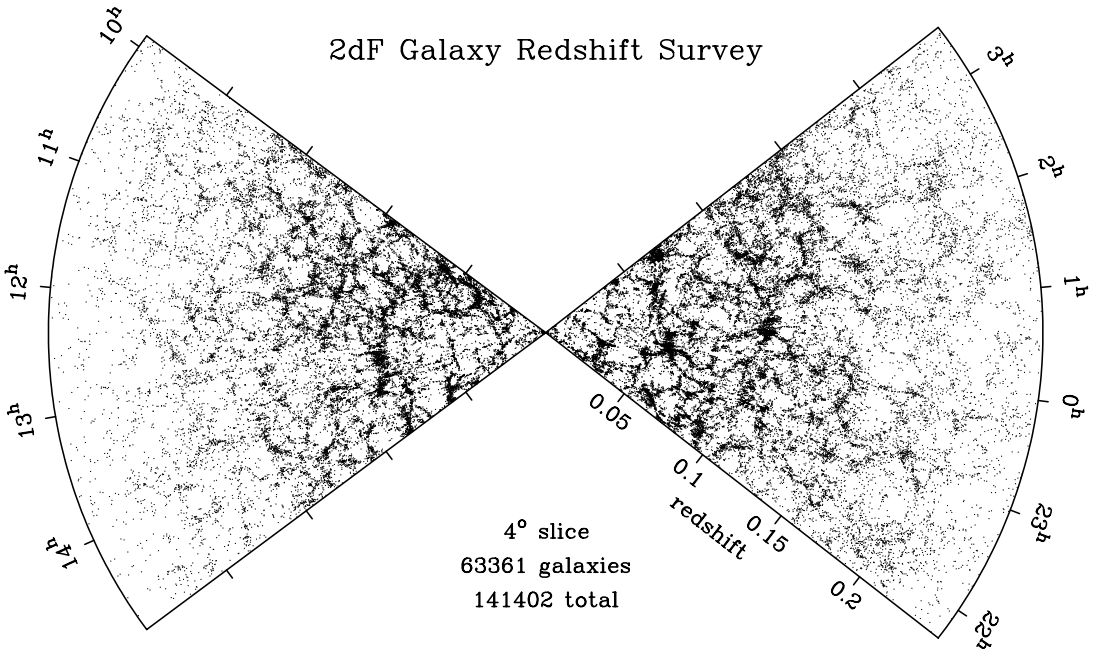

Figure 1.3 *A map of the clustering of galaxies in the local Universe from the Two-degree Field Galaxy Redshift Survey (2dFGRS; taken from Peacock et al. 2001). On large scales, the statistical properties of galaxy clustering agree well with the clustering of dark matter halos in N-body simulations (Fig. 1.2). This shows that dark matter forms the backbone of large-scale structure in the Universe.*

By measuring the redshifts of large numbers of galaxies, their clustering properties can be determined. Fig. 1.3 shows an example of galaxy clustering in the local Universe from the Two-degree Field Galaxy Redshift Survey (2dFGRS; Colless 1999). On large scales ($\gtrsim 10$ Mpc), the clustering properties of galaxies agree well with the clustering of dark matter halos visible in Fig. 1.2. This indicates that dark matter halos form the gravitational potential wells into which the neutral hydrogen-helium gas produced at recombination collapses to form galaxies.

It is also worth noting that significant departures are observed from the clustering predicted by dark matter simulations on certain scales due to baryonic processes. For example, an excess of clustering is observed on ~ 100 Mpc scales due to baryon acoustic oscillations (BAO; e.g. Eisenstein et al. 2005). These are sound waves in the primordial plasma that are frozen in place at recombination, when the sound speed rapidly falls, and can be used as an additional constraint on cosmological parameters. On scales similar to and smaller than the sizes of galaxy clusters ($\lesssim 10$ Mpc), baryonic processes associated with galaxy evolution come to dominate the clustering of matter in the Universe.

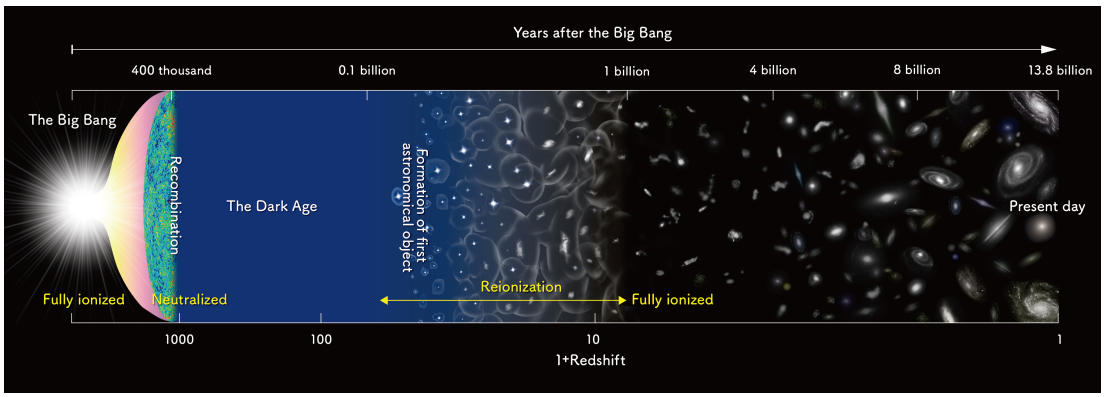

Figure 1.4 *A timeline of important events in cosmic history. Following recombination, dark matter halos begin to collapse, with baryonic matter then collapsing into these halos to form the first galaxies. Ionizing photons again begin to be produced by baryonic processes such as star-formation and AGN activity, which reionize the Universe over* ~ 1 *Gyr, ending at* $z \sim 6$*. (Credit: NAOJ)*

1.1.5 The first galaxies and the reionization of the intergalactic medium

Once dark matter halos begin to become established, the primordial neutral hydrogen-helium gas produced at recombination begins to collapse into them, eventually leading to the formation of the first stars and galaxies. The period between recombination and the formation of the first stars is often referred to as the cosmic dark age, as no sources of light existed, meaning no new photons were produced. Fig. 1.4 shows a timeline of important events in cosmic history, including the cosmic dark age.

The exact time at which the cosmic dark age ended is still a matter of debate, with expanding observational capabilities continuing to extend our knowledge of galaxy formation back to earlier cosmic times. The existence of galaxies within the first billion years of cosmic history was first demonstrated indirectly by Dunlop et al. (1996) and Peacock et al. (1998), who observed massive galaxies with old stellar populations ($\sim$ 3 Gyr) already in place by $z \sim 1.5$. Since approximately the beginning of the 21$^{\rm st}$ Century, it has been possible to directly detect galaxies within the first billion years of cosmic history, at $z \gtrsim 5$ (e.g. Dey et al. 1998; Spinrad et al. 1998; Hu et al. 1999; Ellis et al. 2001; see Dunlop 2013 for a review), with the redshift frontier pushed to $z \sim 10$ within the last decade (e.g. McLure et al. 2010, 2011; Ellis et al. 2013; McLeod et al. 2015, 2016). In parallel, active galactic nuclei (AGN) have now been detected with extremely massive

black holes ($\gtrsim 10^9$ $M_\odot$) at $z > 6$ (e.g. Fan et al. 2006; Mortlock et al. 2011; Carnall et al. 2015; Bañados et al. 2018), suggesting that galaxy formation was already underway at $z \sim 15 - 20$, when the Universe was ~ 200 Myr old.

Gas accreted by dark matter halos must first cool in order to collapse and form stars. Cooling of gas in local galaxies is mostly achieved by radiation of energy during atomic transitions in metal atoms, however this mechanism was not available to the pristine hydrogen-helium gas from which the first stars, often called population III stars, formed. Cooling must instead have been dominated by the atomic and molecular transitions of hydrogen, and, as a consequence, the first stars are thought to have had typically much larger masses than later generations ($\sim 10 - 100$ $M_\odot$; e.g. Loeb 2010; Bromm & Yoshida 2011).

As well as ending the cosmic dark age, these extremely massive population III stars, along with succeeding stellar generations and the first AGN, trigger a final phase transition for the baryonic matter in the Universe. Radiation produced by these sources contains the first hydrogen-ionizing photons to exist since recombination, and consequently, as galaxy formation begins, the primordial gas that remains in the IGM is reionized.

The absorption of hydrogen-ionizing photons by a substantially neutral intergalactic medium at high redshift, leading to troughs in the observed spectra of high redshift objects blue-wards of Lyman-α, was first predicted by Gunn & Peterson (1965). This effect, which was first observed in the spectra of $z \sim 6$ quasars (e.g. Becker et al. 2001), allows us to constrain the epoch at which reionization took place. Contemporary studies, which combine systematic surveys of the high redshift galaxy population with the optical depth to Thomson scattering for CMB photons, suggest that the IGM was primarily reionized by star-forming galaxies between $6 < z < 10$ (e.g. Robertson et al. 2013, 2015).

1.1.6 Dark energy and accelerating expansion

As well as methods that use the CMB, the expansion rate of the Universe can be determined by using a series of local indicators to build up an independent estimate of the distances to extragalactic objects. The first step on this cosmic distance ladder is the measurement of the Earth-Sun distance, or astronomical unit, historically using geometric methods, or more recently using radar observations. Distances to nearby stars can then be obtained by using

Earth's motion around the Sun to measure parallaxes. From this point a variety of indicators can be used, for example period-luminosity relationships for variable stars (e.g. Cepheid variables; Leavitt 1908), and surface brightness fluctuations in local galaxies. Many such methods rely on astronomical objects with fixed peak luminosities, referred to as standard candles.

The brightest class of standard candles are type Ia supernovae (e.g. Kowal 1968), which allow us to extend the distance ladder to cosmological distances. The first such analyses at $z \gtrsim 0.1$ found that distant supernovae were fainter than expected (Riess et al. 1998; Perlmutter et al. 1999), suggesting that the expansion of the Universe has begun to accelerate since $z \sim 0.5$. This is evidence for the final component of our Λ-CDM cosmological model, dark energy. Sometimes called the cosmological constant, current evidence suggests that dark energy takes the form of a homogeneously distributed, time-independent energy density that is a property of otherwise-empty spacetime.

1.2 Astronomical Observations

The vast sizes of galaxies and the extremely long timescales over which they evolve preclude meaningful attempts to understand them through experimentation. Instead we must rely entirely on the radiation emitted by distant galaxies to probe their formation and evolution processes. The first major step towards modern observational techniques came at the beginning of the 17^{th} Century, with the manufacture of the first telescopes, allowing observations to move beyond the sensitivity of the human eye. Further major developments came with the application of photographic techniques to astronomy in the 19^{th} Century and the use of charge-coupled devices (CCDs) from the early 1980s. As well as allowing observations to be recorded, these developments allowed us to begin extending our knowledge beyond the optical wavelength range.

Modern technology now allows us to extend our observations across the electromagnetic spectrum from gamma rays to the radio regime, as well as to progressively fainter objects with ever-larger telescopes. The 21^{st} Century promises further progress, as we move into the era of multi-messenger astronomy, with additional information available through the study of cosmic rays, neutrinos and gravitational waves. In addition, the advent of 30-metre class telescopes on the ground and the 6.5-metre *James Webb Space Telescope* (JWST) will allow us

to see far deeper into the Universe in the 2020s than ever before. In this section I will outline relevant information on the methods by which astronomical studies of galaxies are conducted.

1.2.1 Photometry vs spectroscopy

Astronomical observations are typically taken in one of two ways, often referred to as photometric and spectroscopic observations. For photometric observations, light is collected and passed through a filter that selects a certain wavelength range. For spectroscopic observations, light is dispersed by a prism or grating (or combination of the two, referred to as a grism), such that the intensity of light can be recorded as a function of wavelength.

Photometric observations typically allow all galaxies within the field of view of the telescope to be observed simultaneously, although, as described in Section 1.2.2, a lack of sufficient spatial resolution can cause problems with the blending of sources. Conversely, spectroscopic observations could historically only be directed at single objects. The development of multi-object spectroscopy, beginning in the 1980s, has greatly expanded our ability to conduct large spectroscopic surveys, using fibre-fed, multi-slit or integral-field units.

1.2.2 Spatial and spectral resolution

When measuring the light from galaxies, several important considerations determine how much information can be obtained. Firstly, it is desirable to be able to observe with high spatial resolution, corresponding to small angular scales on the sky. In the first instance this is important for the separation of galaxies that are either physically close, or projected along nearby lines of sight. Even higher resolution allows us to resolve different regions in individual galaxies, affording greater insights into galaxy evolution processes. In the absence of atmospheric turbulence angular resolution is determined by diffraction, which is proportional to the diameter of the telescope collecting area. The minimum angle that can be resolved, $\theta_{\rm min}$, is given by the Rayleigh criterion (Rayleigh 1879), which, for a circular aperture of diameter D, is given by

$$\theta_{\rm min} \;=\; 1.22\,\frac{\lambda}{D} \tag{1.13}$$

where λ is the wavelength of the light being observed. For observations using ground-based telescopes, angular resolution is also limited by turbulence, or seeing, in the atmosphere (typically at levels of $\sim 1''$), although progress is now being made in adaptive optics systems that can help to correct for this.

A second major consideration is spectral resolution. This is the minimum separation in wavelength, $\Delta\lambda$, that can be distinguished. Spectral resolution is often expressed in terms of the resolving power, R, given by

$$R = \frac{\lambda}{\Delta\lambda}. \tag{1.14}$$

In general, higher resolving power allows more detailed information to be extracted from galaxy spectra. An important threshold exists at $R \sim 500 - 1000$, with observations at higher resolution able to distinguish emission and absorption features due to individual atomic species in the optical wavelength range, allowing detailed galaxy physical properties to be determined (e.g. Gallazzi et al. 2005).

1.2.3 Fluxes and magnitudes

When measuring light from an astronomical source, the quantity of interest is the apparent flux observed as a function of wavelength. The flux from a source, f, is defined as the energy received per unit area per unit time. As we are often interested in how the flux varies as a function of wavelength, it is also common to report either the flux per unit wavelength, f_λ, or flux per unit frequency, f_ν. Three commonly used non-SI units for measuring astronomical fluxes are angstroms (Å), equivalent to 10^{-10} metres; ergs, equivalent to 10^{-7} joules; and janskys (Jy), equivalent to 10^{-26} watts per metre squared per hertz, or 10^{-23} ergs per centimetre squared per angstrom.

For historical reasons, it is also common to report the fluxes from objects using the magnitude system. This relates the apparent magnitude of an object, m, to its apparent flux using a standardised zeropoint flux, f_0, which corresponds to a magnitude of $m_0 = 0$, such that

$$m = -2.5 \log_{10}\left(\frac{f}{f_0}\right). \tag{1.15}$$

Historically, f_0 was defined as the observed flux from the star Vega, however this

Vega magnitude system has now been largely superseded by the absolute, or AB magnitude system (Oke & Gunn 1983), that defines $f_0 = 3631$ Jy. AB magnitudes will be used exclusively throughout this thesis.

In practice, observed fluxes (and hence magnitudes) are averages over a range of wavelengths, weighted by the transmission function of the optical system, $T(\lambda)$. Transmission is normally deliberately limited to a relatively narrow range of wavelengths by the use of an optical filter. In the case of a photon counting device such as a CCD, the flux used in Equation 1.15 is in fact given by

$$f = \frac{\int f_\lambda \, \lambda \, T(\lambda) \, \mathrm{d}\lambda}{\int \lambda \, T(\lambda) \, \mathrm{d}\lambda}. \tag{1.16}$$

A final important quantity is the absolute magnitude, M, that relates the apparent magnitude to the intrinsic luminosity of the source. The absolute magnitude is the apparent magnitude that a source would have if it were placed at a distance of 10 parsecs (pc). A parsec is a commonly used non-SI distance measurement in astronomy, defined as the radial distance at which a transverse distance of one astronomical unit subtends an angle of 1 arcsecond on the sky. By introducing the ratio of distances into Equation 1.15, the absolute magnitude can be shown to be given by

$$M \;=\; m - 5 \, \log_{10}\!\left(\frac{d_L}{10 \text{ pc}}\right) \;+\; 5 \;-\; K(z) \tag{1.17}$$

where d_L is the luminosity distance to the object of interest (see Section 1.1.1). The factor $K(z)$ is the k-correction, which accounts for the redshifting of light relative to the range of wavelengths over which the filter is sensitive.

1.2.4 Observing galaxy populations

Typically, when conducting studies of galaxy populations, a series of photometric observations are first taken through a complimentary set of filters at different wavelengths. This allows a low-resolution spectrum to be built up for all objects in the area surveyed. Typically, the design of such surveys is a trade-off between depth, number of filters, area covered and time available (e.g. Kemp et al. 2019). Fig. 1.5 shows an example of a complementary set of filter curves through which photometric observations might be taken. In this case, the filters are those

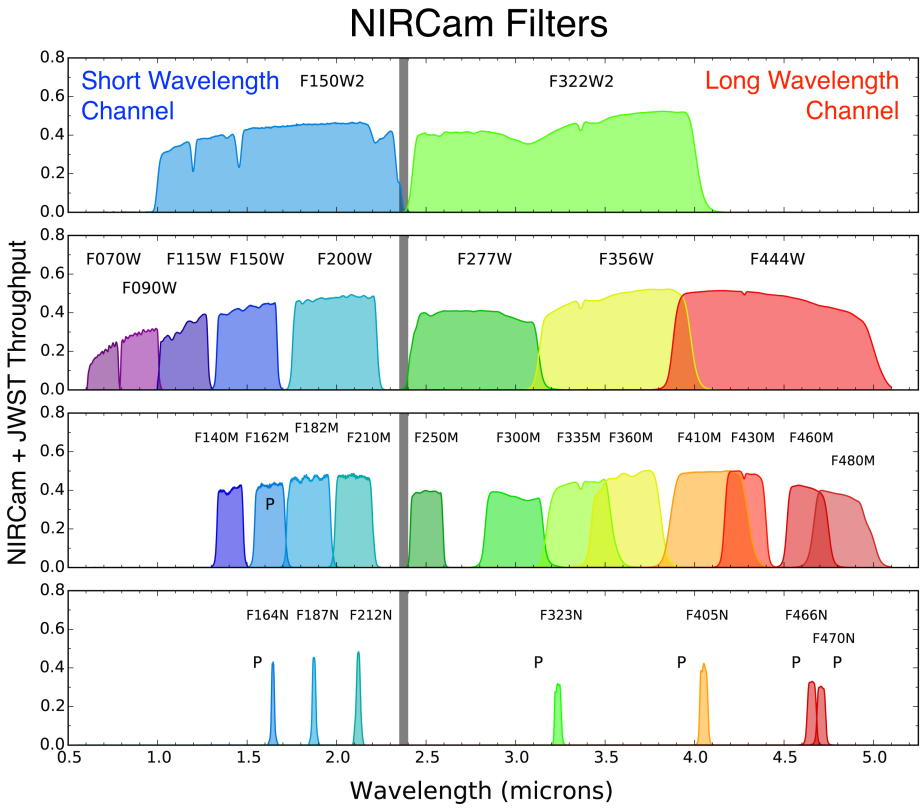

Figure 1.5 *Filter response curves for the NIRCam instrument onboard the James Webb Space Telescope (JWST). Filters are distributed across a broad wavelength range, with a variety of different resolving powers available, from extremely broad-band filters at the top to narrow-band filters at the bottom. (Credit: JWST user documentation)*

available when using the NIRCam instrument onboard JWST. Filters of different widths are available such that trade-offs between resolution and signal-to-noise ratio (SNR) can be made depending on the brightness of the galaxies of interest and time available. The narrowest filters are normally designed to isolate specific emission or absorption features.

Using the information gained from photometric observations (see Section 1.3), a subset of galaxies can be selected for spectroscopic follow-up observations in order to obtain more detailed information about objects of particular interest. This process is necessary as the integration time required to obtain a fixed SNR increases with spectral resolution, and, even with multi-object spectroscopy, it is not normally possible to observe all galaxies in a given field of view. Data obtained via photometry is usually referred to as the spectral energy distribution (SED), whereas data obtained via spectroscopy is referred to as the spectrum.

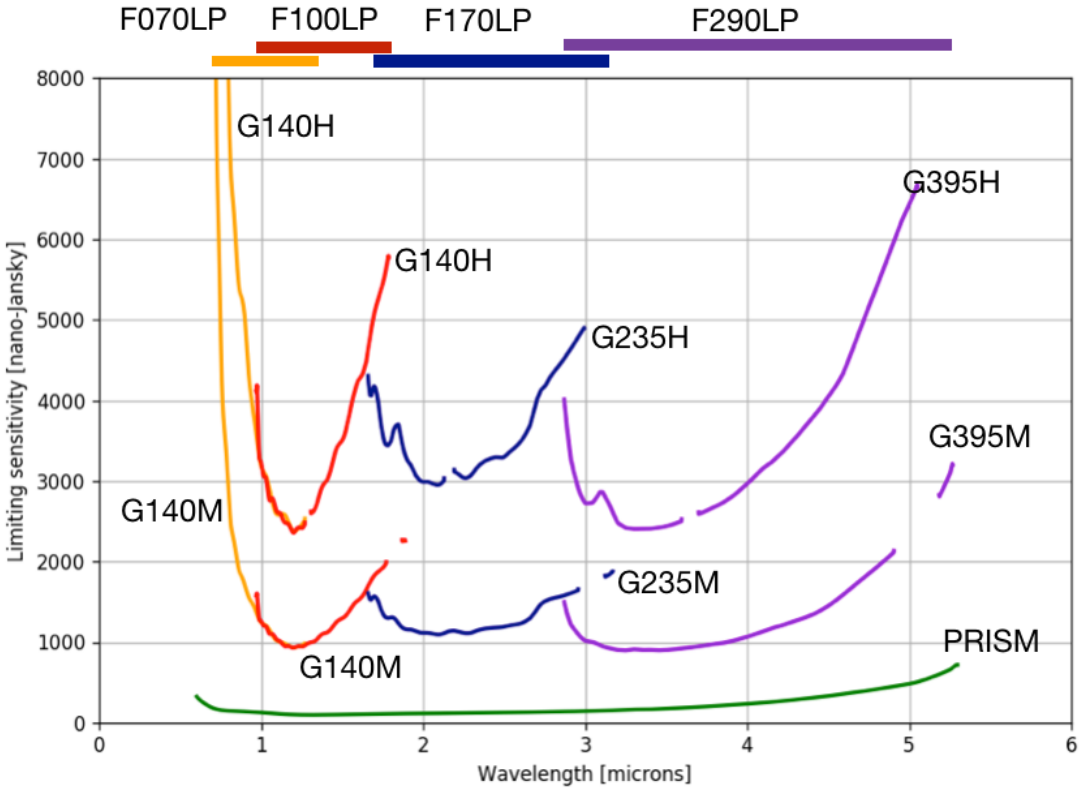

Figure 1.6 *Limiting sensitivities (at SNR = 10 per pixel) for the NIRSpec instrument onboard the James Webb Space Telescope (JWST) in multi-object spectroscopy (MOS) mode. Sensitivities assume a 963 second integration using the different available grisms and prism. H grisms are high resolution, M grisms are medium resolution, and the prism is low resolution. (Credit: JWST user documentation)*

As an example, Fig. 1.6 shows the limiting sensitivities for each of the grisms available on the JWST NIRSpec instrument at fixed integration time. The high resolution grisms (labelled H; $R \sim 2700$) have higher limiting sensitivities than the medium resolution equivalents (labelled M; $R \sim 1000$). The low resolution prism ($R \sim 100$) is more sensitive, and photometric observations using NIRCam ($R \sim 10$) would be more sensitive still, albeit providing less-detailed information.

1.2.5 Data reduction

Once observational data has been obtained, several steps are required to calibrate the signal received by the telescope. This process is usually referred to as data reduction. The main effects that must be accounted for are the throughput of the telescope and instrument (as shown in Fig. 1.5), the sensitivity function of the detector, instrumental signatures and the absolute flux calibration, which is

usually measured by observing a bright standard star with a well-known spectrum. In the case of ground-based observations, one must also account for emission and absorption processes in Earth's atmosphere.

1.3 Modelling galaxy spectra

Whilst the Λ-CDM cosmological model outlined in Section 1.1 is now well established, no such comprehensive model yet exists to describe the formation and evolution of galaxies. Broadly speaking, this is because galaxy formation is mostly driven by baryonic interactions, which are considerably more complex to model than the gravitational interactions that dominate on larger scales.

In order to build a theory of galaxy formation, we must first gain an understanding of the physical properties of the galaxy population, and how these evolve across cosmic time. The first step in this process is to obtain high quality observations of large and representative samples of galaxies across a wide redshift range. Models must then be constructed that relate the spectra (or spectral energy distributions) of galaxies to their underlying physical properties. Finally, these models must be fitted to the available observational data, in order to constrain the distribution of physical parameters in the galaxy population.

In this section I will describe the methods by which models are constructed to interpret observations of galaxy spectra. I will then discuss the fitting of these models to observational data using Bayesian statistics in Section 1.4, before moving on to describe how these methods have been used to inform our current understanding of galaxy evolution in Section 1.5.

1.3.1 Stellar population synthesis

At wavelengths shorter than $\lambda \sim 5$ μm, the spectra of galaxies are dominated by emission that comes directly from the photospheres of stars. In addition, reprocessed stellar light is also responsible for nebular line and continuum emission, as well as the dust emission that dominates at $\lambda \gtrsim 5$ μm. This means that the modelling of the stellar populations of galaxies, commonly referred to as stellar population synthesis (SPS), is the most fundamental ingredient in modelling galaxy spectra. In the following subsections I will present an overview

of SPS modelling. A detailed review is presented in Conroy (2013). Because of the complexities of SPS modelling, studies of galaxy evolution commonly use pre-computed model libraries that have been extensively tested in the literature (e.g. Bruzual & Charlot 2003; Eldridge & Stanway 2009; Falcón-Barroso et al. 2011; Conroy et al. 2018). Throughout this thesis I will use a 2016 updated version of the Bruzual & Charlot (2003) models, which, at the time of writing, are available from `http://www.bruzual.org/~gbruzual/bc03/Updated_version_2016`.

Stellar evolution and the Hertzsprung-Russell diagram

The most important factor in determining the physical and spectral properties of stars are their initial masses. Throughout the majority of their lifetimes, all stars are powered by the burning of hydrogen to produce helium. The intensity of these fusion reactions is determined by the mass, and hence a strong relationship exists between the stellar mass and stellar temperature. To first order, stellar spectra are well approximated as black bodies (see Equation 1.12), meaning that the mass, via the temperature, also strongly determines the observed spectrum.

This relationship was first recognised as a tight correlation between stellar temperature and stellar luminosity, known as the stellar main sequence. A diagram showing these quantities, commonly known as a Hertzsprung-Russell (HR) diagram (e.g. Hertzsprung 1911; Russell 1914), is shown in Fig. 1.7. The main sequence can be seen running from top-left to bottom-right, with stars of progressively higher masses further up the sequence.

Stellar spectra also exhibit different atomic and molecular absorption features depending on their temperatures, and are classified by the features present using the Morgan-Keenan system (Morgan et al. 1943). Seven classes, labelled from hottest to coolest by the letters OBAFGKM, are recognised, depending on the most prominent features present (e.g. strong Balmer absorption in A-type stellar spectra). These classes are also shown in Fig. 1.7. Each class is then further divided into 10 sub-classes, numbered from $0 - 9$. Different luminosity classes, related to stellar densities, are denoted with Roman numerals.

Four example stellar spectra are shown in the top-right panel of Fig. 1.8. It can be seen that the spectra of stars with hotter spectral types peak at shorter wavelengths, and different strong absorption features are present in each case. The strengths of absorption features due to metals also have a secondary

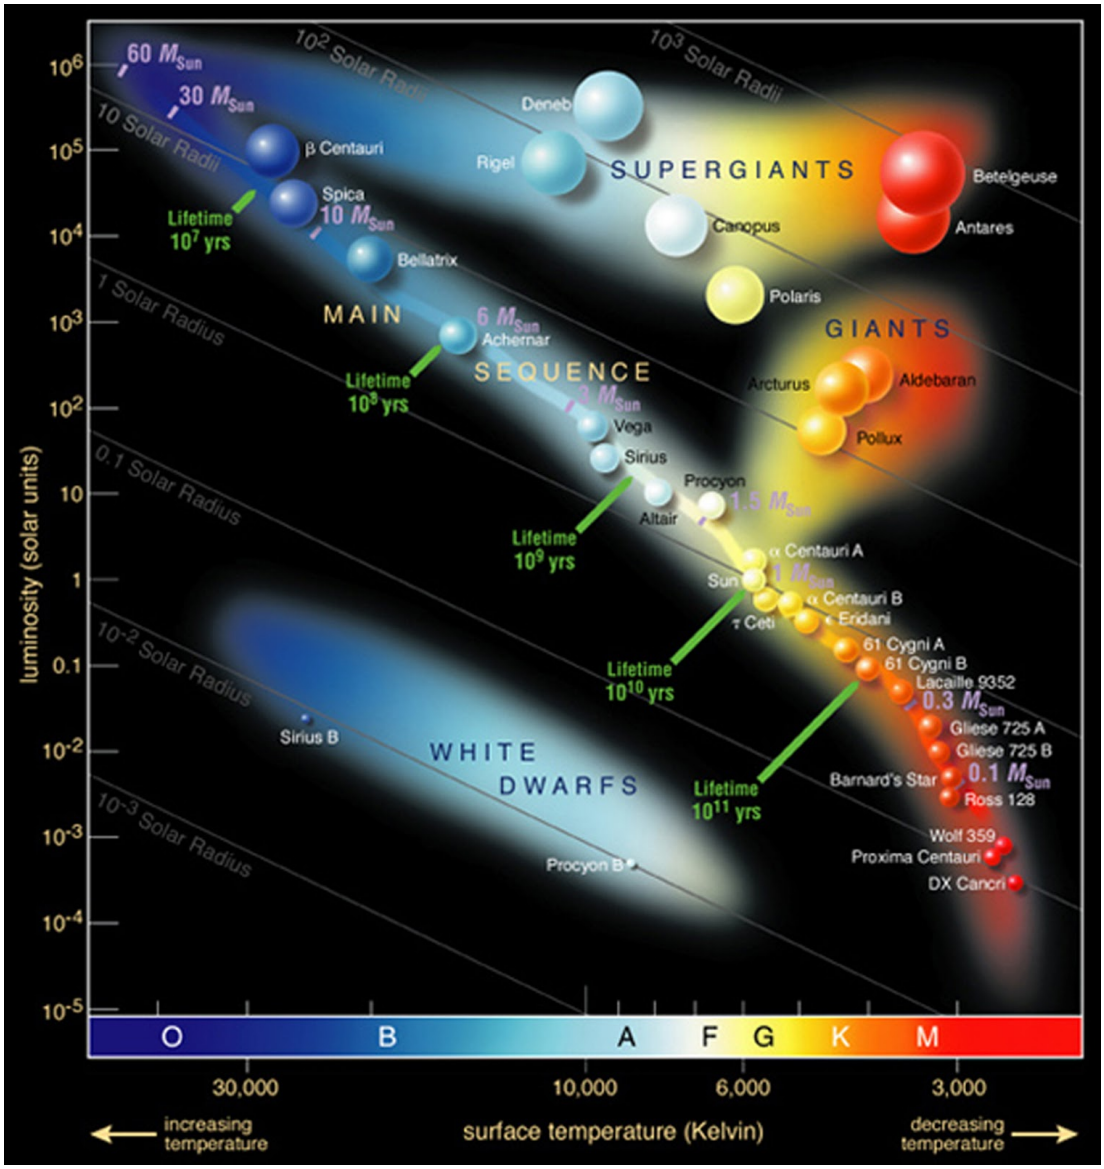

Figure 1.7 *The Hertzsprung-Russell diagram, showing the relationship between stellar temperature and stellar luminosity. The Morgan-Keenan classification scheme is also shown at the bottom. Stars spend the majority of their lifetimes burning hydrogen into helium, and display a tight relationship between temperature and luminosity known as the stellar main sequence. The positions and lifetimes of stars on the main sequence are mass-dependent, with more massive stars higher and further to the left. Once their hydrogen fuel is exhausted, stars begin to evolve away from the main sequence, beginning with the most massive, which are shortest lived. (Credit: ESO)*

dependence on the metal content, or metallicity, of the star, Z_*. This is normally measured relative to the Solar metal content, $Z_\odot$, though the precise value of the Solar metallicity is still a matter of debate (e.g. Asplund et al. 2009).

The position of an individual star on the stellar main sequence does not evolve substantially during its hydrogen-burning lifetime. However, once a star exhausts all of the available hydrogen fuel, it will begin to evolve away from the main sequence, as more complex nuclear reactions that generate heavier elements begin to take place. This evolution is initially towards the top-right of the HR diagram, as can be seen in Fig. 1.7.

Because the nuclear reaction rates in stars are determined by their masses, the time at which stars begin to evolve away from the stellar main sequence, as well as the specifics of their post-main-sequence evolution, are also determined by their masses. As a result of this, stellar evolutionary tracks are often represented by isochrones on the HR diagram. Individual isochrones show the positions occupied by stars of a fixed age as a function of their initial masses.

An example HR diagram with isochrones plotted for a range of ages is shown in the top-middle panel of Fig. 1.8. The stellar main sequence can again be seen running from top-left to bottom-right. For the youngest age, 10^6 yr, even the most massive stars remain on the main sequence. However, for older stellar populations, all stars above a certain mass threshold have evolved away from the main sequence, towards the top-right of the HR diagram. This threshold evolves towards lower masses with increasing age. Whilst post-main-sequence stars are typically more luminous than main sequence stars of the same mass, they are short lived, eventually exploding as supernovae, or shedding their material as planetary nebulae, and ceasing to radiate significant quantities of light. It is during the post-main-sequence evolution of stars that almost all of the metals in the Universe are produced (e.g. Hoyle 1946, 1954; Burbidge et al. 1957).

Simple stellar population models

Armed with a working knowledge of stellar evolution, it is possible to begin to construct models for the light from populations of stars in galaxies. The first step is to build models for the spectra of populations of stars that were all formed at the same time with the same metallicity, known as simple stellar population (SSP) models.

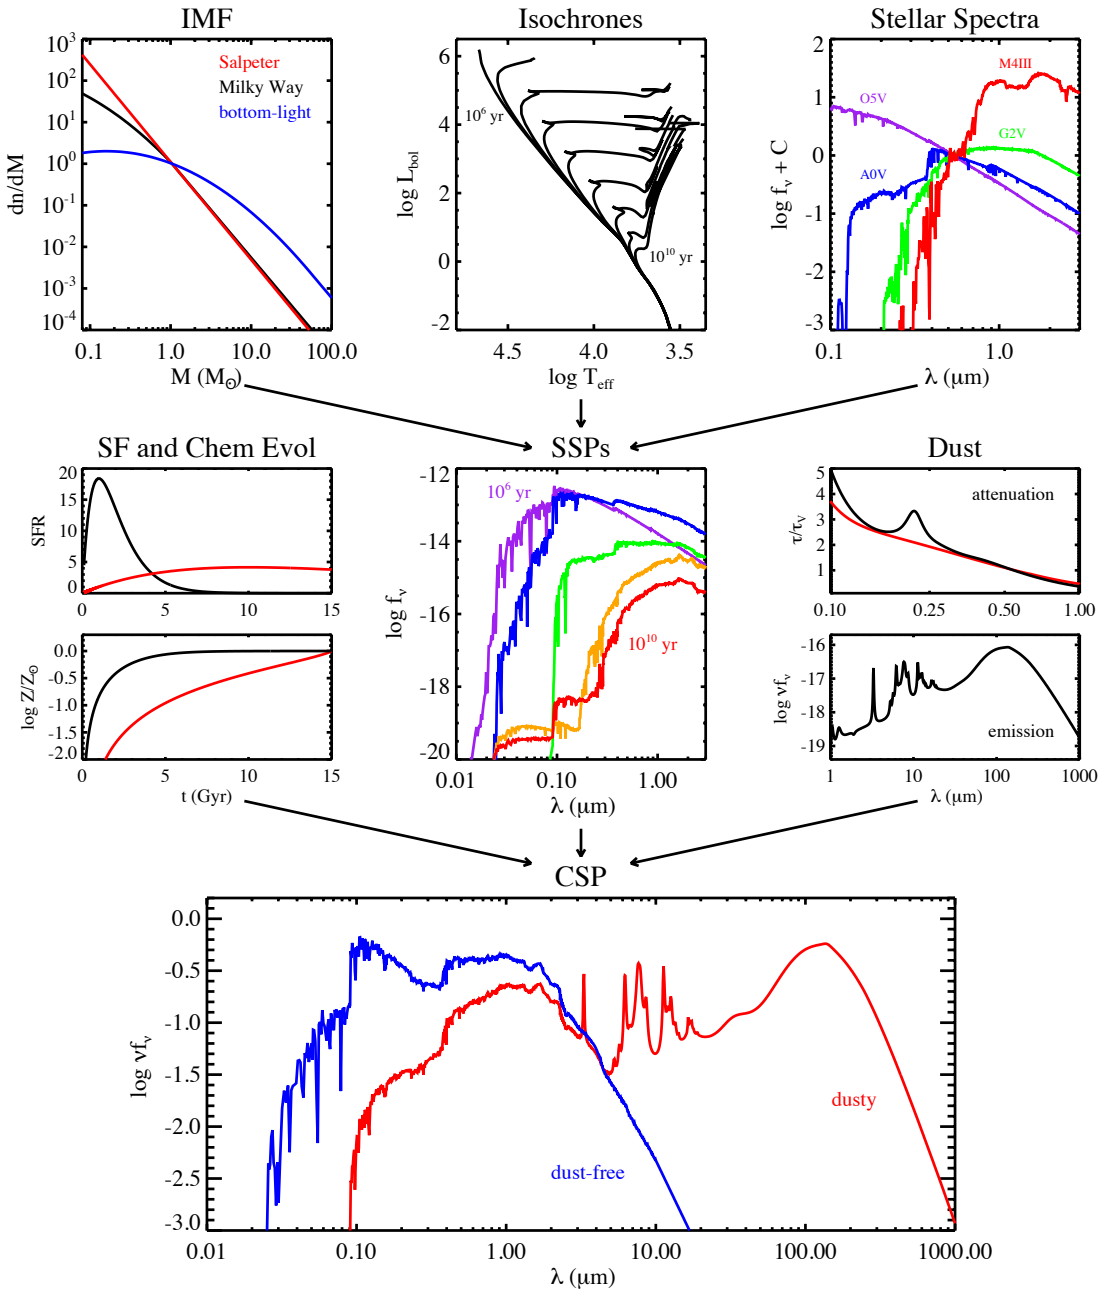

Figure 1.8 *A schematic diagram of stellar population synthesis modelling, taken from Conroy (2013). The three main inputs are the stellar initial mass function (IMF), isochrones describing the distribution of stars in luminosity and temperature as a function of age and mass, and individual stellar spectra. These are combined to produce simple stellar population (SSP) models for stellar populations with single ages and metallicities. SSPs can then be combined with a star-formation history (SFH) and chemical enrichment history (CEH) to produce a composite stellar population (CSP) model. The effects of dust attenuation and emission on a CSP model are also shown.*

SSP models are constructed by firstly observing the spectra of large numbers of stars across the HR diagram. It is also desirable to observe examples of stars with a range of metallicities at a given position on the HR diagram, such that SSP models with different metallicities can be constructed. However, this is often challenging, particularly in the high-mass, low-metallicity regime, as few local examples of these kinds of stars exist.

Once these spectra have been collected, the isochrone for each age at which a SSP model is desired can be overlaid on the HR diagram, and appropriate spectra that sample the isochrone selected. These spectra can then be added together to form the SSP model. The weights associated with different spectra must reflect the relative abundances of stars with different masses. These relative abundances are described by the stellar initial mass function (IMF; e.g. Salpeter 1955; Kroupa 2001; Chabrier 2003).

The three ingredients of SSP models (stellar spectra, isochrones and the IMF) are shown in the top row of Fig. 1.8, with example SSPs shown in the central panel. Because the luminosities of stars are strongly determined by their initial masses (in general, $M \sim L^{3.5}$), SSPs tend to be dominated by the most massive stars that remain on the main sequence and in post-main-sequence phases at that age. This means that SSPs display significant evolution with age, in the same sense as individual stellar spectra display strong evolution with mass. This effect becomes increasingly pronounced as one moves to shorter wavelengths, as the emission from hotter stars more strongly dominates the observed spectrum at young ages.

A secondary effect is observed as a function of stellar metallicity, with more-metal-rich populations displaying stronger individual metal absorption features, as well as a redder overall shape due to the blanketing effect of large numbers of blended metal absorption features.

Star-formation and chemical enrichment histories

Galaxies are assembled gradually from many generations of stars formed at different times. Therefore, in order to construct detailed models for galaxy stellar populations, it is necessary to combine SSPs with a range of different ages and metallicities to produce composite stellar populations (CSPs). A CSP is constructed as a weighted sum of different SSPs, with weights derived from models describing the star-formation and chemical-enrichment histories of the

galaxy. The star-formation history (SFH) is a function that describes how the star-formation rate (SFR) varied across cosmic time up to the time of observation. The chemical-enrichment history (CEH) describes the metal contents of stars formed at different times.

Several methods are in common usage for modelling galaxy SFHs, including simple parametric forms, multiple age bins with constant SFRs in each bin, and linear combinations of different SSP models. Conversely, galaxy CEHs are almost universally modelled with simple delta functions, assuming all stars in the galaxy have the same metallicity. This is because it is rarely possible to differentiate between CEH models using current data, and because of the computational expense of fitting complex models (see Section 1.4.3). The bottom panel of Fig. 1.8 shows an example CSP model in blue, built from the SSP models shown in the central panel, and SFH and CEH models shown in the middle-left panels.

1.3.2 Dust extinction, attenuation and emission

When stars end their lives, a fraction of their material is returned to the interstellar medium (ISM), including the metals they formed during their lifetimes. Certain elements, in particular carbon and silicon, then condense to form particles up to ~ 1 μm in diameter. These particles are usually referred to as dust, and their main effect is the absorption and scattering of light in the ultraviolet-optical wavelength range, and the re-radiation of energy as infrared photons. This significantly impacts the observed spectrum, and is the second-most-important effect that must be considered, after emission coming directly from the stellar population, when modelling galaxy spectra.

The effects of dust were first recognised through interstellar extinction of the spectra of stars in the Milky Way. Extinction occurs when light is absorbed and scattered out of the line of sight to an object, and results in the spectrum appearing fainter and redder, as light at shorter wavelengths is more strongly affected by dust. This wavelength-dependence is normally parameterised by A_λ, the total extinction in magnitudes at wavelength λ, which is related to the fraction of transmitted light, T_λ, by

$$A_\lambda = -2.5 \log_{10}(T_\lambda). \tag{1.18}$$

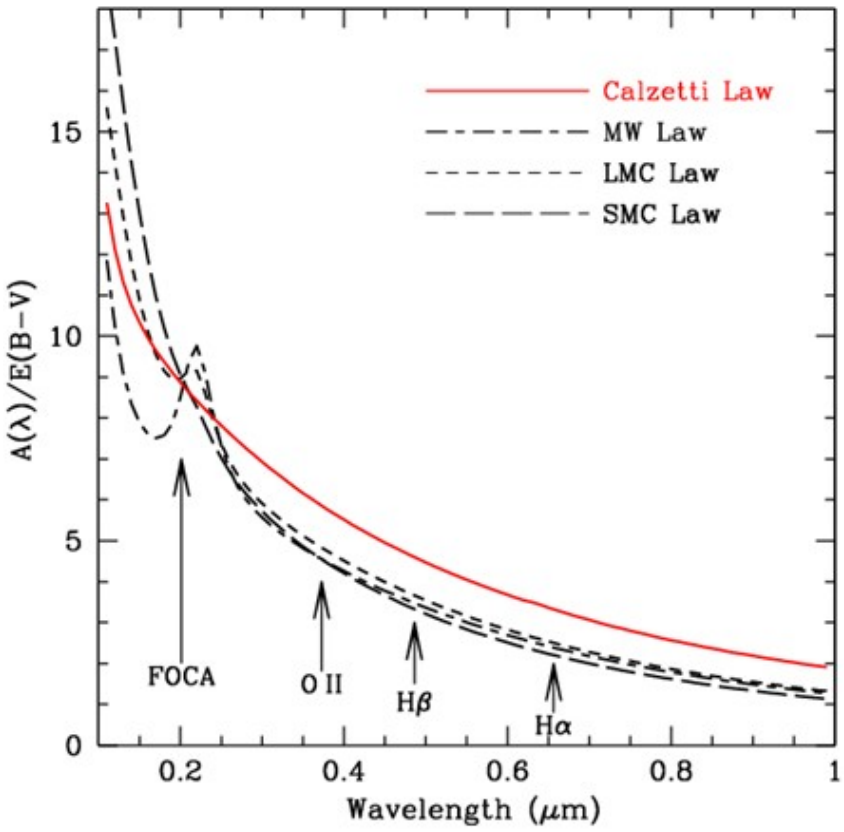

Figure 1.9 *A comparison of several commonly used dust attenuation and extinction laws, taken from Ellis (2008). The black lines show extinction curves, the red line shows the Calzetti et al. (2000) attenuation law for local star-forming galaxies. The extinction curves are similar in the optical, however they diverge in the ultraviolet, both in steepness and in the strength of the 2175Å bump.*

This is usually expressed relative to A_V, the total extinction in the V-band ($\lambda \sim$ 5500Å). The slope of the extinction law is often parameterised by R_V, the total-to-selective extinction ratio, given by

$$R_V = \frac{A_V}{E_{B-V}} = \frac{A_V}{A_B - A_V} \tag{1.19}$$

where E_{B-V} is the colour excess: the difference in extinction between the B-band and V-band. Smaller values of R_V correspond to steeper dust curves and more reddening of the light that passes through the dust.

Interstellar extinction was first studied in the local Universe: in the Milky Way (MW; e.g. Cardelli et al. 1989), Large Magellanic Cloud (LMC; e.g. Fitzpatrick 1986) and Small Magellanic Cloud (SMC; e.g. Prevot et al. 1984). A comparison

of the three resulting extinction laws is shown in Fig. 1.9. The three black lines can be seen to be similar in the optical, however the LMC and SMC laws are steeper in the ultraviolet when compared to the MW law. Another difference lies in the bump centred on 2175Å, which is strongest in the MW law, weaker in the LMC law, and absent in the SMC law. The precise mechanism that leads to this bump is still unclear, however it is thought to be related to polycyclic aromatic hydrocarbon (PAH) molecules.

When extending such analyses to unresolved stellar populations in distant galaxies, the geometry of the problem is subtly different. When observing individual stars in the MW, SMC or LMC, dust is positioned between the source and observer, whereas in unresolved galaxies the stars and dust are mixed together. This means that the effects of light being scattered back into the line of sight become important. The result is attenuation, rather than extinction, of the light emitted by stars. Early measurements of the average attenuation curve for star-forming galaxies in the local Universe were made by Calzetti et al. (1994, 2000). The resulting attenuation curve is often known as the starburst law, or Calzetti law, and is shown in red in Fig. 1.9. This can be seen to be significantly shallower than the three extinction curves shown, in part due to the scattering of light back into the line of sight.

When modelling the spectra of distant galaxies, it is common to assume one of the attenuation or extinction curves show in Fig. 1.9, allowing the normalisation (either A_V or E_{B-V}) to vary. However, evidence exists for significant object-by-object variations in galaxy attenuation curves (e.g. Kriek & Conroy 2013; Narayanan et al. 2018). This is largely due to galaxies having different orientations relative to our line of sight, but also potentially due to variations in dust composition. This is particularly true at high redshift, when dust and stars are less well mixed than at low redshift. One approach for dealing with this is to use a dust attenuation model with a variable slope, commonly in the form of a power-law with a variable spectral index (e.g. Charlot & Fall 2000).

The light absorbed by dust in the ultraviolet-optical is re-radiated in the infrared. This re-radiation can be reasonably well described (in the local Universe) by a black-body spectrum with a temperature of ~ 30 Kelvin at longer wavelengths ($\gtrsim 50$ μm), however at shorter wavelengths there is a significant contribution from hotter dust in star-forming regions, as well as broad emission lines due to PAH compounds at ~ 10 μm.

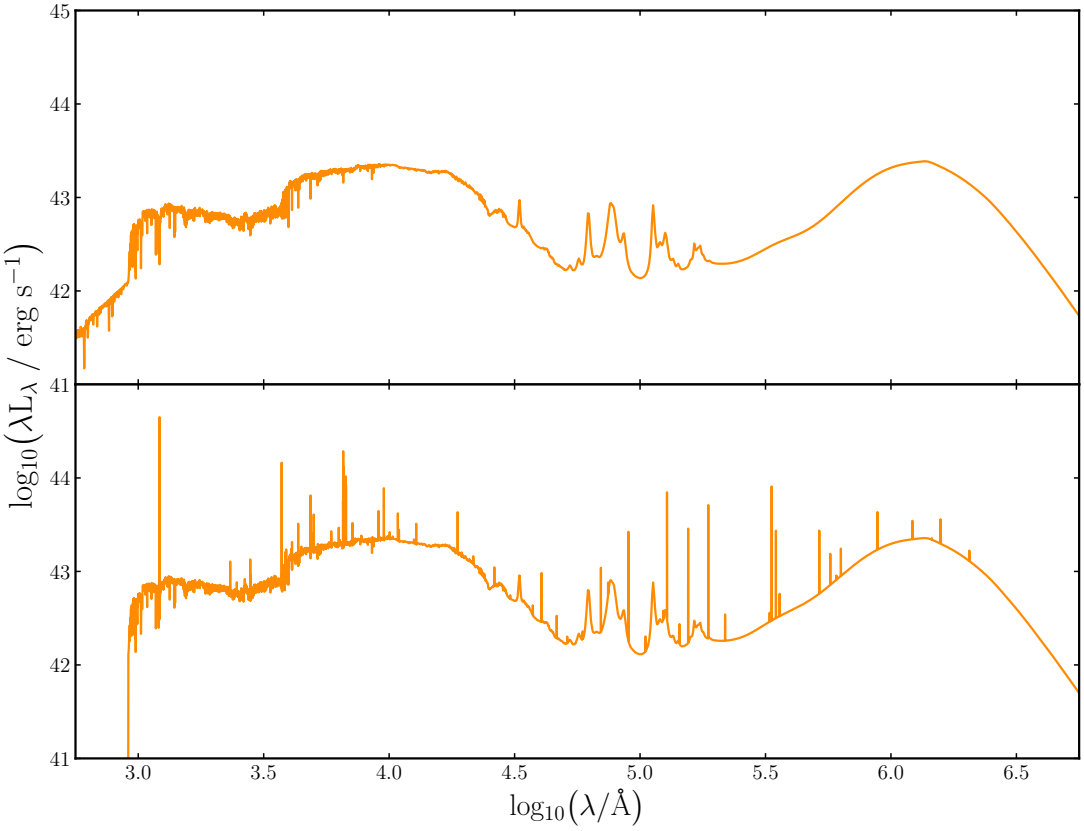

Figure 1.10 *An example galaxy spectral model before (top) and after (bottom) being processed through the* Cloudy *photoionization code. Light below 912Å photoionizes the gas in the surrounding* H II *region, which then re-radiates at specific wavelengths corresponding to atomic transitions of elements in the gas, such as hydrogen, oxygen and nitrogen. The ratios of these emission lines can be used to probe the physical conditions of the gas in star-forming regions.*

A widely used set of models to describe dust emission from galaxies were presented by Draine & Li (2007). These models parameterise the dust emission spectrum by $Q_{\rm pah}$, the percentage of the dust mass in PAH molecules; $U_{\rm min}$, the minimum starlight intensity to which the dust is exposed; and $\gamma_{\rm e}$, the fraction of the incident starlight that is at $U_{\rm min}$.

An example of the effects of dust attenuation and emission is shown in Fig. 1.8. The centre-right panel shows example dust attenuation curves and an example dust emission spectrum. The bottom panel shows the change to a purely stellar CSP model when dust attenuation and emission are considered. Light can be seen to be attenuated at $\lambda \lesssim 5$ μm, and redistributed to longer wavelengths.

1.3.3 Nebular emission

When stars are born, they remain enshrouded in the clouds of molecular gas and dust from which they formed for several million years. Radiation from the young stars photoionizes the surrounding gas, giving rise to H II regions, often referred to as nebulae or stellar birth clouds.

Recombination of electrons with hydrogen nuclei in these regions gives rise to strong line emission, most famously Balmer alpha, or Hα, at 6563Å, which gives these regions their characteristic red glow. Additionally, collisionally excited atomic transitions in metal ions give rise to further emission lines, with the strongest collisionally excited optical emission lines associated with forbidden transitions in oxygen and nitrogen ions. Finally, free-bound transitions and free-free (Bremsstrahlung) processes give rise to nebular continuum emission.

Both emission lines and nebular continuum emission can contribute strongly to the broad-band photometric fluxes observed from galaxies, meaning that it is important to include their effects in spectral models (e.g. Anders & Fritze-v. Alvensleben 2003; Byler et al. 2017). Additionally, the relative strengths of these emission lines can be used to probe the physical properties of the gas in these regions, in particular metallicity, temperature and density (e.g. Osterbrock 1989). However, separation of these subtle parameters is challenging, and requires extensive calibration.

Several pieces of software exist for physical photoionization modelling, which allow predictions to be made for the expected emission line ratios and strengths based upon a set of physical conditions and input stellar spectrum (e.g. Cloudy, Ferland et al. 2017; Mappings, Allen et al. 2008). These codes can be used to post-process galaxy spectral models in order to include the effects of H II regions on galaxy spectra. In Fig. 1.10, an example galaxy spectral model, including stellar emission and the effects of dust, is shown before (top panel) and after (bottom panel) being processed through the Cloudy photoionization code.

1.3.4 Intergalactic medium attenuation

Even after reionization (see Section 1.1.5), residual neutral hydrogen in the IGM still attenuates light short-wards of the wavelength of Lyα from distant galaxies. This absorption is observed to be split into discrete features at different

wavelengths, which are labelled as part of the Lyα forest (LAF), Lyman limit systems (LLS), or damped Lyα (DLA) systems depending on the strength of the absorption. These absorption features are thought to be associated with circumgalactic gas, as well as gas flowing along filaments, such as those shown in Fig. 1.2, towards galaxies (e.g. Carilli & Walter 2013).

The average transmission through the IGM can be modelled as a function of wavelength and observed redshift in order to correct galaxy spectral models for IGM attenuation. An analytic model was presented by Madau (1995), with separate contributions due to LAF, LLS and DLA systems. This model was updated to take into account improved statistical constraints on the distributions of these systems by Inoue et al. (2014).

In general terms, IGM attenuation is observed to increase with increasing redshift up to $z = 6$, at which point almost all of the light is attenuated. The average attenuation becomes higher as one moves to shorter wavelengths, as attenuation takes place due to an increasing number of atomic transitions in the Lyman series of hydrogen. This reaches a maximum at 912Å, corresponding to the Lyman limit: the energy required to ionize hydrogen from its ground state.

1.3.5 The age-metallicity-dust degeneracy

In the previous subsections, the main physical effects that determine the observed spectra of galaxies have been introduced. It is important to note that changes to several of these physical parameters result in very similar changes to the observed spectra of galaxies. Specifically, increasing the age, metallicity and dust contents of galaxies all act to redden their observed spectra, leading to the age-metallicity-dust degeneracy (e.g. Papovich et al. 2001; Lee et al. 2009).

This problem is most pronounced when only photometric data are available, as the broad-band colours of galaxies are particularly susceptible to the age-metallicity-dust degeneracy. Spectroscopic data are less susceptible, as the ages and metallicities of galaxies affect individual absorption features in galaxy spectra, as well as their overall shapes, allowing these effects to be separated out (e.g. Gallazzi et al. 2005). Historically, the age-metallicity-dust degeneracy has limited our ability to constrain these physical parameters, with photometric analyses limited to the determination of galaxy stellar masses and star-formation rates.

1.3.6 Monochromatic star-formation-rate indicators

As well as full models for the spectra of galaxies as a function of their physical parameters, a range of simpler models have historically been used to infer galaxy SFRs. These monochromatic indicators rely on observing the amount of light emitted at a specific wavelength, and empirical calibrations for converting these luminosities directly into SFRs (e.g. Kennicutt 1998; Kennicutt & Evans 2012).

These monochromatic indicators typically use either ultraviolet light that comes directly from the photospheres of young, high-mass O and B-type stars, or emission from such stars that is reprocessed into the infra-red through dust in the interstellar medium or into emission lines (typically Hα) by H II regions. The lifetimes of stellar birth clouds are relatively short ($\lesssim$ 10 Myr) compared to the lifetimes of the majority of O and B-type stars ($\sim$ 100 Myr). This means that Hα traces star-formation on shorter timescales than these other two indicators. SFRs based on ultraviolet and infrared light are complimentary, as they trace star-formation that is unobscured and obscured by dust respectively. The SFRs obtained through these two indicators are often added together to give an approximation of the total SFR.

1.4 Bayesian statistics and computational methods

The aim of scientific inquiry is to develop and refine of a set of beliefs about how the Universe works based upon observations. Bayesian statistics provides a rigorous theoretical framework within which this process can be conducted. One of the main innovations I seek to make in this thesis is the application of rigorous Bayesian statistical methods to the study of galaxy evolution. Bayesian statistics differs from frequentist statistics in the way probabilities are interpreted. In frequentist statistics, probabilities only have meaning in the context of repeated events: a probability is the fraction of trials that will result in a certain outcome. Bayesian statistics instead interprets probabilities as degrees of belief in hypotheses.

The law of conditional probability states that the probabilities of two events, A and B, are related by

$$P(A \mid B) = \frac{P(B \mid A)\, P(A)}{P(B)} \tag{1.20}$$

where $P(A)$ is the probability of event A, $P(B)$ is the probability of event B, $P(A \mid B)$ is the probability of A given that B has happened, and $P(B \mid A)$ is the probability of B given that A has happened.

Bayes' theorem (Bayes & Price 1763) is the same formula as Equation 1.20, with the important difference being that probabilities are now assigned to hypotheses, rather than events. First, let us assume we have some model, $\mathcal{M}$, which we believe describes some physical process. The predictions made by this model are a function of some parameters, $\{\theta_i\}$, for example the density of matter or dark energy in the Universe. Two questions naturally arise: firstly, what are the most probable values for the parameters $\{\theta_i\}$, and, secondly, how good is our model at describing the physical process of interest? These correspond to the two main objectives of Bayesian analyses: parameter estimation and model selection.

Equation 1.20 can be used to address these questions as follows. Firstly, let event A be that a given set of parameter values are true: $P(A) = P(\{\theta_i\})$. Now suppose we conduct some experiment that returns a number of measurements, $\{y_j\}$, relevant to the model $\mathcal{M}$. Let event B be the measurement of a given set of experimental results, so that $P(B) = P(\{y_j\})$. We will initially assume that the model $\mathcal{M}$ is correct, meaning that all of our probabilities are conditional on $\mathcal{M}$ (this will become important in Section 1.4.2). We can now write

$$P(\{\theta_i\} \mid \{y_j\},\ \mathcal{M}) \ = \ \frac{P(\{y_j\} \mid \{\theta_i\},\ \mathcal{M})\ P(\{\theta_i\} \mid \mathcal{M})}{P(\{y_j\} \mid \mathcal{M})}. \tag{1.21}$$

It is however common to suppress the conditionality on $\mathcal{M}$ and write this as

$$P(\{\theta_i\} \mid \{y_j\}) \ = \ \frac{P(\{y_j\} \mid \{\theta_i\})\ P(\{\theta_i\})}{P(\{y_j\})}. \tag{1.22}$$

We now have an expression for $P(\{\theta_i\} \mid \{y_j\})$, the probability of a given set of model parameters being true given a set of experimental measurements. This is commonly referred to as the posterior probability distribution. This depends on several quantities. Firstly $P(\{y_j\} \mid \{\theta_i\})$, the probability of obtaining the data given that a certain set of model parameters are true. This is commonly referred to as the likelihood, and often written as $\mathcal{L}(\{y_j\} \mid \{\theta_i\})$. Secondly, $P(\{\theta_i\})$, the prior probability of a given set of model parameters before the experiment has taken place. This could be based on previous experiments or other relevant knowledge of the physical process being examined. Finally, $P(\{y_j\})$,

the probability of obtaining the data, which is often referred to as the Bayesian evidence, or marginal likelihood. In parameter estimation the evidence can be treated simply as a normalisation factor, however it has a central role in Bayesian model selection, as will be discussed in Section 1.4.2.

1.4.1 Parameter estimation and marginalisation

The aim of parameter estimation is to use some observational data to obtain constraints on the physical parameters of interest. For example, one might wish to use the CMB data shown in Fig. 1.1 to estimate the Hubble constant, H_0, discussed in Section 1.1.1. This information is contained within the posterior probability distribution, obtained by solving Bayes' equation. For a discussion of how this is achieved in practice, see Section 1.4.3.

Typically, models have several free parameters, meaning that the posterior probability distribution is multivariate. Often, the set of parameters will include several that are of physical interest, along with nuisance parameters that describe uninteresting effects, for example the noise properties of the data being examined. It is often of interest to extract the probability distribution for individual parameters, such that confidence intervals can be placed on their values. Alternatively, one might wish to extract a small subset of parameters, such that the correlations between these parameters can be examined.

In order to obtain the probability distribution for a parameter, or parameters of interest, it is necessary to marginalise over all possible values of the other parameters of the distribution. This is achieved by integrating the joint posterior probability distribution with respect to the uninteresting parameter(s). Consider a simple case, in which the set of model parameters $\{\theta_i\}$ consists of two elements, θ_1 and θ_2. We can obtain the marginal posterior probability distribution for θ_1, $P(\theta_1 \mid \{y_j\})$ by integrating θ_2 out of the joint distribution,

$$P(\theta_1 \mid \{y_j\}) = \int P(\{\theta_i\} \mid \{y_j\}) \, \mathrm{d}\theta_2. \tag{1.23}$$

For the case in which one has a series of samples from the posterior distribution instead of an analytic expression (see Section 1.4.3), this process becomes trivial as the values of uninteresting parameters can simply be ignored.

1.4.2 Bayesian model selection

Whilst the values of certain physical parameters are often of considerable interest, in most cases the most interesting questions relate to whether or not a model is appropriate for describing the available data. Continuing with the example of Section 1.4.1, this is the difference between using the CMB data shown in Fig. 1.1 to constrain H_0, and using it to assess whether Λ-CDM is an appropriate cosmological model. The latter class of questions can be addressed using Bayesian model selection.

A frequentist approach: the minimum reduced chi-squared value

To illustrate the advantages (and disadvantages) of Bayesian model selection, it is useful to first consider a more traditional, frequentist method of assessing the ability of a model to describe data (alternatively the quality, or goodness of fit). We will consider Pearson's chi-squared test (Pearson 1900).

We first assume that the uncertainties on our measurements, which we will call $\{\sigma_j\}$, are known, and are independently Gaussian distributed. Under these assumptions, the log-likelihood of a particular set of model parameters being true is given by

$$\ln \mathcal{L}(\{y_j\} \mid \{\theta_i\}) = -\frac{1}{2} \sum_j \left(\frac{m_j(\{\theta_i\}) - y_j}{\sigma_j} \right)^2 \tag{1.24}$$

where $m_j(\{\theta_i\})$ is the prediction made by the model for the datum y_j, assuming parameter values $\{\theta_i\}$. We now vary the model parameters in order to find the values, $\{\theta_i\}_{\max}$, which correspond to the maximum likelihood. A variety of computational methods exist for finding these maximum likelihood values (e.g. Byrd et al. 1995).

Assuming that the model, with its parameters set to their maximum likelihood values, correctly describes our data, the value of the sum on the right hand side of Equation 1.24 should follow a chi-squared distribution across a large number of repeat experiments. We therefore identify the maximum likelihood parameters

with the minimum chi-squared value, χ^2_{min}, given by

$$\chi^2_{\text{min}} = \sum_j \left(\frac{m_j(\{\theta_i\}_{\text{max}}) - y_j}{\sigma_j} \right)^2. \tag{1.25}$$

This chi-squared distribution has ν degrees of freedom, where ν is the number of data-points minus the number of free parameters of the model. It is common to divide through χ^2_{min} by ν to give the minimum reduced chi-squared value

$$\chi^2_{\nu,\ \text{min}} = \frac{\chi^2_{\text{min}}}{\nu}. \tag{1.26}$$

As the right-hand side of Equation 1.25 is chi-squared distributed, we can calculate the probability of obtaining a higher value of $\chi^2_{\nu,\ \text{min}}$ by comparing with a chi-squared distribution with ν degrees of freedom. This gives a rough idea of the probability that our model correctly describes our data.

For example, consider the case in which 10 experimental data-points are being fitted with a model containing 3 free parameters, and a minimum reduced chi-squared value of $\chi^2_{\nu,\ \text{min}} = 1.1$ is obtained. There is a ~ 37 per cent chance of obtaining $\chi^2_{\nu,\ \text{min}} > 1.1$ from a chi-squared distribution with 7 degrees of freedom, so the model can be said to provide a reasonable description of the data. This percentage, expressed as a decimal, is often called the p-value. Conversely, if $\chi^2_{\nu,\ \text{min}} = 2.7$, the probability of a larger $\chi^2_{\nu,\ \text{min}}$ is only ~ 1 per cent, corresponding to a lower probability that the model correctly describes the data. Typically some p-value threshold is chosen beyond which the model (or null hypothesis) is rejected, for example 0.05 or 0.01.

A Bayesian approach: the evidence ratio

When constructing a model it is desirable not only that it produces a good fit to the data, but also that it is predictive. A more complex model, typically one with a larger number of free parameters, is likely to be able to match a larger range of observational data than a simpler model. This means that the simpler model makes stronger predictions about what should be observed: it is more predictive. This is a statement of Ockham's razor, a heuristic argument that the simplest explanation consistent with the observations should be favoured over others. It is for this reason that the minimum chi-squared value is typically divided by the

number of degrees of freedom, disfavouring models with larger numbers of free parameters when $\chi^2_{\nu,\,\mathrm{min}}$ values for different models are compared.

The quality of fit and predictivity of a model are both encapsulated within the denominator of the right hand side of Equation 1.21, the Bayesian evidence, $P(\{y_j\} \mid \mathcal{M})$. This is often denoted by Z. As the posterior probability distribution must be normalised to 1, the Bayesian evidence is the integral of the likelihood multiplied by the prior across the whole parameter space of $\{\theta_i\}$. Alternatively it is the average value of the likelihood across the parameter space (or prior volume). Models that include large prior volumes that do not correspond to the data (those that are not predictive) are naturally disfavoured by this, as well as models that fail to reproduce the data. For a more detailed discussion see Trotta (2008).

Unfortunately, the Bayesian evidence value for one model does not have a clear interpretation in terms of the quality of fit, as is the case for the minimum reduced chi-squared value. However, the evidence ratio between two models is a more principled way of constructing a likelihood ratio, indicating which of two models is preferred, than comparing minimum reduced chi-squared values. Consider two models, $\mathcal{M}_1$ and $\mathcal{M}_2$, seeking to explain the same data, $\mathcal{D}$. We can use the Bayesian evidence values, $P(\mathcal{D} \mid \mathcal{M}_1)$ and $P(\mathcal{D} \mid \mathcal{M}_2)$, along with Equation 1.20 to obtain the ratio of posterior probabilities for the two models given the data,

$$\frac{P(\mathcal{M}_1 \mid \mathcal{D})}{P(\mathcal{M}_2 \mid \mathcal{D})} = \frac{P(\mathcal{M}_1)}{P(\mathcal{M}_2)} \frac{P(\mathcal{D} \mid \mathcal{M}_1)}{P(\mathcal{D} \mid \mathcal{M}_2)}, \tag{1.27}$$

where $P(\mathcal{M}_1 \mid \mathcal{D})$ and $P(\mathcal{M}_2 \mid \mathcal{D})$ are the posterior probabilities of each model given the data, and $P(\mathcal{M}_1)$ and $P(\mathcal{M}_2)$ are their prior probabilities. The ratio of these prior probabilities is usually assumed to be one, unless previous experiments have demonstrated a preference for one model over the other.

Bayesian evidence vs minimum reduced chi-squared

Comparing the two approaches we have considered, the minimum reduced chi-squared value has the advantage of being simple to calculate and providing a clear method of evaluating the quality of fit for an individual model. However, it relies on a very narrow set of assumptions, in particular the assumption of well-known, independent Gaussian uncertainties. The p-value, which is the end product of such analyses, is a measure of the probability of obtaining a worse fit to the data

assuming the model is correct. Whilst this is a helpful indicative value in a wide range of circumstances, it is not a measurement of the probability that the model being examined is the correct one for describing the data. The choice of p-value threshold for rejection of the model is somewhat arbitrary, and does not take into account the role of the prior.

Conversely, the Bayesian evidence can be used to rigorously obtain the ratio of probabilities for two models given the observational data. This is the specific quantity of interest in the majority of circumstances. Bayesian evidence is also applicable to a larger range of problems that do not conform to the assumptions necessary for the calculation of p-values based on the minimum reduced chi-squared value. However, the computational challenges inherent in solving the multi-dimensional integral by which Z is computed make evidence ratios challenging to obtain in most circumstances. It is also not possible to use Bayesian evidence to evaluate whether a single model produces an acceptable quality of fit to data: it is only useful in the context of comparing how well two or more models perform.

1.4.3 Solving Bayes' equation: computational methods

To apply Bayesian methods for parameter estimation and model selection as discussed in Sections 1.4.1 and 1.4.2, the posterior distribution (and potentially the Bayesian evidence) must be obtained through solving Equation 1.21. This is usually a complex multi-dimensional problem that must be solved numerically. Typical methods rely on drawing a representative set of samples from the posterior distribution across the prior volume. This means that the local density of samples is proportional to the value of the posterior distribution at that point in parameter space. The development of efficient algorithms for this kind of sampling is a highly active area of ongoing research.

Markov chain Monte Carlo methods

Markov chain Monte Carlo (MCMC) methods are a class of algorithms used for sampling from probability distributions. In general terms, a chain of samples is constructed by jumping randomly around parameter space, accepting or rejecting points depending on the value of some function at that point. The Metropolis-Hastings algorithm (Metropolis et al. 1953) provides a general framework for

sampling in this way, and represents the foundation of much work in the field.

The process of sampling from a posterior probability distribution $P(\{\theta_i\} \mid \{y_j\})$ using the Metropolis-Hastings algorithm can be summarised as follows:

1. Choose some starting parameter values $\{\theta_i\}_0$ and calculate $P(\{\theta_i\}_0 \mid \{y_j\})$.

2. Using some pre-defined proposal function (see below), propose a new point $\{\theta_i\}_1$ and calculate $P(\{\theta_i\}_1 \mid \{y_j\})$.

3. Draw a random number, a from a uniform distribution between 0 and 1.

4. If $\dfrac{P(\{\theta_i\}_1 \mid \{y_j\})}{P(\{\theta_i\}_0 \mid \{y_j\})} > a$, accept the point $\{\theta_i\}_1$, otherwise stay at $\{\theta_i\}_0$.

5. Repeat steps $2-4$ until the desired number of accepted samples, N, from the parameter space have been generated.

The chain of points in parameter space $(0, \ldots, N)$ will be draws from the probability distribution $P(\{\theta_i\} \mid \{y_j\})$. In practice, care must be taken to ensure these are not strongly correlated, for example by selecting only every tenth point. The proposal function is some probability distribution used to generate new points. In the simplest case this is a Gaussian distribution with a fixed variance in each parameter. Provided that the algorithm is allowed to converge, the choice of proposal function does not affect the final outcome of the calculation. However the choice of proposal function will strongly determine the efficiency of the sampling process, and hence the time taken to reach convergence.

Modern MCMC methods make use of highly optimised proposal functions and multiple chains run in parallel to improve the efficiency of the sampling process (e.g. Goodman & Weare 2010; Foreman-Mackey et al. 2013). However, MCMC methods still struggle to characterise posterior distributions with multiple, widely separated peaks or strong curving degeneracies, particularly in higher-dimensional spaces (e.g. Feroz & Hobson 2008). Additionally, calculation of the Bayesian evidence with MCMC (for example using thermodynamic integration techniques) is extremely computationally expensive, although the development of new and more efficient methods is ongoing (e.g. Heavens et al. 2017).

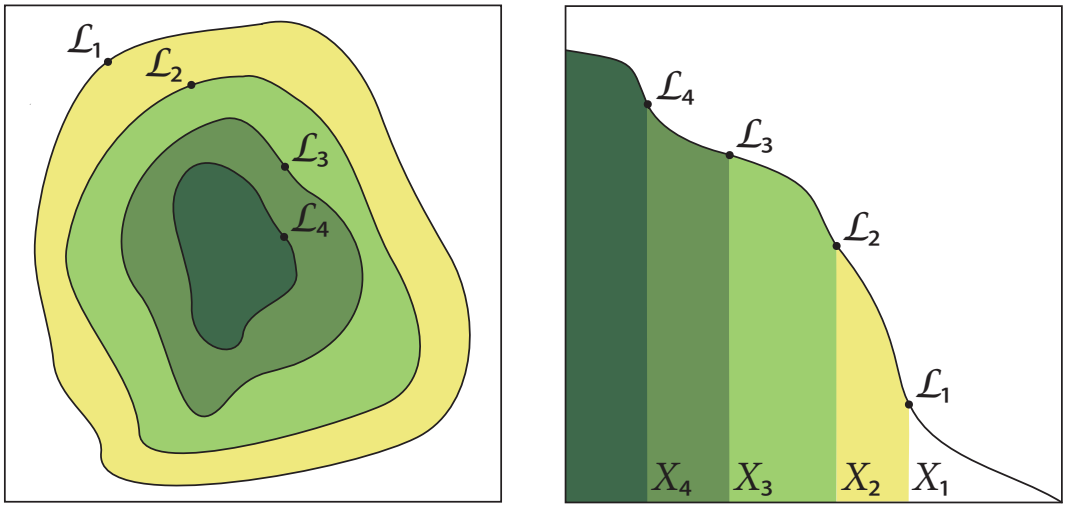

Figure 1.11 *Schematic diagram of the nested sampling process, taken from Feroz et al. (2013). To the left an example two-dimensional parameter space is shown with four example points, $\mathcal{L}_{1-4}$. Contours have been drawn enclosing all points with higher likelihoods. To the right the parameter space is shown mapped into one dimension: the fraction of the prior volume, X, within the likelihood contour. By sampling points from the likelihood-bounded prior and statistically determining the associated X values, the multi-dimensional evidence integral can be simply evaluated as a one-dimensional integral over X.*

Nested sampling methods

Introduced by Skilling (2006), nested sampling is a method that allows efficient calculation of both the posterior distribution and Bayesian evidence for problems with low to moderate dimensionalities ($\lesssim 15-20$). Importantly, nested sampling can be used to efficiently sample from posterior distributions that include multiple widely separated maxima and severe multi-parameter degeneracies.

The central idea of nested sampling is that the multi-dimensional integral of the likelihood function over the prior volume that returns the Bayesian evidence can be represented as a one-dimensional integral over $X(\lambda)$, the fraction of the prior volume with likelihood greater than λ,

$$Z = \int_{\Theta} \mathcal{L}(\{y_j\} \mid \{\theta_i\})\, P(\{\theta_i\})\, \mathrm{d}\Theta = \int_0^1 \mathcal{L}(X)\, \mathrm{d}X, \tag{1.28}$$

where Θ represents the multi-dimensional parameter space $\{\theta_i\}$. This transform-

ation is represented schematically in Fig. 1.11.

In order for this to be helpful however, a series of points, $\mathcal{L}_k$, must be drawn from the prior that can be associated with known fractional volumes, X_k. The evidence can then be obtained using simple, one-dimensional numerical integration techniques such as the trapezium rule.

Skilling (2006) demonstrated that, if the point $\mathcal{L}_k$ can be sampled at random from within the region bounded by the likelihood contour associated with $\mathcal{L}_{k-1}$, the value of X_k can be determined statistically, and is given by

$$X_k = e^{-k \pm \sqrt{k}}, \tag{1.29}$$

where $\mathcal{L}_0 = 0$ and $X_0 = 1$. This exponential relationship is desirable, as the bulk of the evidence typically occupies a very small fraction of the prior volume. The evidence can then be computed from this set of points using the trapezoidal rule.

The final missing component is an efficient method of sampling points at random from within the likelihood-bounded prior. In the simplest case, this can be achieved by rejection sampling, with points drawn at random from the whole prior volume. However this approach scales very poorly with dimensionality, and becomes increasingly inefficient as the process moves towards higher likelihoods, corresponding to smaller fractional prior volumes.

The process described above can be readily generalised to a scheme involving more than one "live point", in which the point with the lowest likelihood is replaced at each step. This is advantageous, firstly because it provides improved statistical constraints on the evidence, but also because the current set of live points, which are randomly distributed within the current likelihood contour, can be used to construct a boundary from within which new samples can be drawn.

This is usually achieved by calculating the covariance matrix for the positions of the current set of live points, then using this to construct an ellipsoidal surface that encloses all of the points. This bound is typically expanded by some factor to allow for the incomplete sampling of the likelihood contour by the current set of live points. In order to increase the efficiency of the process, an optional clustering step can be implemented, to allow for the drawing of multiple ellipsoids around separated clusters of live points. This has been shown to significantly increase the efficiency of sampling, in particular when dealing with strong degeneracies

and multi-modal posterior distributions (Feroz & Hobson 2008).

As well as being used to calculate the evidence, the set of points $\mathcal{L}_k$ represent a set of weighted samples from the posterior distribution, meaning that this is obtained for free as a by-product of the evidence calculation. All posterior distributions and Bayesian evidence values reported in this thesis were calculated using the MultiNest nested sampling algorithm (Feroz & Hobson 2008; Feroz et al. 2009, 2013), which uses the multi-ellipsoidal approach outlined above.

1.5 Towards a theory of galaxy formation

In this section I will summarise our current understanding of galaxy evolution across cosmic time. This understanding has been developed by observing galaxies using the methods described in Section 1.2, building physical models to describe these observations as described in Section 1.3, and fitting these models to data using methods such as the Bayesian statistical approach described in Section 1.4.

In parallel with the models for interpreting galaxy spectra discussed in Section 1.3, a wide variety of models have been developed that attempt to describe the physical properties of galaxies from first principles. These models begin with descriptions of the underlying physical mechanisms, and work up to properties that can be inferred from galaxy spectra. By attempting to match observational results in this way, an understanding can be gained of the important physical mechanisms that lead to the galaxy population we observe.

Because of the complexity of galaxy formation physics, connecting underlying mechanisms with observed properties is usually a significant computational exercise, involving semi-analytic modelling or numerical hydrodynamic simulations. A full discussion of such simulations is outside of the scope of this thesis, however in this section I will refer to certain key results that inform our interpretation of observations. For a detailed review, see Somerville & Davé (2015).

1.5.1 The galaxy luminosity and stellar-mass functions

As discussed in Section 1.1.4, the clustering properties of galaxies are a key prediction of the CDM model. Another key probe is the number density of galaxies as a function of luminosity, or galaxy luminosity function. However,

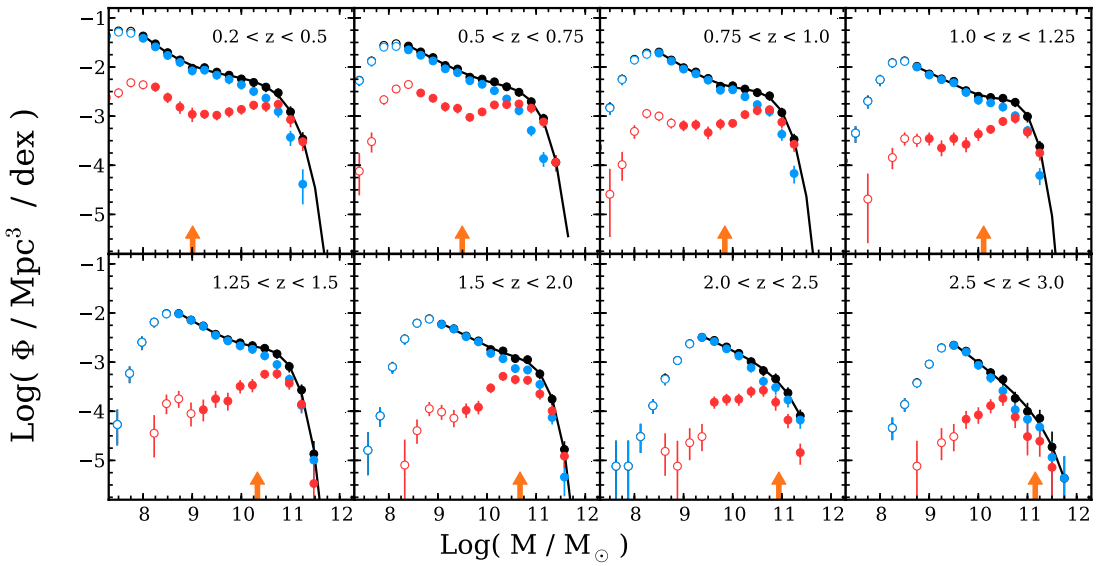

Figure 1.12 *The evolution of the galaxy stellar-mass function since $z = 3$, taken from Tomczak et al. (2014). Results for the whole galaxy population are shown in black, with star-forming galaxies in blue and quiescent galaxies in red. Open circles are below the mass-completeness limits for each redshift bin.*

since the majority of the optical light emitted by galaxies comes from stars, the galaxy luminosity function is also a key probe of the physics of galaxy formation and evolution (e.g. White & Rees 1978; Frenk et al. 1988).

Early work focussed on local galaxies (e.g. Schechter 1976, Shanks et al. 1984; Efstathiou et al. 1988), with the luminosity function found to be well described by a power-law at low luminosities, with an exponential cut-off at high luminosities. This parametric form is commonly referred to as the Schechter function. As deeper photometric and spectroscopic observations became feasible, these studies were extended to higher redshifts, allowing constraints to begin to be placed on the build-up of stellar mass in galaxies across cosmic time (e.g. Broadhurst et al. 1988; Colless et al. 1990; Lilly et al. 1995; Ellis et al. 1996). A key development came at the beginning of the 21$^{\rm st}$ Century, when robust measurements of stellar mass-to-light ratios began to allow for galaxy luminosity functions to be translated directly into galaxy stellar-mass functions (e.g. Bell & de Jong 2001; Bell et al. 2003; Dickinson et al. 2003b, Faber et al. 2007).

Contemporary studies allow strong statistical constraints to be placed on the evolution of the galaxy luminosity and stellar mass functions out to $z \sim 4$ (e.g. Muzzin et al. 2013; Tomczak et al. 2014; Mortlock et al. 2017), demonstrating that approximately half of the current stellar-mass density of the Universe had

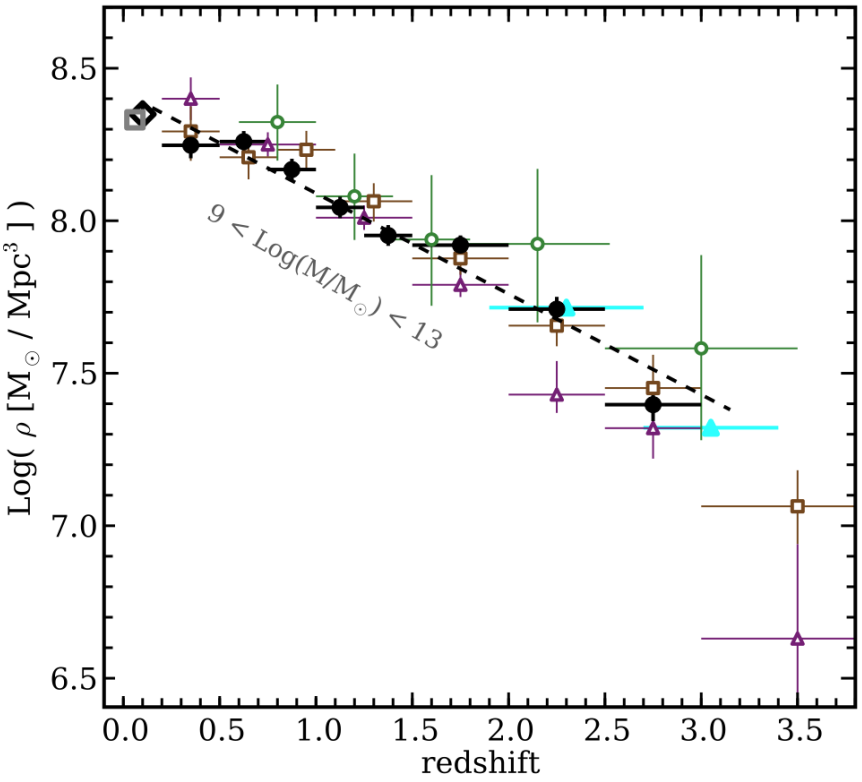

Figure 1.13 *A compilation of results from recent studies showing stellar mass assembly across cosmic time, taken from Tomczak et al. (2014). These results are obtained by integrating the observed galaxy stellar-mass function (Fig. 1.12) from* $9 < \log_{10}(M_*/\mathrm{M}_\odot) < 13$.

formed by $z \sim 0.75$, with only $\sim$ 10 per cent formed by $z = 2$, and only $\sim$ 1 per cent by $z = 3.5$. Fig. 1.12 shows an example of the evolution of the galaxy stellar-mass function from $0.2 < z < 3$ from Tomczak et al. (2014), whereas Fig. 1.13 shows the build-up of stellar mass obtained by integrating over the galaxy stellar-mass function at different observed redshifts. Efforts are ongoing to extend such studies out to the highest redshifts (e.g. McLure et al. 2013; Bowler et al. 2014, 2015, McLeod et al. 2015, 2016).

The galaxy luminosity and stellar-mass functions can be used to inform models of galaxy formation in a variety of ways. On a basic level, matching the observed functions is a simple and widely applied test of any galaxy formation simulation. Deeper insights can be obtained by considering the connections between galaxies and the dark matter halos in which they reside (e.g. Wechsler & Tinker 2018). One approach is to associate galaxies with dark matter halos from an N-body simulation (see Section 1.1.4), assuming the most massive galaxies correspond to the most massive halos, a process known as abundance matching. Fig. 1.14 shows

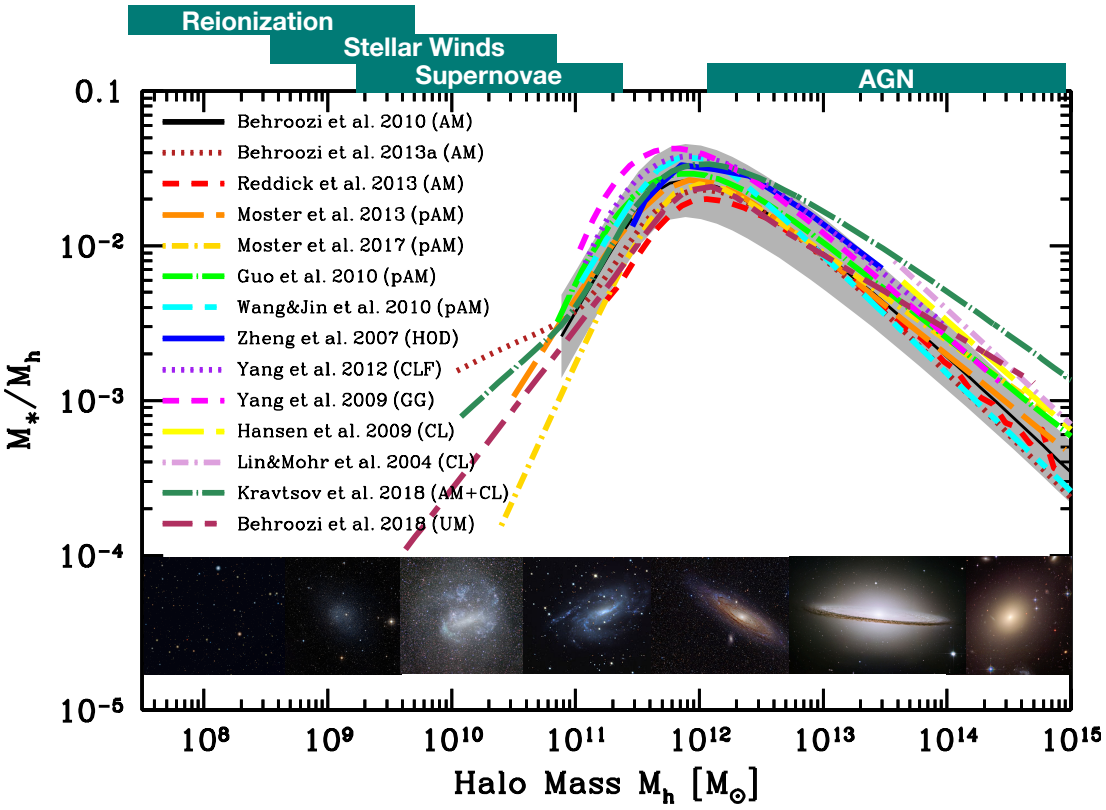

Figure 1.14 *Stellar-mass to halo-mass ratios for central galaxies at $z = 0$ as a function of halo mass, taken from Wechsler & Tinker (2018), adapted from Behroozi et al. (2018). Processes that play a significant role in inhibiting star-formation at different masses are indicated at the top. Example images are shown at the bottom.*

the results of a large number of such analyses, compiled by Behroozi et al. (2018) and presented by Wechsler & Tinker (2018). Stellar-mass to halo-mass ratio for halos of different masses are shown, which is a measure of how efficient halos of different masses are at converting gas into stars. Halos of mass $\sim 10^{12}$ $M_\odot$, corresponding to stellar masses of $\sim 10^{10.5}$ $M_\odot$, can be seen to be the most efficient.

The factors that affect the star-formation efficiencies of galaxies have been a subject of widespread research. The emerging picture is one in which the formation of stars, fuelled by gravitationally driven accretion of gas, is regulated by the feedback of energy from baryonic processes. At the lowest masses, energy input from early galaxies as part of reionization (see Section 1.1.5) is thought to prevent gas from ever beginning to condense to form stars. At higher masses, star-formation occurs, but this is self-regulating, with energy input from stellar winds and supernovae moderating further star-formation. This kind of energy input is however insufficient to eject gas from the deeper potential wells of the

most massive halos. Instead, AGN activity, powered by accretion of gas onto supermassive black holes, is thought to provide the energy necessary to inhibit star-formation (e.g. Croton et al. 2006; Davé et al. 2016).

1.5.2 Cosmic star-formation history

Another key probe of galaxy evolution across cosmic time is the cosmic star-formation history. This is the time differential of the stellar mass density, and is related to the time-evolution of the galaxy stellar-mass function, since it determines how much new stellar mass is added to galaxies as a function of time.

The most common method of constraining this evolution is by constraining the star-formation rates of individual galaxies across a wide range in redshift (see Section 1.3.6). However attempts can also be made based on inferring the star-formation histories (see Section 1.3.1) of representative samples of local galaxies, which can then be summed to give the cosmic star-formation history (e.g. Heavens et al. 2000, 2004; Panter et al. 2003, 2007, Tojeiro et al. 2007; Leitner 2012).

Early results by Lilly et al. (1996) found that the cosmic star-formation-rate density (SFRD) appeared to increase by approximately an order of magnitude moving back in time from the present day to $z = 1$. This analysis was extended to higher redshifts ($2.0 < z < 4.5$) by Madau et al. (1996, 1998) and Sawicki et al. (1997), who found a SFRD higher than the local Universe, but lower than at $z = 1$, suggesting that the SFRD rose and then fell across cosmic time, peaking somewhere around $1 < z < 2$.

Modern studies use a wide variety of observational data to constrain the SFRD, based upon both rest-frame ultraviolet emissions, which come directly from young stars, and infrared observations, which probe star-formation obscured by dust. Madau & Dickinson (2014) compiled the results of many such studies, resulting in a consensus best-fitting SFRD evolution function, $\psi(z)$, given by

$$\psi(z) \;=\; 0.015\,\frac{(1+z)^{2.7}}{1+[(1+z)/2.9]^{5.6}}\;\mathrm{M}_{\odot}\;\mathrm{yr}^{-1}\;\mathrm{Mpc}^{-3}. \tag{1.30}$$

This best-fitting function, along with the data analysed by Madau & Dickinson (2014), are shown in Fig. 1.15. It can be seen that cosmic star-formation rose rapidly within the first ~ 4 Gyr following the Big Bang to a peak at $z \sim 2$. At

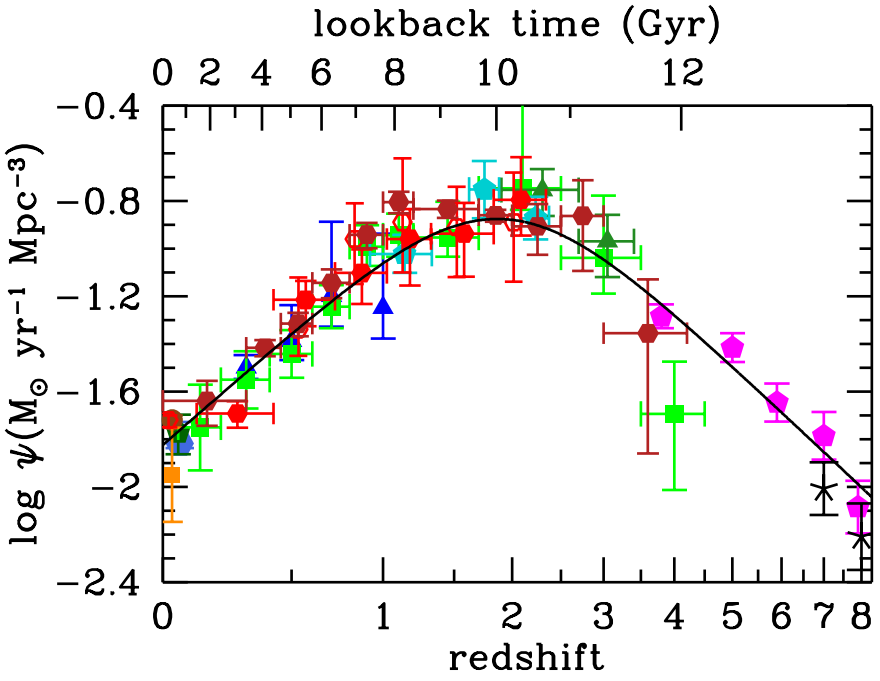

Figure 1.15 *The history of the cosmic star-formation-rate density, taken from Madau & Dickinson (2014). The black line shows the best-fitting function describing the data, given in Equation 1.30. The SFRD rose rapidly over the first ∼ 4 Gyr of cosmic history to a peak at $z \sim 2$, then declined by approximately a factor of ten to the present.*

this point, star-formation activity levelled off, then began to fall, with current levels approximately a factor of ten lower.

1.5.3 The star-forming sequence

Whilst both the galaxy stellar-mass function and cosmic star-formation rate density describe the integrated properties of the galaxy population as a function of time, they offer relatively little information about the evolutionary paths followed by individual galaxies. For example, the global evolution shown in Fig. 1.15 could be dominated by the steady rise and fall of star-formation activity in all galaxies, or it could be dominated by a small number of galaxies going through periods of extremely elevated star-formation activity, with more galaxies in this state at $z = 1$ than at other times.

In order to obtain a more nuanced understanding of galaxy evolution, it is necessary to consider the relationship between the stellar masses and star-formation rates of individual galaxies as a function of cosmic time. An early

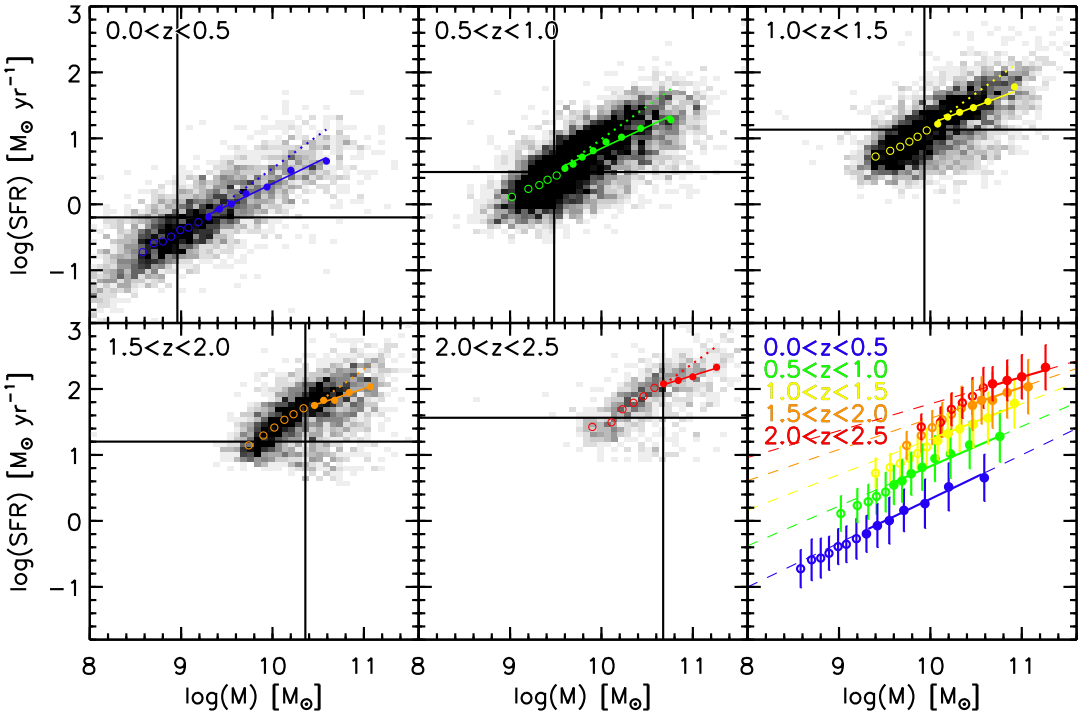

Figure 1.16 *The star-forming sequence from $0 < z < 2.5$, taken from Whitaker et al. (2012). A tight relationship can be seen, in this case with a scatter of 0.34 dex, which moves to higher SFRs at higher redshifts.*

analysis of this relationship was conducted by Brinchmann et al. (2004), who reported a clear correlation between the stellar masses and star-formation rates of galaxies in the local Universe ($z < 0.2$). This analysis was extended to $0.2 < z < 1.1$ by Noeske et al. (2007), who found evidence for a star-forming sequence, with SFR approximately proportional to stellar mass (SFR $\propto M_*$). This sequence was observed to be fairly tight, with a scatter of ~ 0.3 dex, and to evolve towards higher SFRs at fixed stellar mass with increasing redshift.

Fig. 1.16 shows the star-forming sequence reported by Whitaker et al. (2012), who extended earlier analyses to $z = 2.5$, again finding a tight relationship that evolves towards higher SFRs at higher redshifts. A large number of results from the literature across a redshift range from $0 < z < 6$ were compiled by Speagle et al. (2014), finding good agreement between different studies once selection effects had been accounted for. Speagle et al. (2014) report a best-fitting, time-dependent relationship between stellar mass and SFR of

$$\log_{10}(\mathrm{SFR}) \;=\; (0.84 - 0.026\,t)\,\log_{10}(M_*) \;-\; (6.51 - 0.11\,t) \qquad (1.31)$$

where t is the age of the Universe in Gyr. This could be thought of as describing the SFH of the average galaxy, with several studies reporting a scatter of ~ 0.3 dex in SFR, which remains relatively constant across cosmic time.

This strong evolution in typical galaxy star-formation rates across cosmic time is not thought to be related to changes in the feedback processes that inhibit star-formation, with the relationship shown in Fig. 1.14 found to be relatively constant across cosmic time (e.g. Conroy & Wechsler 2009). Instead, it is thought that the decline in star-formation efficiencies towards lower redshifts is a consequence of the accretion of gas onto halos becoming less efficient over time (e.g. Wechsler et al. 2002; Behroozi et al. 2013a). Overall, the process of star-formation is relatively inefficient, with < 20 per cent of baryons having been incorporated into stars by the present day.

1.5.4 The galaxy colour bimodality and quenching

Whilst it had long been realised that galaxies appeared to be broadly separated into two different kinds (e.g. Hubble 1926), this was not put onto a firm quantitative basis until the advent of large, systematic sky surveys at the beginning of the 21$^{\rm st}$ Century. The rest-frame optical colours of a large and representative sample of local galaxies were first studied by the *Sloan Digital Sky Survey* (SDSS) in the early 2000s, revealing a clear bimodality (e.g. Strateva et al. 2001; Blanton et al. 2003; Kauffmann et al. 2003).

As discussed in Section 1.3.1, the optical colours of galaxies are largely determined by the level of ongoing star-formation. The discovery of a clear bimodality therefore suggests that there is a distinction between two different types of galaxies: blue galaxies that are experiencing ongoing star-formation, and red galaxies that are no longer forming new stars. Galaxies that are not forming stars are said to have quenched their star-formation, and are consequently referred to as quenched, quiescent or passive galaxies.

As well as being distinguished by their redder colours, quiescent galaxies are also found to be typically more massive (e.g. Kauffmann et al. 2003; Brinchmann et al. 2004), more likely to have spheroidal rather than disk-like morphologies (e.g. Papovich et al. 2012; Bell et al. 2012; van der Wel et al. 2014), and more likely to inhabit dense environments (e.g. Baldry et al. 2006; Hartley et al. 2010). One projection of this bimodality is shown in Fig. 1.17. These findings are often

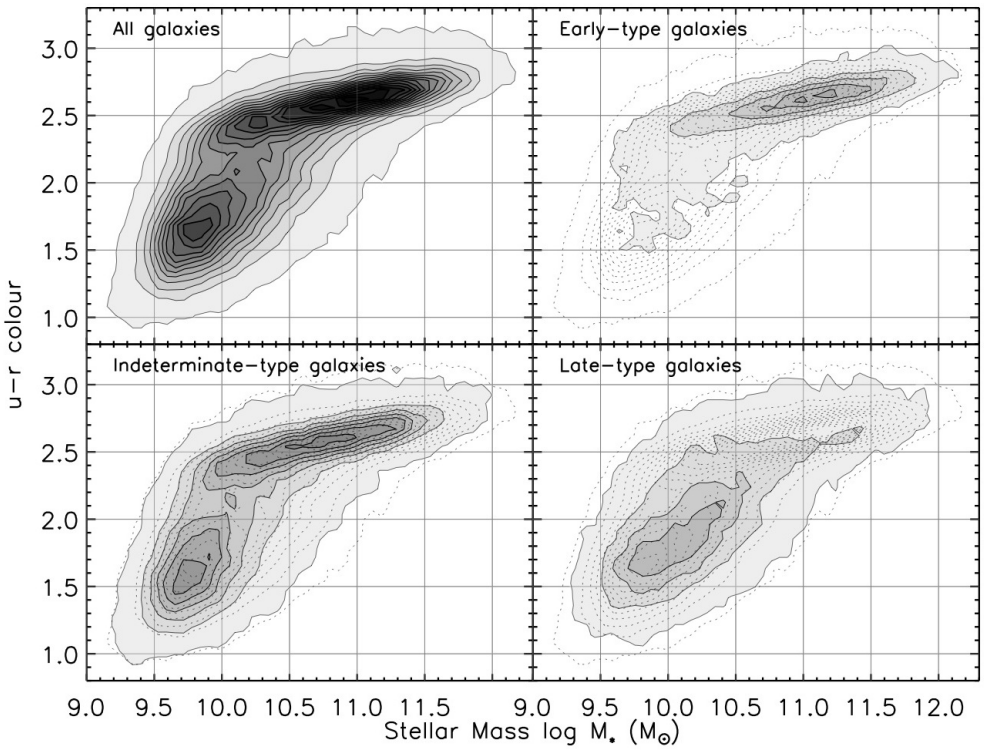

Figure 1.17 *The galaxy colour bimodality, shown on a plot of $u - r$ colour vs stellar mass. The top-left panel shows all galaxies, with a clear red sequence at the top, blue cloud at the bottom, and sparsely-populated green valley in between. The other panels show galaxies split by morphology, with early-type (elliptical) galaxies more likely to be on the red sequence, and late-type (spiral) galaxies more likely to be in the blue cloud. (Credit: C. Mihos)*

taken to support the idea that quiescent galaxies are subjected to some specific physical effect, or effects, which result in their star-formation shutting down. These effects are referred to as quenching mechanisms.

To study the differences between star-forming and quiescent galaxies, a robust means of dividing the two populations is necessary. Williams et al. (2009) introduced a set of selection criteria based upon the rest-frame UVJ colours of galaxies. This selection diagram is shown in Fig. 1.18 with the distribution of galaxies shaded and the proposed selection criteria shown with solid lines. The Williams et al. (2009) selection criteria are redshift-dependent, with the diagonal line moving downwards (towards bluer colours) with increasing redshift.

Cosmological simulations can be used to probe quenching physics by attempting to reproduce the observed properties of the quiescent population, such as observed trends in colour, metallicity and age with stellar mass. Significant attention has focussed on understanding the contribution of AGN feedback, which has been

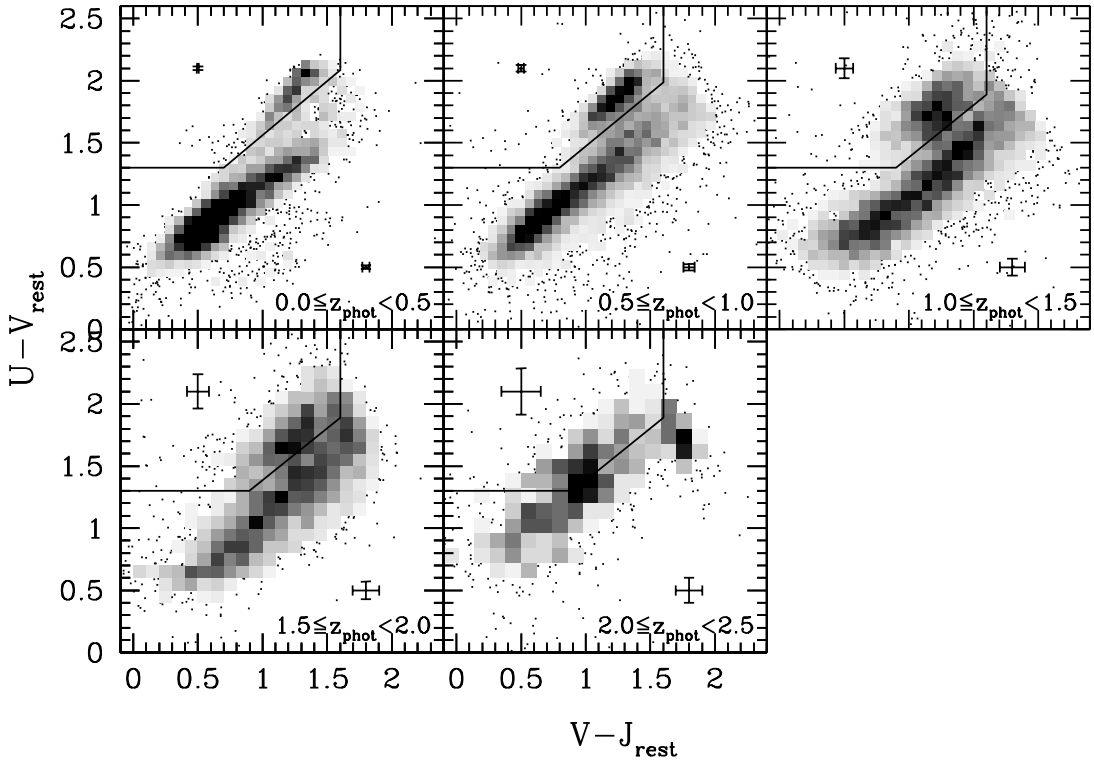

Figure 1.18 *The rest-frame UVJ colour selection diagram, or UVJ diagram, taken from Williams et al. (2009). The shading shows the density of galaxies, which can be seen to be bimodal. Star-forming galaxies fall onto a sequence running from the bottom-left to top-right. Quiescent galaxies are above and to the left of the star-forming sequence. The solid lines show the proposed UVJ selection criteria.*

shown to be important for reproducing many of these observed properties (e.g. Croton et al. 2006; Somerville et al. 2008; Trager & Somerville 2009).

1.5.5 Metal enrichment

As discussed in Sections 1.1.2 and 1.3.1, the gas from which galaxies initially begin to form is almost entirely pristine hydrogen and helium, which is later polluted by metals from successive generations of star-formation. Two main methods exist for directly probing this metal enrichment process. Firstly, the average metallicity of the whole stellar population can be probed by observing stellar absorption features in galaxy spectra (e.g. Worthey et al. 1994). Secondly, the metallicity of stars being formed at the epoch of observation can be probed by considering emission line ratios, since these lines are produced by hot gas in star-forming regions (see Section 1.3.3).

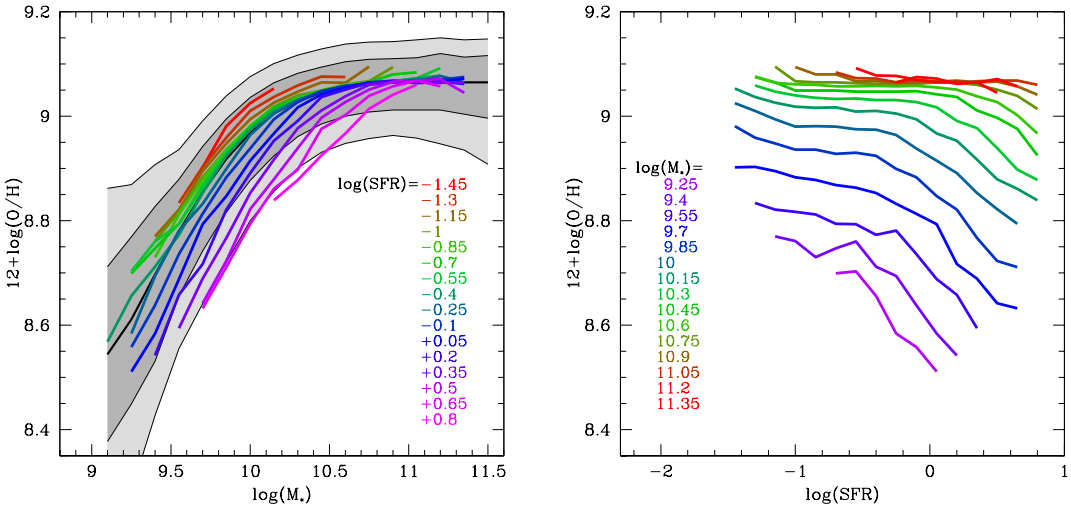

Figure 1.19 *Two projections of the fundamental metallicity relationship between stellar mass, star-formation rate and gas-phase metallicity, taken from Mannucci et al. (2010). This relationship is found to be redshift-independent, and is thought to be fundamentally linked to the underlying physics of galaxy formation, specifically the depth of the gravitational potential well and quantity of inflowing gas.*

These two metallicities are often referred to as stellar and nebular (or gas-phase) metallicities respectively. Historically, stellar metallicities have most often been derived for quiescent galaxies, as spectra of local high-metallicity, low-mass stars are readily available for comparisons with old stellar populations. Conversely, gas-phase metallicities can only be derived for galaxies with detectable levels of ongoing star-formation. A detailed review is provided by Maiolino & Mannucci (2019).

Stellar and nebular metallicities were first studied on a large scale in the local Universe using SDSS data by Tremonti et al. (2004) and Gallazzi et al. (2005). The key finding was a strong relationship between stellar mass and metallicity, with steep evolution found at lower masses, which becomes flatter towards higher masses. This is interpreted as evidence that higher-mass galaxies retain more of the metals they produce, whereas lower-mass galaxies lose their metals to the IGM before they can be incorporated into new stars. These studies have now begun to be extended towards higher redshifts (e.g. Maiolino et al. 2008; Gallazzi et al. 2014; Cullen et al. 2014; Steidel et al. 2016), with evolution found towards lower metallicities at higher redshifts.

An alternative perspective on the mass-metallicity relationship was proposed by

Mannucci et al. (2010), who identified a key additional dependence on star-formation rate. They found that galaxies occupy a narrow plane in mass-metallicity-SFR space, with more-highly star-forming galaxies displaying lower gas-phase metallicities. This 3D relationship is referred to as the fundamental metallicity relation (FMR). Two projections of the FMR are shown in Fig. 1.19.

Crucially, Mannucci et al. (2010) find that this relationship does not evolve with observed redshift, at least as far as $z \sim 2$, with the decrease in gas-phase metallicities found with increasing redshift a consequence of higher SFRs at earlier epochs. This relationship is interpreted as being related to fundamental underlying physical mechanisms in galaxy formation. Higher star-formation rates are thought to be associated with greater inflows of low-metallicity gas from the IGM, which dilute the metal content of star-forming gas in galaxies, whereas higher stellar masses are thought to be associated with deeper gravitational potential wells that are better able to retain their metals. These ideas are sometimes referred to as the gas-regulator, or bathtub model (Lilly et al. 2013).

1.5.6 Dust attenuation

Another potential probe of the metallicities of galaxies comes from the levels of dust attenuation their stellar populations are subjected to. Whilst dust is a less-direct tracer of metallicity, levels of dust attenuation also have important consequences for other key questions in astronomy. For example, as described in Section 1.5.2, a complete census of the cosmic SFRD requires an understanding of how much ongoing star-formation is obscured by dust, and how this evolves as a function of redshift.

In the local Universe, the dependencies of dust attenuation on parameters such as stellar mass, SFR and metallicity are well studied (e.g. Heckman et al. 1998; Hopkins et al. 2001; Brinchmann et al. 2004). A detailed analysis of all of these factors was conducted using SDSS data by Garn & Best (2010), who found that the main factor determining dust attenuation is stellar mass, with little secondary dependence on SFR or metallicity.

Current research focusses on extending our understanding of dust attenuation to high redshifts. Several recent studies have claimed to observe steeper dust attenuation curves at high redshift than in the local Universe (see Section 1.3.2), consistent with the SMC extinction curve (e.g. Reddy et al. 2015, 2018). However

this methodology has recently been called into question, with other studies finding average attenuation curves consistent with the Calzetti law out to $z = 5$ (e.g. Cullen et al. 2017, 2018; McLure et al. 2018b).

1.6 Thesis outline

The following three chapters constitute the scientific contents of this thesis. Each chapter is based upon a separate journal article published during the course of my studies, all of which are focussed on modelling observational data in terms of galaxy physical parameters for the purpose of star-formation history recovery. Chapters 2 and 4 have a particular emphasis on massive quiescent galaxies.

In Chapter 2, I introduce Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation (Bagpipes), a state of the art code for generating realistic model galaxy spectra and fitting these to spectroscopic and photometric observations within a Bayesian statistical framework. I then use Bagpipes to fit a large sample of multi-band photometric data for $\sim 10,000$ massive quiescent galaxies from the UltraVISTA survey across a redshift range from $0.25 < z < 3.75$, analysing the properties of their star-formation histories as a function of stellar mass and observed redshift.

In Chapter 3, I conduct a detailed investigation into the potential biasing effects of four different parametric star-formation-history models on galaxy physical parameters inferred from multi-band photometry. As part of this analysis, Bagpipes is used to fit a volume-complete sample of galaxies with multi-band photometric data from the GAMA survey at redshifts of $0.05 < z < 0.08$, with the aim of reconstructing the redshift evolution of the cosmic star-formation-rate density.

In Chapter 4, I develop a framework within Bagpipes for the joint fitting of photometric and spectroscopic data, including a model to account for wavelength-dependent systematic uncertainties in spectroscopic data. I then apply this model to a sample of 75 massive quiescent galaxies with redshifts from $1.0 < z < 1.3$, fitting both multi-band photometric data and extremely deep rest-frame near-ultraviolet spectroscopy from VANDELS. This higher-quality data allows a ed investigation of the star-formation histories and quenching properties of these galaxies.

Following these chapters, I present a summary of my conclusions and discuss potential extensions to my work in Chapter 5.

Chapter 2

Inferring the star-formation histories of massive quiescent galaxies with BAGPIPES: Evidence for multiple quenching mechanisms

The material in this chapter was originally published in Carnall et al. (2018).

2.1 Introduction

Recent years have seen rapid advances in the quantity and quality of data available for the study of galaxy evolution. Where once studies of large samples of objects relied solely on photometric data, the advent of multi-object and integral-field spectrographs, and the advancing near-infrared capability of both spectroscopic and photometric instruments mean that increasingly detailed studies can now be performed on statistical samples of objects. This presents new challenges, as ever more complex models must be developed to accurately reproduce observed properties, and more advanced statistical techniques are necessary to deal with larger sample sizes and the higher-dimensional, often highly degenerate parameter spaces of more complex models. Techniques from the field of data science for dealing with large datasets, such as machine learning, are also becoming increasingly valuable to astronomers.

One of the most striking observations in the field is the bimodal distribution in rest-frame colours that galaxies have been observed to display across at least the $10 - 11$ billion years since redshift, $z = 2 - 3$ (e.g. Strateva et al. 2001; Bell et al. 2004; Faber et al. 2007; Williams et al. 2009; Brammer et al. 2009; Whitaker et al. 2011, 2013; Straatman et al. 2014, 2016). When plotted on a colour-magnitude diagram, this bimodality manifests as a diffuse blue cloud and tight red sequence, with a sparsely populated green valley in between. In addition, the same bimodality has been observed in several other properties, such as environment (e.g. Baldry et al. 2006; Hartley et al. 2010; Chuter et al. 2011; Law-Smith & Eisenstein 2017) and morphology (e.g. Papovich et al. 2012; Bell et al. 2012; Strazzullo et al. 2013; van der Wel et al. 2014), with blue, star-forming galaxies exhibiting a more disk-like morphology and being less strongly clustered than red, quiescent galaxies, which typically exhibit early-type morphologies.

The observation of this bimodality has led to the search for a quenching mechanism, or mechanisms, capable of shutting down star-formation activity in galaxies. Many potential quenching mechanisms have been proposed, falling broadly into two categories (e.g. Peng et al. 2015). These can be described as ejective, associated with the expulsion of the gas reservoir that fuels star formation, and preventative, associated with the supply of new gas to the galaxy being shut down. The former results in a rapid halting of star formation, whereas the latter leads to a more gradual decline, or strangulation, as the existing reservoir is depleted.

The most commonly discussed process related to the quenching of massive galaxies is active galactic nucleus (AGN) feedback, in which the accretion of gas onto a galaxy fuels an AGN. This, in turn, either ejects the gas reservoir from the galaxy, or prevents further gas from being accreted, depending on the strength of the feedback. Strong AGN-driven outflows, sometimes called quasar-mode feedback, have been observed at high redshifts (e.g. Maiolino et al. 2012; Förster Schreiber et al. 2014), and are thought, in some cases, to be triggered by major-merger events. However, it has been shown that galaxies quenched in this way are likely to re-ignite star formation (e.g. Gabor et al. 2010), rather than remaining quiescent.

Sustained quiescence is thought to require a continuing input of energy to heat the circumgalactic medium, preventing new gas from falling onto the galaxy (e.g. Gabor et al. 2011). Low-accretion-state AGN feedback, sometimes called radio-mode or jet-mode, has been proposed (Croton et al. 2006) as the source of this

energy input. This mechanism has been shown to significantly improve agreement with the observed properties of the red sequence when implemented in modern hydrodynamic simulations, such as MUFASA (Davé et al. 2017) and IllustrisTNG (Nelson et al. 2018).

Another example of an ejective quenching process is the tidal/ram-pressure stripping experienced by galaxies falling into clusters. This process is also capable of causing very rapid quenching of star formation, and is thought to to give rise to the environmental dependence of the galaxy colour bimodality (e.g. Peng et al. 2012). Rapid quenching of satellites is also observed in simulations such as Illustris (e.g. Diemer et al. 2017).

However, recently Abramson et al. (2016) have shown that no qualitatively different processes are required to have acted on different galaxies to explain the galaxy colour bimodality. In their model, quiescence is the end-point of a single evolutionary track along which all galaxies progress at different rates. Star formation in individual galaxies dies down naturally in the same way as the cosmic star formation rate (SFR) density. Under this paradigm, we should not necessarily expect to observe the imprints of quenching processes in the properties of quiescent galaxies.

In order to understand whether, and to what extent, different potential quenching processes contribute to the termination of star formation, much tighter observational constraints on the way quenching behaves will be necessary. Two fundamental questions I will seek to address in this chapter are:

1. When did quenched galaxies form their stellar masses?
2. How long did the process of quenching take?

It is already known that a complete answer to these questions must include a description of the dependencies on physical properties such as size and mass, and the effects of different environments. Differences in the properties of quenched galaxies observed at different redshifts may also provide useful insights. In this chapter I will focus on observed-redshift and stellar-mass dependencies for massive quenched galaxies ($\boldsymbol{M}_* > 10^{10}$ M$_\odot$).

Question (i) has been studied by extracting the star formation histories (SFHs) of local galaxies from the fossil records imprinted in their spectra (e.g. Heavens et al. 2004; Panter et al. 2007; Thomas et al. 2010; Carson & Nichol 2010; Citro

et al. 2016). More recently this analysis has been extended to higher redshifts (e.g. Moresco et al. 2010; Onodera et al. 2012, 2015; Jørgensen & Chiboucas 2013; Gallazzi et al. 2014; Choi et al. 2014; Lonoce et al. 2014; Fumagalli et al. 2016; Pacifici et al. 2016; Siudek et al. 2017) to connect local galaxies with their precursors at earlier epochs, and build up a consistent picture of quenching across cosmic time.

Whilst individual results for the derived ages of quenched galaxies are difficult to compare due to differences in the definition of age, a coherent picture has been established, often referred to as downsizing or mass-accelerated evolution, with more massive galaxies typically forming their stellar masses at earlier epochs than their less massive counterparts. An average trend towards later formation for samples observed at lower redshifts is also seen, indicating that the assembly of the red sequence is still ongoing, in agreement with the observed evolution of the galaxy stellar mass function (e.g. Tomczak et al. 2014).

Question (ii) however is less well understood. Barro et al. (2013) and Schawinski et al. (2014) present evidence for a scheme in which there are both fast and slow tracks towards quiescence, with different mechanisms driving quenching. Massive quiescent galaxies at earlier epochs ($z \gtrsim 1.5$) are denser systems, colloquially known as 'red nuggets', which quench rapidly through strong AGN-driven outflows, whereas a second population of less-dense star-forming galaxies quench more gradually as a result of strangulation processes. The denser, earlier-quenching systems are theorised to gradually expand through processes such as minor mergers to leave one population of relatively large quiescent galaxies by the present epoch (e.g. McLure et al. 2013).

The results of Abramson et al. (2016) however show that it is not necessary to assume a change in quenching mechanism to obtain early-time rapid and late-time slower quenching. This property is reproduced by the model of Gladders et al. (2013), which assumes no physics, as part of a more general relationship, in which galaxies cross the green valley between star-forming and quiescent states in $\simeq 20$ per cent of the age of the Universe at the epoch of their quenching. A similar result is also obtained by Pacifici et al. (2016), although they also argue that their results are inconsistent with a single quenching mechanism acting in all circumstances.

Finally, Peng et al. (2015) argue that a comparison of metallicities between local quiescent and star-forming galaxies strongly favours a slow strangulation of star

formation over a timescale of $\simeq$ 4 Gyr. Clearly, a direct method that could be demonstrated to reliably constrain in detail the SFHs of individual quiescent galaxies would greatly assist in reconciling these disparate results.

Methods employed to obtain this information rely on modelling the light emitted by galaxies as a function of wavelength in terms of the physical parameters of the system (e.g. Cid Fernandes et al. 2005; Thomas et al. 2017; Wilkinson et al. 2017). These models are then fitted to observational data, which may consist of spectroscopy and/or photometry (sometimes referred to as the spectral energy distribution; SED). Once the posterior distribution for model parameters has been obtained, nuisance parameters may be marginalised out to obtain constraints on the parameters of interest, such as the SFH. A common approach has been to fit spectroscopic data indirectly, either by data compression (e.g. Heavens et al. 2000), or by using pre-calibrated spectral indices (e.g. Worthey et al. 1994).

Recently, the development of new Bayesian statistical techniques, such as advanced Markov chain Monte Carlo (MCMC; e.g. Goodman & Weare 2010; Foreman-Mackey et al. 2013) and nested sampling algorithms (Skilling 2006, Feroz & Hobson 2008; Feroz et al. 2009, 2013) has begun to enable the efficient exploration of higher-dimensional parameter spaces to obtain posterior distributions for the parameters of complex models.

A new generation of modern spectral modelling and fitting tools has been developed in order to exploit this e.g. Beagle (Chevallard & Charlot 2016) and Prospector (Leja et al. 2017; Johnson et al. in prep). These codes allow on-the-fly generation and fitting of complex, self-consistent models to describe galaxies across continuous parameter spaces. Their models include emission and absorption processes due to the stellar population, ionized gas in H II regions, diffuse dust in the interstellar medium (ISM) and neutral gas in the intergalactic medium (IGM) along our line-of-sight. The ability to explore higher-dimensional parameter spaces also opens up the ability to fit more complex SFHs. This is important because it has been shown that the rigid, exponentially declining SFHs typically employed in SED fitting techniques can introduce significant biases into SFH parameter estimates (see Section 2.4).

The combination of new statistical techniques with the ever increasing volume and quality of data available means that it is now possible to study the behaviour of quenching processes in unprecedented detail. In this chapter I present Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation,

or Bagpipes, a new public galaxy spectral modelling framework and fitting tool written in the Python programming language. Bagpipes provides a highly intuitive application programming interface (API) for rapid, on-the-fly generation (up to hundreds of models per second) of complex, physically realistic model galaxy spectra across continuous parameter spaces. It also provides a tool, built around the MultiNest nested sampling algorithm (Feroz & Hobson 2008; Feroz et al. 2009, 2013), which allows for the direct fitting of these models to arbitrary combinations of spectroscopic and photometric data.

Using Bagpipes I then analyse photometric data for a large sample of quenched galaxies from the UltraVISTA Survey (McCracken et al. 2012). The recovered SFHs for these objects will be considered in the context of Questions (i) and (ii) posed above in order to understand galaxy quenching properties as a function of stellar mass and observed redshift across the majority of cosmic time, from $0.25 < z < 3.75$. In Chapter 4 I will extend this analysis using Bagpipes to the direct fitting of spectroscopic data from VANDELS (McLure et al. 2018a; Pentericci et al. 2018) to obtain stronger constraints.

The structure of this chapter is as follows. In Section 2.2 I introduce the UltraVISTA dataset I will use to explore quenching. In Section 2.3 I describe the model generation and fitting methodologies employed by the Bagpipes code. In Section 2.4 I consider different models within Bagpipes, in particular different SFH parameterisations, and test their ability to recover realistic SFHs by fitting mock observations of simulated quenched galaxies from the Mufasa suite of cosmological hydrodynamic simulations (Davé et al. 2016). By this process I define a model that is capable of recovering unbiased estimates of the SFH properties of large samples of quenched galaxies from photometric data, making use of a double-power-law SFH parameterisation. In Section 2.5, I apply this Bagpipes model to select a sample of 9289 quenched galaxies with $M_* > 10^{10}$ $M_\odot$ and redshifts in the range $0.25 < z < 3.75$ from UltraVISTA. In Section 2.6 I analyse the properties of their SFHs to constrain the epoch and duration of their quenching as a function of stellar mass and observed redshift. I present the conclusions of this chapter in Section 2.7.

Throughout this chapter, I will distinguish between times, t, which are measured forwards from the beginning of the Universe (i.e. $t(z)$ is the age of the Universe at redshift z), and ages, a, which are measured backwards in time from the redshift of observation, $t(z_{\rm obs})$.

Table 2.1 *Mean 5σ limiting magnitudes within 2″ apertures for the 1 deg^2 twelve band catalogue of Mortlock et al. (2017).*

Region	Percentage of area	u^*	g'	r'	i'	z'	$z'_{\rm Subaru}$
Deep strips	40 per cent	27.0	27.1	26.6	26.3	25.4	26.4
Wide strips	60 per cent	27.0	27.1	26.6	26.3	25.4	26.4
		Y	J	H	K_s	IRAC1	IRAC2
		25.1	24.9	24.6	24.8	25.3	25.1
		24.7	24.4	24.1	23.9	25.3	25.1

2.2 The UltraVISTA data

The UltraVISTA Survey (McCracken et al. 2012) is an ultra-deep imaging programme over a 1.5 deg^2 contiguous area of the Cosmological Evolution Survey (COSMOS) field in the near-infrared $YJHK_s$ bands on the Visible and Infrared Survey Telescope for Astronomy (VISTA). The area is divided into deeper and shallower stripes (UltraVISTA deep and UltraVISTA wide, respectively) that each account for half of the total area. The data utilised here comes from the UltraVISTA DR3 release.

In this chapter I make use of the K-band selected catalogue compiled by Mortlock et al. (2017), which covers the 1 deg^2 overlap region between UltraVISTA and the Canada-France-Hawaii Telescope Legacy Survey (CFHTLS; Hudelot et al. 2012) T0007 release, which provides optical $u^*g'r'i'z'$ band imaging. These datasets are also combined with deep z'-band imaging from the Subaru telescope (Furusawa et al. 2016) and 3.6 μm + 4.5 μm imaging from *Spitzer*/IRAC (Ashby et al. 2013; Steinhardt et al. 2014) to produce a final twelve-band catalogue. The process by which the catalogue is compiled is described in detail in sections 2.1, 2.3 and 2.5 of Mortlock et al. (2017). Table 2.1 gives the mean depths for each of the twelve bands across the wide and deep regions, which respectively make up 60 per cent and 40 per cent of the 1 deg^2 overlap region.

Mortlock et al. (2017) also calculate highly robust photometric redshifts for their whole catalogue using the median values from five different photometric-redshift codes. These median photometric redshifts have $\sigma_{dz} \simeq 0.02$, where $dz = (z_{spec} - z_{phot})/(1 + z_{spec})$, and a catastrophic outlier fraction ($|dz| > 0.15$) of $\sim 1\%$ (see their section 3 for more details). They also calculate mass-completeness limits as a function of redshift and clean the catalogue by performing star-galaxy separation,

Table 2.2 *List of input parameters for the example* Bagpipes *model shown in Fig. 2.1. All parameters are discussed in Section 2.3 except for α, β and τ for the double-power-law SFH parameterisation, the functional form for which is given in Equation 2.11. Here I assume Solar metallicity to take the value $Z_\odot = 0.02$.*

Global	Double-power-law	Dust	Nebular
$\sigma_{\rm vel} = 300$ km s^{-1}	$\log_{10}(M_{\rm formed}/{\rm M}_\odot) = 11$	Calzetti	$\log_{10}(U) = -3$
$a_{\rm BC} = 0.01$ Gyr	$Z = 0.8\ Z_\odot$	$A_V = 0.2$	
$z = 0$	$\tau = 12$ Gyr	$\epsilon = 3$	
	$\beta = 0.5$	$T = 30$ K	
	$\alpha = 30$	$\beta = 1.5$	

and by matching to radio and X-ray datasets to remove AGN contaminants. The combination of all of these high-quality datasets provides an extremely clean catalogue of unrivalled scope and depth for studying galaxy evolution.

2.3 The BAGPIPES code

Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation, or Bagpipes, is a Bayesian spectral fitting code, designed to model the emission from galaxies from the far-ultraviolet to the microwave regimes, and to fit these models to arbitrary combinations of spectroscopic and photometric observational data using the MultiNest nested sampling algorithm (Feroz & Hobson 2008; Feroz et al. 2009, 2013). Bagpipes is written purely in the Python programming language and considerable effort has been made to make the API as intuitive and user-friendly as possible. This section describes the Bagpipes code. In Section 2.3.1 I describe how models are generated by the code, then in Section 2.3.2 I describe the fitting of these models to observational data. Bagpipes is publicly available and fully documented at https://bagpipes.readthedocs.io.

2.3.1 Model generation

The first major aspect of Bagpipes is its ability to rapidly generate physically realistic model galaxy spectra from the far-ultraviolet to the microwave regime. Models can be built up to the desired level of complexity by specifying a number of model components, for example dust and nebular-emission prescriptions and SFH components. Bagpipes uses these components (in practice these are passed

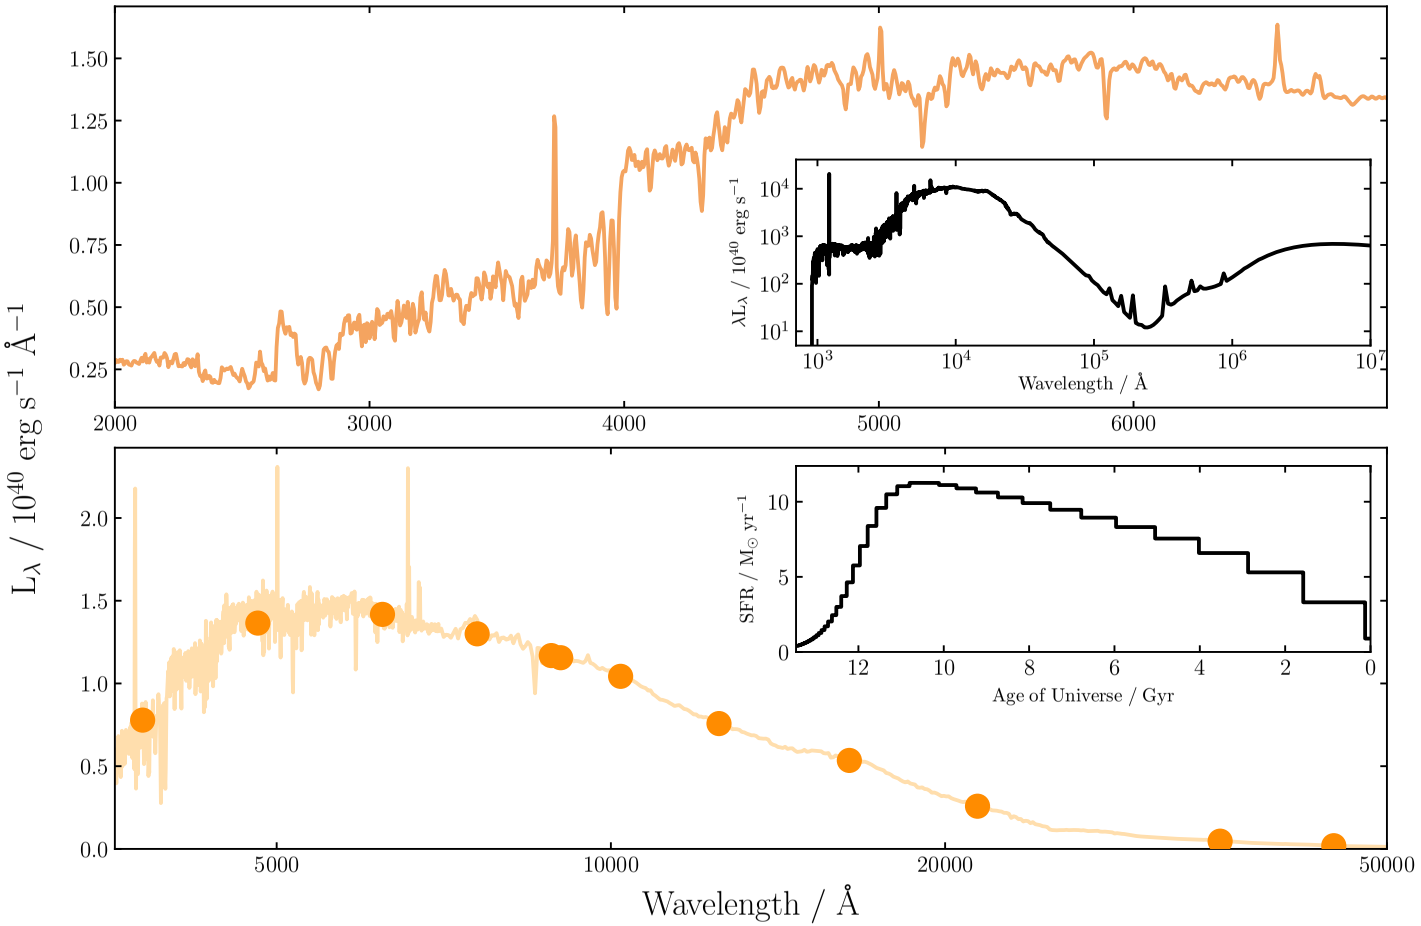

Figure 2.1 *An example of a model galaxy built with* Bagpipes. *The parameters of the model are given in Table 2.2 and are designed to be representative of a present-day quenched galaxy. The top panel shows the spectral output, in this case over the range* 2000–7000 *Å with a sampling of 5 Å. The top inset panel shows the full spectrum for the model. The bottom panel shows the photometric output, in this case fluxes observed through the twelve filters of the UltraVISTA catalogue (orange circles) described in Section 2.2. The bottom inset panel shows the SFH of the galaxy, which has been modelled as a double-power-law.*

to the API as PYTHON dictionaries containing parameter values) to generate an internal model of a galaxy spectrum. From this model, the user can request spectral data covering a given wavelength range at a given sampling, photometric fluxes through a series of user-defined filters, and emission-line fluxes. An example model is shown in Fig. 2.1. The parameters of this model are listed in Table 2.2.

A model galaxy spectrum in BAGPIPES observed at redshift $z_{\rm obs}$ has a luminosity per unit rest-frame wavelength, λ of $L_\lambda(\lambda)$. The luminosity is constructed as a sum involving four ingredients:

1. Simple stellar-population models, SSP(a, λ, Z), which are a function of λ, the age of the stellar population, a, stellar metallicity, Z and the initial mass function (described in Section 2.3.1).

2. The star-formation history, SFR(t), which is composed of a sum over one or more SFH components (described in Section 2.3.1).

3. The transmission function of the ionized ISM, $T^+(a, \lambda)$, as defined by Charlot & Longhetti (2001), including absorption, line emission, ionized continuum emission and emission from warm dust within H II regions (described in Section 2.3.1).

4. The transmission function of the neutral ISM, $T^0(a, \lambda)$, due to diffuse dust attenuation and emission (described in Section 2.3.1).

These four ingredients are described in more depth in the following four sections. The ingredients are then combined to give the galaxy luminosity according to the relationship

$$L_\lambda(\lambda) \; = \; \sum_{j=1}^{N_c} \sum_{i=1}^{N_a} \mathrm{SFR}_j(t_i) \; \mathrm{SSP}(a_i, \lambda, Z_j) \; T^+(a_i, \lambda) \; T^0(a_i, \lambda) \; \Delta a_i \tag{2.1}$$

where i runs across the age bins used in BAGPIPES (see Section 2.3.1), Δa_i are the widths of these bins, j runs across SFH components (see Section 2.3.1), N_c is the number of SFH components and N_a the number of age bins. The distinction between times and ages is explained at the end of Section 2.1; here $t_i = t(z_{\rm obs}) - a_i$.

Once the galaxy luminosity has been calculated, it is redshifted and converted into an observed frame flux density, $f_{\lambda_{\rm obs}}(\lambda_{\rm obs})$, where $\lambda_{\rm obs} = (1 + z_{\rm obs})\lambda$, using the relationship

$$f_{\lambda_{\rm obs}}(\lambda_{\rm obs}) = \frac{L_\lambda(\lambda)}{4\pi D_L(z_{\rm obs})^2(1 + z_{\rm obs})} T_{\rm IGM}(\lambda, z_{\rm obs}), \tag{2.2}$$

where $D_L(z_{\rm obs})$ is the luminosity distance to redshift $z_{\rm obs}$ and $T_{\rm IGM}(\lambda, z_{\rm obs})$ is the transmission function of the IGM, described in Section 2.3.1. Finally, if an output spectrum is requested (as opposed to photometry or emission-line fluxes), $L_\lambda(\lambda)$ can be convolved with a Gaussian kernel (before applying Equation 2.2) to model the effects of velocity dispersion, as described in Section 2.3.1.

Stellar population synthesis

Stellar Population Synthesis (SPS) is not implemented directly in Bagpipes. Instead, the code is designed to accept pre-defined SPS models in the form of grids of simple stellar-population (SSP) models of different ages across a range of metallicities.

The SPS models currently implemented in the code are the 2016 version of the Bruzual & Charlot (2003) (hereafter BC03) models. These differ from earlier versions by their use of the Medium-resolution Isaac Newton Telescope library of empirical spectra (MILES; Falcón-Barroso et al. 2011) in the UV-optical spectral region. The models implemented within the code are constructed using a Kroupa (2001) initial mass function (IMF). Bagpipes does not currently include the option to vary element abundance patterns (e.g. alpha enhancement), with the model set currently implemented having scaled Solar abundances. A version of the the code that includes the Binary Population and Stellar Synthesis (BPASS, Eldridge & Stanway 2009) models is also available.

When a set of SPS models is loaded by Bagpipes, the models are resampled in age (using a weighted summation method) onto a grid of ages, a_i, from $\log_{10}(a_i/{\rm Gyr}) = 6.0$ to 10.2 with uniform width in $\log_{10}(\Delta a_i/{\rm Gyr})$. The default spacing is 0.1 dex, which sets the value of N_a in Equation 2.1 to 43.

Star-formation histories

Star-formation histories in BAGPIPES are constructed from one or more components, j, each of which specifies some functional form for star-formation rate as a function of time, $\mathrm{SFR}_j(t)$. The total SFR is given by summing over these components, i.e.

$$\mathrm{SFR}(t) \; = \; \sum_{j=1}^{N_c} \mathrm{SFR}_j(t). \tag{2.3}$$

When $L_\lambda(\lambda)$ is calculated using Equation 2.1, the $\mathrm{SFR}_j(t)$ are evaluated for all $t_i = t(z_{\mathrm{obs}}) - a_i$ where a_i is less than $t(z_{\mathrm{obs}})$, i.e. the age of the stellar population is less than the age of the Universe at z_{obs}. For all t_i corresponding to ages greater than the age of the Universe, $\mathrm{SFR}_j(t_i)$ is set to zero.

One may specify an unlimited number of SFH components, each with an individual functional form (see Section 2.4.1 for further discussion). The options currently implemented are:

- Delta function
- Constant
- Exponentially declining (Equation 2.10)
- Delayed exponentially declining
- Lognormal
- Double-power-law (Equation 2.11)
- Custom (directly input an array of SFR values).

In this way, BAGPIPES can be used to generate a huge parameter space of possible SFHs, encompassing, for example, the non-parametric SFHs used by Leja et al. (2017) by the use of multiple constant components, or SFHs drawn from simulations of galaxy formation as in Pacifici et al. (2016). SFHs from simulations may be inputted either by assigning a burst component to each star particle (as in Section 2.4.3), or by loading tabulated SFHs from simulations directly into the code as custom SFH components.

Each component also requires the specification of a total mass of stars formed, $M_{\rm formed}$, by that component over its whole history, and a metallicity value, Z_j, which is generated by linearly interpolating the SSP(a_i, λ, Z_j) between different grids of models. Thus, metallicity evolution can be modelled by specifying different metallicities for multiple SFH components covering different epochs of cosmic time. Bagpipes also includes basic functionality for specifying a distribution function for the metallicities of stars in a galaxy as an alternative to linear interpolation between grids, as in Leja et al. (2017). In the future I will extend this scheme to allow full chemical-evolution histories to be specified, permitting metallicity distribution functions that are also a function of cosmic time.

Nebular emission

The nebular-emission model implemented in Bagpipes is constructed following the methodology of Byler et al. (2017), using the latest (2017) version of the Cloudy photoionization code (Ferland et al. 2017). H II regions are modelled with a spherical shell geometry of fixed radius. The nebular emission from a galaxy is assumed to be the sum of emission from H II regions of different ages, as in Charlot & Longhetti (2001). The metallicity of the ionized gas is assumed to be the same as that of the stars that produce the ionizing photons. The metallicity is scaled relative to the Solar abundances of Anders & Grevesse (1989), using the ISM depletion factors and Helium and Nitrogen scaling relations of Dopita et al. (2000). All models include dust grains using the "ISM" prescription within Cloudy, which has a grain size distribution and abundance pattern designed to reproduce the observed extinction properties for the ISM of the Milky Way.

Bagpipes includes functionality for computing grids of nebular-emission models corresponding to grids of input SSP models (these have been pre-computed for BC03). Cloudy is run using each SSP in turn as the input spectrum whilst varying the logarithm of the ionization parameter, U, in steps of 0.5 between $\log_{10}(U) = -4$ and -2. The ionization parameter is the ratio of the number of hydrogen-ionizing photons, $Q_{\rm H}$, to the number of hydrogen atoms,

$$U = \frac{Q_{\rm H}}{4\pi r^2 n_{\rm H}} \tag{2.4}$$

where $n_{\rm H}$ is the number density of hydrogen atoms and r is the radius of the

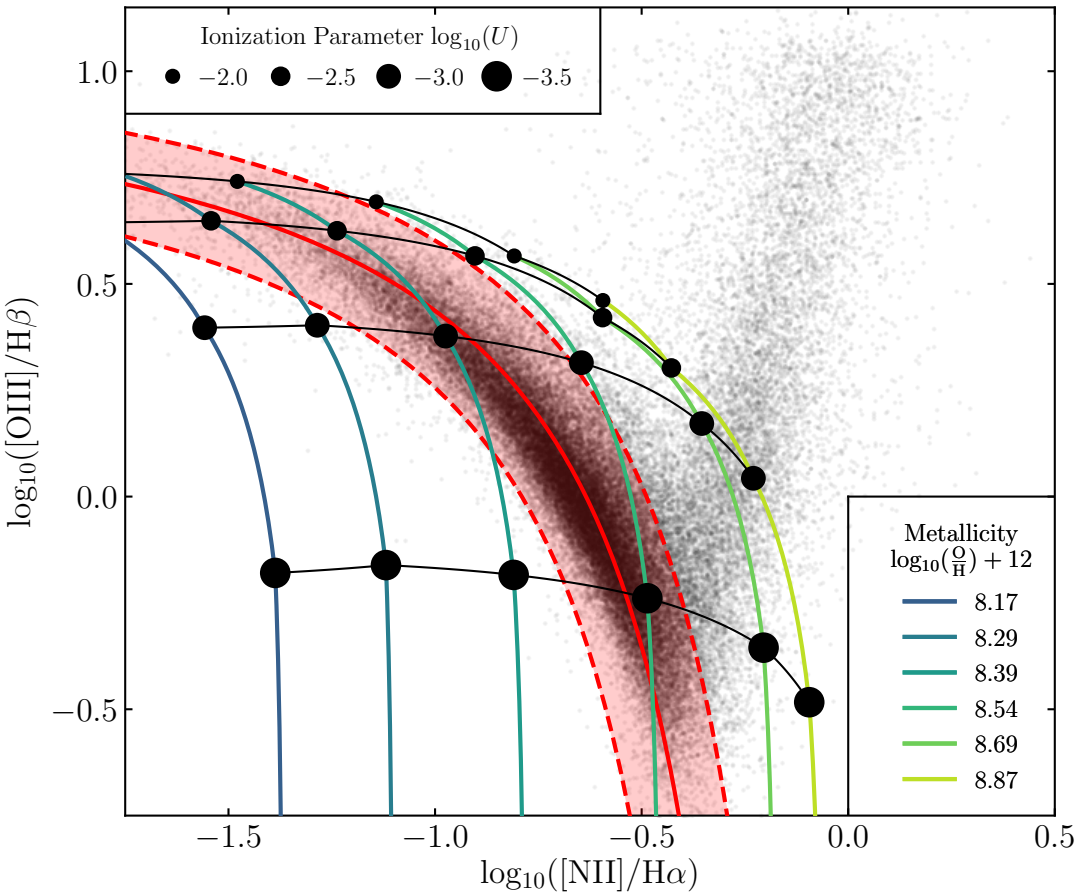

Figure 2.2 *BPT diagram with the positions of* Bagpipes *models overlaid as a function of metallicity and ionization parameter. Black lines are lines of constant ionization parameter and coloured lines are lines of constant metallicity. The theoretically calibrated star-forming sequence for galaxies at $z \sim 0$ of Kewley et al. (2013) is shown in red and a sample of local SDSS galaxies from Brinchmann et al. (2013) is shown in grey. The* Bagpipes *models are BC03-based, with constant SFR over the last 10 Myr. The models show good overall agreement with the models of Kewley et al. (2013) and the results of Moustakas et al. (2006), who find local star-forming galaxies have $-3.8 < \log_{10}(U) < -2.9$ and $8.15 < 12 + \log(\frac{O}{H}) < 8.7$.*

volume being considered.

Adjustments to U are made by varying $Q_{\rm H}$ at a fixed $n_{\rm H}$ of 100 atoms cm^{-3}. At each value of $\log_{10}(U)$, a_i and metallicity the diffuse continuum is recorded, which includes contributions from ionized gas and warm dust. Fluxes for a series of emission lines are also recorded.

The base list of lines employed is that given in table 3 of Byler et al. (2017), containing 128 separate features. The features have been renamed in the new version of Cloudy, due to slight shifts in wavelength, so it was necessary to manually match-up new and old labels. In addition, the five narrowly spaced C II lines around 2326Å were replaced with the total flux for a blend of these lines, as identified by 'Blnd 2326.00A', and the two narrowly spaced He I lines around 1.083 μm were replaced with 'TOTL 1.08303m'. Finally the He II feature at 4686Å, denoted in Cloudy by 'He 2 4685.64A' was added to the list, meaning that, in total, 124 separate emission features are tracked.

When including nebular emission in a Bagpipes model, one must specify the ionization parameter and lifetime of H II regions (or stellar birth-clouds), $a_{\rm BC}$. These are denser regions of leftover gas that surround newly formed stars, resulting in extra attenuation of the light emitted. The nebular-continuum and line models corresponding to all SSP(a_i, λ, Z_j) with $a_i < a_{\rm BC}$ are then linearly interpolated on-the-fly in $\log_{10}(U)$ and metallicity, then added to the corresponding SSP(a_i, λ, Z_j) model. Contrary to Equation 2.1 this is done by addition to, rather than multiplication of SSP(a_i, λ, Z_j), however the effect is the same as setting $T^+(a_i, \lambda) > 1$.

At this stage, the total energy in the combined (emission lines plus continuum) nebular model is set to the same as the total energy in the hydrogen-ionizing photons from all SSPs with $a_i < a_{\rm BC}$. The hydrogen-ionizing continuum is then removed from these SSPs (I set $T^+(a_i, \lambda) = 0$ for $a_i < a_{\rm BC}$ and $\lambda < 911.8$ Å; effectively assuming an escape fraction of zero) in order to maintain energy conservation. As well as being inserted into the output spectrum, observed fluxes for each emission feature are propagated through Bagpipes separately, meaning line fluxes can be accessed directly by the user.

Two consistency checks were performed on the grid of BC03-based nebular-emission models to demonstrate that they are in agreement with similar predictions from the literature. Firstly the optical line ratios as a function of metallicity and ionization parameter were plotted on the Baldwin-Phillips-

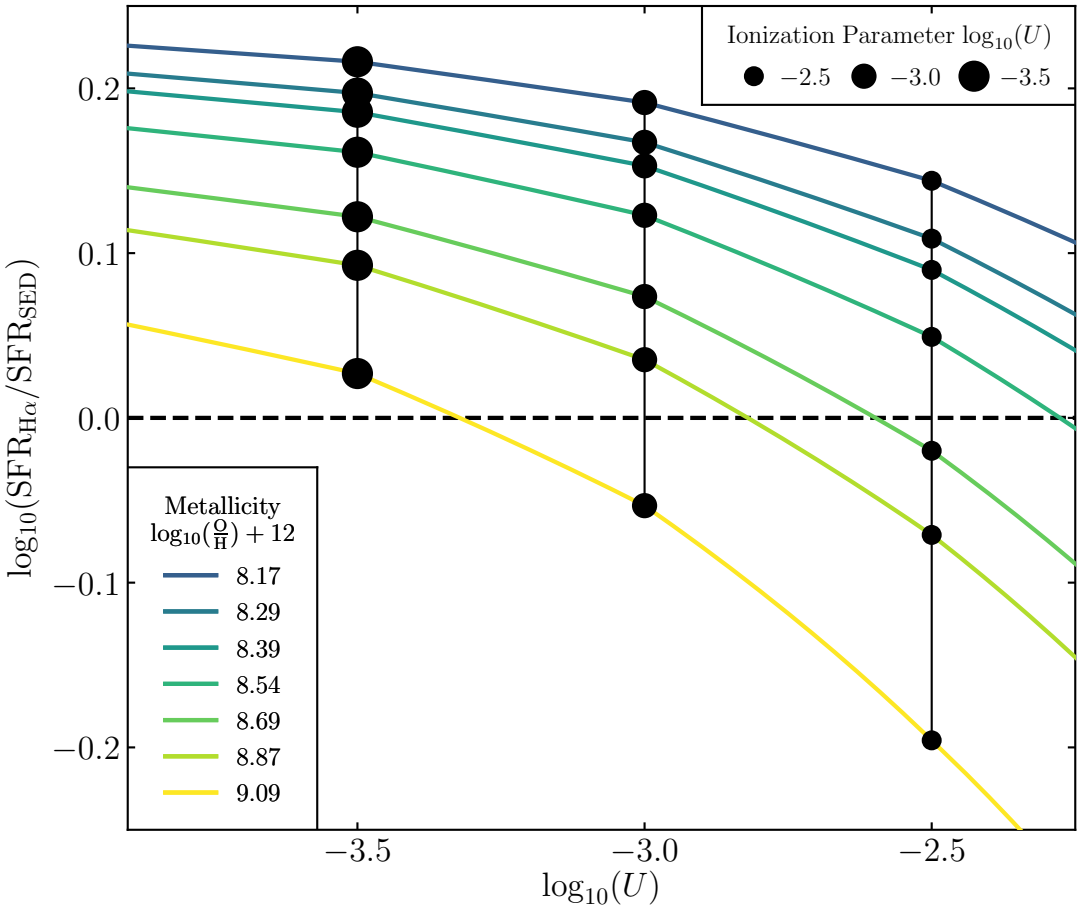

Figure 2.3 *The offset between input and* Hα*-derived SFRs for* Bagpipes *models as a function of metallicity and ionization parameter. The* Hα*-derived SFRs were calculated using the Kennicutt & Evans (2012) calibration. A slight excess of* Hα *flux is found in the* Bagpipes *models, compared to the fluxes predicted by Kennicutt & Evans (2012), except at supersolar metallicities or log-ionization parameters above* −2.5*. This is likely to be due to the assumption of an escape fraction of zero. The models are consistent to within the observed scatter on the Kennicutt & Evans (2012) relationship.*

Terlevich (BPT) diagram (Baldwin et al. 1981), shown in Fig. 2.2. The predictions can be seen to be consistent with those of Kewley et al. (2013) by comparison with their figure 1. It can also be seen that the range of metallicities and ionization parameters measured by Moustakas et al. (2006) for local galaxies would be correctly reproduced by these models.

Having confirmed this, the line strengths are also checked in Fig. 2.3 by comparing input star-formation rates, $\mathrm{SFR}_{\mathrm{SED}}$, with those derived from the strengths of the Hα lines in each model, using the calibration of Kennicutt & Evans (2012), $\mathrm{SFR}_{\mathrm{H}\alpha}$. It can be seen that that the Hα fluxes tend to be slightly higher than those predicted by Kennicutt & Evans (2012), e.g. by 0.08 to 0.19 dex for metallicities of $12 + \log\left(\frac{\mathrm{O}}{\mathrm{H}}\right) = 8.69$ and 8.17 respectively at $\log_{10}(U) = -3.0$. This is consistent to within the observed scatter on the relationship. The overestimation is likely to be due to the assumption of an escape fraction of zero; an escape fraction of $\sim$ 20 per cent would bring the $12 + \log\left(\frac{\mathrm{O}}{\mathrm{H}}\right) = 8.69$, $\log_{10}(U) = -3.0$ model into line with Kennicutt & Evans (2012).

Dust attenuation and emission

Dust attenuation in Bagpipes is designed in a modular fashion such that different dust curves may be easily implemented. Each dust-attenuation curve takes the overall functional form

$$\log_{10}\left(T^0(a_i, \lambda)\right) = \begin{cases} -\frac{0.4\ \epsilon\ A_V\ k(\lambda)}{R_V} & a_i < a_{\mathrm{BC}} \\ -\frac{0.4\ A_V\ k(\lambda)}{R_V} & a_i > a_{\mathrm{BC}} \end{cases} \tag{2.5}$$

where a_{BC} is the lifetime of stellar birth-clouds as described in Section 2.3.1, A_V is the attenuation in the V band ($\sim$ 5500 Å) in magnitudes, ϵ is a constant that can be used to control the extra attenuation towards H II regions and $k(\lambda)$ and R_V are specific to the dust model being used.

Three dust models are currently implemented: the Calzetti et al. (2000) law for local star-forming galaxies, the Cardelli et al. (1989) Milky Way dust law, and a flexible model based on that of Charlot & Fall (2000). For the first two of these models the form of $k(\lambda)$ and the value of R_V are specified in the paper cited. For use in Bagpipes, these are extrapolated to shorter wavelengths using a power-law fit to the dust curve in the near-ultraviolet. For the Charlot & Fall (2000)

model I define

$$\frac{k(\lambda)}{R_V} = (\lambda/5500\text{Å})^{-n} \tag{2.6}$$

where n is the slope of the attenuation law. The ϵ parameter is the reciprocal of the parameter μ often used in the literature (e.g. Cullen et al. 2017) for models of this type.

Dust emission from the neutral ISM is modelled as a single-temperature greybody, with flux per unit frequency, $S_{\rm gb}(\nu)$ (e.g. Hildebrand 1983; Younger et al. 2009) given by

$$S_{\rm gb}(\nu) \propto \nu^{\beta} B_\nu(T) = \frac{2h}{c^2} \frac{\nu^{\beta+3}}{e^{\frac{hc}{\lambda kT}} - 1} \tag{2.7}$$

where T is the temperature of the greybody, β is the spectral emissivity index and the dust is assumed to be optically thin. The normalisation of this component is set such that the total energy removed from the spectrum by dust attenuation is the same as the total energy emitted by the greybody.

Dust emission in Bagpipes is hence modelled by two separate components. Firstly, the hot-dust component included in the Cloudy output diffuse continua for H II regions (described in Section 2.3.1), and secondly the greybody component due to cold, diffuse dust emission. This kind of two-component approach has been shown to be successful in modelling observed infrared SEDs (e.g. Casey 2012), and work is ongoing to test this model against observational data at these wavelengths. In the future I hope to implement the option of a more complex physical model for dust emission.

IGM attenuation

For $T_{\rm IGM}(\lambda, z_{\rm obs})$, Bagpipes incorporates the IGM attenuation model of Inoue et al. (2014), an updated version of the Madau (1995) model. The analytic expression presented in their section 4 is calculated and tabulated for rest-frame wavelengths between 911.8 Å and 1215.7 Å, then interpolated by the code for use in model generation. I assume $T_{\rm IGM}(\lambda, z_{\rm obs}) = 0$ for any flux at $\lambda < 911.8$ Å, the majority of which is already removed when applying the nebular-emission prescription (see Section 2.3.1).

Velocity dispersion and spectral sampling

BAGPIPES is designed to perform spectroscopic fitting as well as fitting to photometry, meaning the effects of velocity dispersion, $\sigma_{\rm vel}$, must be included in the model to match observed spectral features. To facilitate this, BAGPIPES converts the wavelength sampling of all input SPS models to constant spectral resolution, $R = \frac{\lambda}{\Delta\lambda}$. By default, a low resolution of $R = 100$ is used over regions that will be used only for calculation of output photometry, and a higher resolution of $R = 1000$ is used over regions that will be used to generate the output spectrum (in practice, R is doubled to achieve Nyquist sampling at the resolutions quoted). The latter region is then convolved with a Gaussian kernel in velocity space to model the effect of velocity dispersion within the observed galaxy, which is assumed to be the same for the stellar and gas components.

2.3.2 Model fitting

The second major aspect of BAGPIPES is its ability to fit the models described in Section 2.3.1 to observational data within the framework of Bayesian inference. These data can take the form of an observed spectrum and/or any number of photometry points, all of which can be fitted simultaneously by the code. In this section I begin by outlining the principles of the BAGPIPES fitting method as applied to photometric observations. Additional considerations when fitting spectroscopic observations will be addressed in Chapter 4. I then outline how fitting is performed using the MULTINEST nested sampling algorithm. I finally provide a brief description of the process of specifying prior distributions within BAGPIPES.

Bayesian inference methodology

Bayesian inference methods use Bayes' theorem to update prior knowledge about the probability of a hypothesis, or model, based on new data. This can take the form of parameter estimation, in which constraints on the parameters of a given model are updated, or model selection, in which the relative probabilities of two or more different models are assessed. Bayes' theorem states that for some new

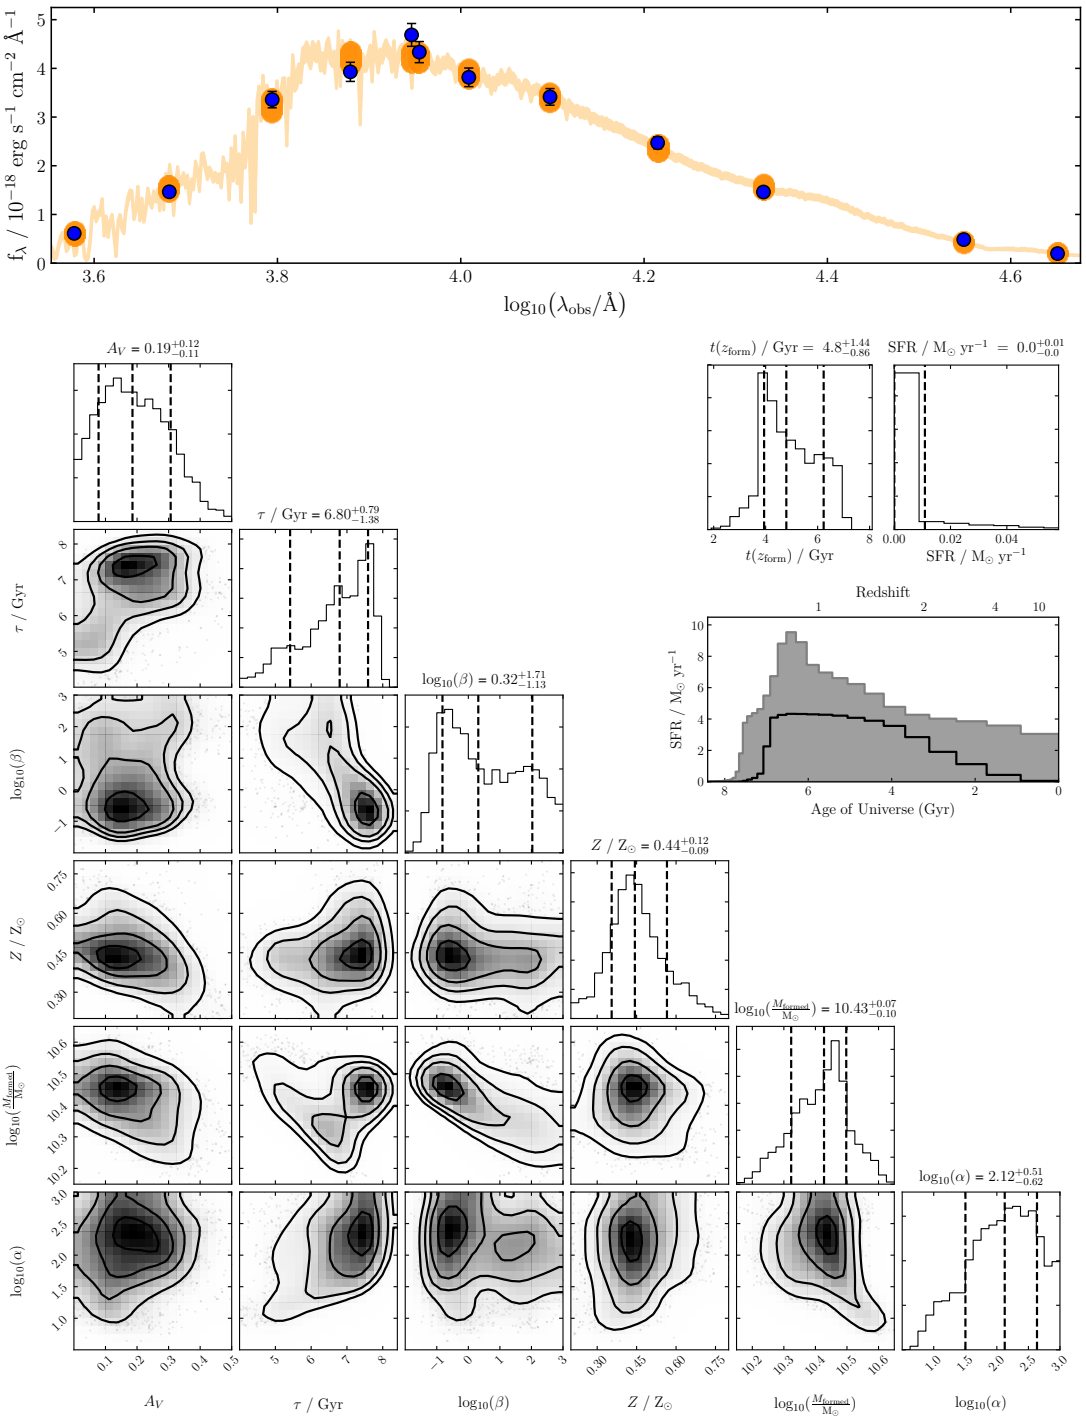

Figure 2.4 *Example of* Bagpipes *output when fitting mock photometry from the* Mufasa *simulation for an object at at* $z_{obs} = 0.5$ *using a double-power-law SFH. The mock photometry is shown on the top panel in blue (the bands are those of Table 2.1). The 16th to 84th percentile range for the posterior spectrum and photometry are shaded orange. A corner plot showing the posterior for fitted parameters is shown below, with posterior SFH information shown to the right.*

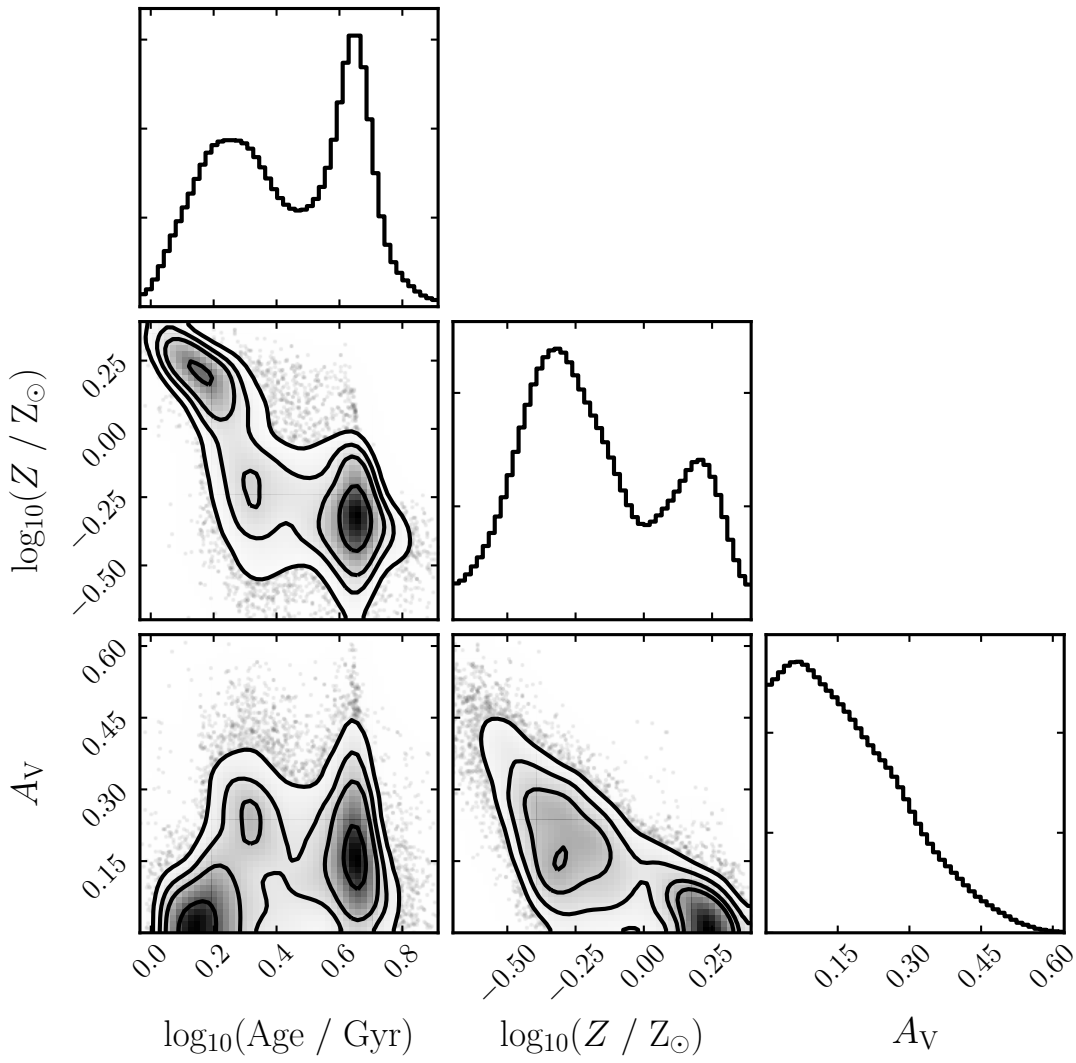

Figure 2.5 *The age-metallicity-dust degeneracy for the example* MUFASA *object shown in Fig. 2.4.*

data, $\mathcal{D}$ and a hypothesis, $\mathcal{H}$ with parameter vector $\mathbf{\Theta}$,

$$P(\mathbf{\Theta} \mid \mathcal{D},\ \mathcal{H}) \ = \ \frac{P(\mathcal{D} \mid \mathbf{\Theta},\ \mathcal{H})\ P(\mathbf{\Theta} \mid \mathcal{H})}{P(\mathcal{D} \mid \mathcal{H})}. \tag{2.8}$$

Here, $P(\mathbf{\Theta} \mid \mathcal{H})$ is the prior probability distribution for the model parameters: what is known before the new data are considered. $P(\mathcal{D} \mid \mathbf{\Theta}, \mathcal{H}) = \mathcal{L}(\mathcal{D} \mid \mathbf{\Theta},\ \mathcal{H})$ is the likelihood: the probability of obtaining the new data under the assumption of a certain set of parameter values. $P(\mathbf{\Theta} \mid \mathcal{D},\ \mathcal{H})$ is the posterior probability distribution: what is known once the new data have been considered. Finally, $P(\mathcal{D} \mid \mathcal{H})$ is the Bayesian evidence, or marginal likelihood: how good the hypothesis is overall at explaining the data.

For our purposes, the hypothesis consists of a certain model parameterisation within Bagpipes (e.g. type of dust model, number and type of SFH components, whether nebular emission is included etc). Once the model has been constructed, the user then defines prior probability distributions for its parameters (see Section 2.3.2 for details), and provides observational data, f_i with associated uncertainties σ_i, which will be used to constrain those parameters.

The observational data are incorporated into the calculation through the likelihood function, which is constructed assuming that uncertainties are Gaussian and independent. When fitting only photometric data, the log-likelihood function used by Bagpipes is given by (e.g. Hogg et al. 2010)

$$\ln(\mathcal{L}) \ = \ -0.5 \sum_i \ln\left(2\pi\sigma_i^2\right) \ - 0.5 \sum_i \frac{\left(f_i - f_i^{\mathcal{H}}(\mathbf{\Theta})\right)^2}{\sigma_i^2} \tag{2.9}$$

where $f_i^{\mathcal{H}}(\mathbf{\Theta})$ is the model prediction corresponding to the observed flux f_i, which is obtained by generating a Bagpipes model as described in Section 2.3.1.

Nested Sampling implementation

Bagpipes makes use of the nested sampling algorithm of Skilling (2006) to obtain the posterior distribution and evidence value given some model, prior distribution and observational data. Nested sampling is implemented in Bagpipes using MultiNest (Feroz & Hobson 2008; Feroz et al. 2009, 2013), accessed through the PyMultiNest interface (Buchner et al. 2014).

Nested sampling allows for efficient exploration of higher-dimensional, multimodal and highly degenerate parameter spaces. This is invaluable in many circumstances relevant to spectral fitting, notably when dealing solely with broad-band photometric observations. In this case, the age-metallicity-dust degeneracy often leads to, at best, poorly constrained and highly degenerate parameter estimates, or at worst several widely spaced local minima in parameter space (e.g. McLure et al. 2011), any of which can trap traditional numerical functional minimisation routines or MCMC methods.

Once the posterior distribution is returned by MultiNest, it can be post-processed by Bagpipes to obtain other posterior information of interest, such as the posterior spectrum and SFH. A variety of visualisations, such as corner plots (e.g. Foreman-Mackey 2016) can also be constructed. An example corner plot and posterior spectral plot from a fit to mock photometry for one of the Mufasa objects introduced in Section 2.4 is shown in Fig. 2.4. The age-metallicity-dust degeneracy is demonstrated in Fig. 2.5, which shows the correlations between these parameters for this object.

In the future, I intend to diversify the fitting options available in Bagpipes, giving the user access to more advanced versions of the nested sampling algorithm. This is necessary for faster and more reliable sampling of parameter spaces with higher dimensionalities, such as those encountered when fitting spectroscopic data. In particular I intend to include the option to implement dynamic nested sampling through the Dynesty package (Speagle 2019).

Specifying prior distributions on model parameters

Any parameter that can be specified in the construction of a Bagpipes model, as described in Section 2.3.1, can be fitted using the code. When fitting a parameter, a prior probability distribution must be specified, consisting of an upper and lower limit on the parameter value, and a functional form for the prior probability density between these limits. In certain cases, these probability density functions will include hyperparameters that must also be specified. Bagpipes includes a separate "priors" module, which allows the user to specify their own prior distributions, or pick from a number of options that have already been implemented.

2.4 Testing star-formation-history models with MUFASA

Bagpipes is a powerful tool for studying galaxy evolution, however, before applying it to real data it is important to test whether models defined within it produce unbiased estimates of the parameters of interest. For example, it is often possible to obtain a good fit to photometric observations of a galaxy by modelling the whole SFH as a single burst of star formation (SSP model). This is likely to produce a reasonably accurate and unbiased estimate of, for example, the redshift of the galaxy. However, if this method were used to estimate the average time at which stars in the galaxy formed (see Section 2.4.2), it would clearly be biased towards younger ages, as the luminosities of galaxies are dominated by the youngest stars they contain. This is an example of model misspecification, where the chosen functional form for a model, or aspect of a model (in this case the SFH), is adequate for obtaining a fit to the data, however fails to accurately represent the underlying data-generating process, leading to biases in derived parameters.

Recently, Leja et al. (2017) performed tests to validate the Prospector inference framework by comparing physical properties derived from fitting only broad-band photometry against well-calibrated indicators derived from aperture-matched spectroscopy for a sample of local galaxies. They show that their model is capable of unbiased estimation of the SFR, dust-attenuation properties and stellar metallicities of their galaxies.

In this thesis, however, the focus is on the SFH properties of galaxies. This presents an additional challenge, as no independent, well-calibrated methods exist to obtain, for example, the average time at which the stars in a galaxy formed. Instead it is necessary look to simulations of galaxy formation to provide realistic SFHs that can be converted into mock photometric observations. These can then be fitted with Bagpipes using a variety of different model parameterisations with different priors, in order to find a model that accurately reproduces the properties of the input SFHs.

In this section I use the Mufasa suite of cosmological hydrodynamic simulations (see Section 2.4.3) to construct a catalogue of mock observations for simulated quenched galaxies with realistic SFHs, chemical-enrichment histories and dust properties. I tailor the mock observations to match, as closely as possible, the

UltraVISTA data introduced in Section 2.2. I then test the ability of a variety of models within Bagpipes to reproduce the SFH properties of quenched Mufasa galaxies.

I begin in Section 2.4.1 by discussing the advantages and disadvantages of different commonly used representations of the SFHs of galaxies. Then, in Section 2.4.2, I introduce a set of parameters that can be used to compare how well different SFH models reproduce the important features of the input Mufasa SFHs. I then give details of the Mufasa simulation and the process of generating mock observations in Section 2.4.3. Finally, in Section 2.4.4 I fit these mock observations using Bagpipes, and compare the abilities of the exponentially declining and double-power-law SFH parameterisations to recover the SFH properties of Mufasa quenched galaxies. I demonstrate that, by using a double-power-law SFH model with the correct priors (see Table 2.3), unbiased estimates of the underlying properties of the Mufasa SFHs can be recovered. In contrast, I show that the exponentially declining model returns small but significant biases in these parameters. Based on these results, I proceed to use the double-power-law SFH model to fit the UltraVISTA dataset in Section 2.5.

2.4.1 Common SFH parameterisations

The traditional approach to SED fitting has been to use a simple functional form to parameterise the SFH, most commonly the exponentially declining function

$$\mathrm{SFR}(t) \propto \begin{cases} e^{-\frac{t-T_0}{\tau}} & t > T_0 \\ 0 & t < T_0 \end{cases} \tag{2.10}$$

where T_0 and τ are parameters to be fitted. The main advantage of this rigid parameterisation is the speed of fitting, however Maraston et al. (2010) and Reddy et al. (2012) have shown by that it becomes less effective at higher redshifts, and Wuyts et al. (2011) and Pforr et al. (2012) have shown that the biases on the estimated SFH parameters (e.g. the ongoing SFR) are highly dependent on the permitted ranges of (alternatively, the priors on) model parameters.

It is possible to define more flexible parameterisations with the same number of parameters, such as the lognormal form of Gladders et al. (2013), which has been shown by Diemer et al. (2017) to produce good fits to SFHs from the cosmological

simulation Illustris (Vogelsberger et al. 2014). This form is clearly more physical, as star formation is required to increase smoothly from zero at the beginning of the Universe, rather than jumping from zero to its maximum value several billion years later. However, when using the lognormal form, star formation must always decline more slowly than it rises, meaning that this parameterisation struggles to model rapidly quenching systems.

Attempts have also been made to improve the flexibility of model SFHs by using more complex parameterisations, at the cost of increased computational expense. One option is the double-power-law form, which has previously been used to fit the evolution of the cosmic SFRD (e.g. Behroozi et al. 2013b),

$$\mathrm{SFR}(t) \propto \left[\left(\frac{t}{\tau}\right)^{\alpha} + \left(\frac{t}{\tau}\right)^{-\beta}\right]^{-1} \tag{2.11}$$

where α and β are the falling and rising slopes respectively, and τ is related to the time at which star formation peaks. The major advantage of the double-power-law SFH model is the decoupling of the rising and falling slopes of the SFH, allowing, in contrast to the lognormal form, a slow rise and rapid cutoff to star formation, as demonstrated in Fig. 2.1. The double-power-law form is also discussed by Diemer et al. (2017) as a possible improvement on their lognormal parameterisation. They confirm that the double-power-law parameterisation produces better agreement with Illustris, particularly for rapidly quenching galaxies.

I also note that progress has been made in SED fitting using libraries of simulated SFHs from cosmological simulations (e.g. Pacifici et al. 2012), and using non-parametric SFHs (see Section 2.3.1). However these approaches are significantly more computationally intensive due to the inclusion of many free parameters, and are hence best suited to smaller sample sizes, rather than the large photometric catalogue I consider in this chapter.

2.4.2 Comparisons between SFH parameterisations

When attempting to compare results obtained using different SFH parameterisations, or to compare fitted SFHs with the input SFHs from Mufasa, it is necessary to define common parameters that can be derived from any SFH. One such parameter that is readily available is the stellar mass at the redshift of

observation, $M_*(z_{\rm obs})$, or simply M_*.

To answer Question (i), posed in Section 2.1, we require a method of quantifying when this mass was assembled. One common approach (e.g. Thomas et al. 2017) is to calculate the mass-weighted age, or in this case the mass-weighted time (measured forwards from the beginning of the Universe; see the note at the end of Section 2.1) of a galaxy, $t_{\rm MW}$, given by

$$t_{\rm MW} = \frac{\int_0^{t_{\rm obs}} t\ {\rm SFR}(t)\ {\rm d}t}{\int_0^{t_{\rm obs}} {\rm SFR}(t)\ {\rm d}t} = t(z_{\rm form}) \tag{2.12}$$

where $t_{\rm obs} = t(z_{\rm obs})$. I then define the redshift of formation, $z_{\rm form}$, based on this mass-weighted time by setting $t_{\rm MW} = t(z_{\rm form})$.

In order to answer Question (ii), we require knowledge of the redshifts at which galaxies quenched, $z_{\rm quench}$. In order to define this we first require a definition of a quenched galaxy. This will also be necessary for selecting quenched galaxies from the UltraVISTA catalogue. Two common methods exist for selecting samples of quenched galaxies, firstly based on their rest-frame UVJ colours (see Section 1.5.4), and secondly using a cut in specific SFR (sSFR; SFR divided by M_*), often evolving with observed redshift (e.g. Pacifici et al. 2016).

Relating the SFH of a galaxy to its rest-frame UVJ colours at earlier times requires several assumptions, such as the time-evolution of dust attenuation. This makes generalising UVJ selection to the past SFHs of galaxies extremely challenging. I therefore use a method similar to a specific SFR cut, by defining the dimensionless normalised SFR, which is the current SFR as a fraction of the time-averaged SFR over the whole SFH, given by

$$\frac{{\rm SFR}(t)}{<{\rm SFR}>} = {\rm SFR}(t)\frac{t}{\int_0^t {\rm SFR}(t')\ {\rm d}t'} = \frac{t\ {\rm SFR}(t)}{M_{\rm formed}} \tag{2.13}$$

where $\int_0^t {\rm SFR}(t')\ {\rm d}t' = M_{\rm formed}$ is the total mass of stars formed by the galaxy up to the time t, which is closely related to the living stellar mass.

I will define quenched galaxies as those that have normalised star-formation rate < 0.1, meaning the SFR at the time of observation must be less than 10 per cent of the average SFR across the history of the galaxy. This is effectively similar to an evolving cut in specific SFR, however it is not biased by the epoch at which

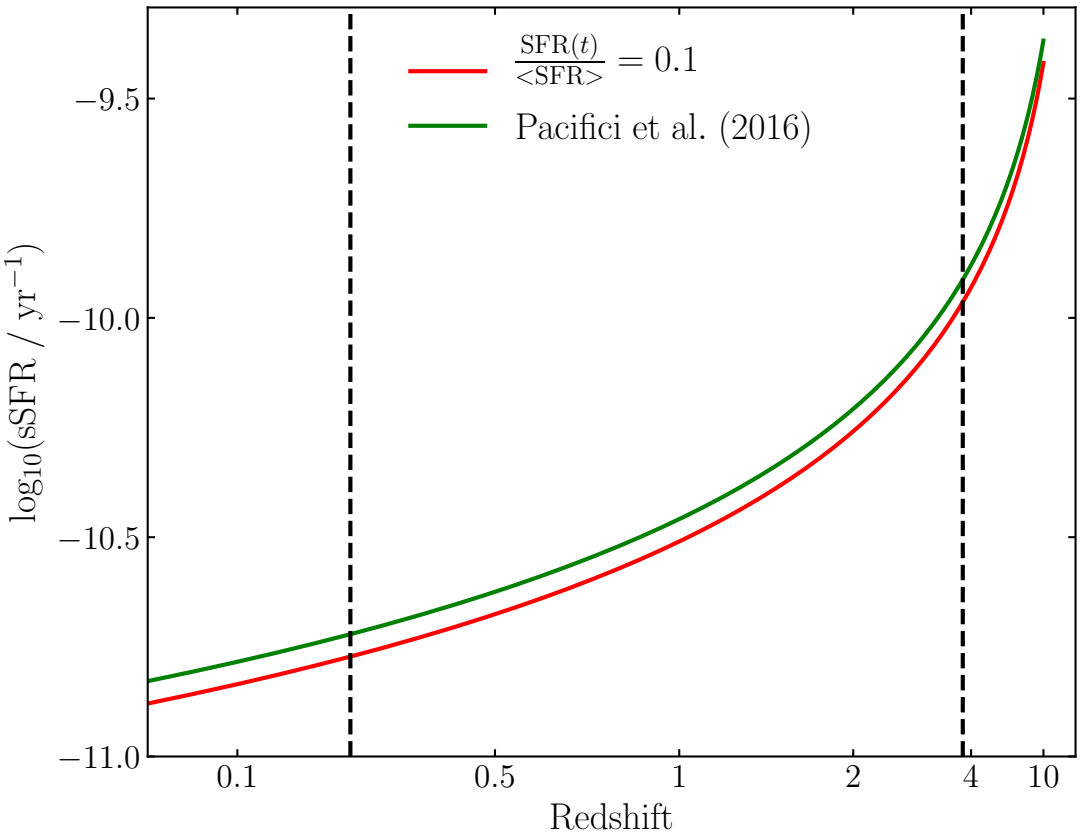

Figure 2.6 *Selection criteria for quiescent galaxies plotted as cuts in specific SFR as a function of redshift. My cut in normalised SFR at 0.1 (see Section 2.4.2) is shown in red, assuming that the living stellar mass, M_* is 0.25 dex less than the total stellar mass formed (in reality this is dependent on the shape of the SFH). The selection criterion employed by Pacifici et al. (2016) is shown in green. These can be seen to be very similar, and both produce good agreement with the UVJ selection criteria of Williams et al. (2009). The black dashed lines denote the edges of the redshift range spanned by the UltraVISTA sample.*

the galaxy's stellar mass was assembled, as it depends on the total stellar mass formed rather than the mass in stars at the time of observation. The value of 0.1 was chosen to preserve good agreement with UVJ selection (as demonstrated in Section 2.5). I show in Fig. 2.6 how the cut in normalised SFR corresponds to a cut in sSFR, confirming that this selection criterion is very similar to that of Pacifici et al. (2016), which they also find to be in good agreement with UVJ selection. I hence define $z_{\rm quench}$ as the redshift at which the normalised star-formation rate falls below 0.1. In the case of the MUFASA SFHs, which can be bursty and therefore fall below this threshold several times, I select the latest time at which this happens.

Finally, for the analysis presented in Section 2.6 it will be useful to define the quenching timescale, $\Delta t_{\rm quench}$ for quenched galaxies. I define this as being the duration between the time of formation $t(z_{\rm form})$, or mass-weighed time, and the time of quenching $t(z_{\rm quench})$, such that

$$\Delta t_{\rm quench} \;=\; t(z_{\rm quench}) - t(z_{\rm form}). \tag{2.14}$$

This definition of $\Delta t_{\rm quench}$ is designed to trace the timescale over which star-formation declines from its peak to a normalised SFR of 0.1, at which point I define the galaxy as quenched.

Figure 2.7 shows how successful $\Delta t_{\rm quench}$ is in capturing this for a range of different SFH shapes. As can be seen, a rapidly quenched galaxy that has an extended rising wing will have a higher $\Delta t_{\rm quench}$ than one that quenches equally quickly, but also forms very rapidly. Thus $\Delta t_{\rm quench}$ traces both the speed of quenching as intended, but also has a secondary dependence on how extended the SFH is prior to the onset of quenching.

As the dynamical timescales of galaxies evolve with redshift, a parameter of significant interest is $\Delta t_{\rm quench}$ as a fraction of the age of the Universe when quenching occurs, $t(z_{\rm quench})$. This allows us to directly compare quenching timescales for galaxies at different observed redshifts. I define this normalised quenching timescale as $\tau_{\rm quench}$, where

$$\tau_{\rm quench} \;=\; \frac{\Delta t_{\rm quench}}{t(z_{\rm quench})}. \tag{2.15}$$

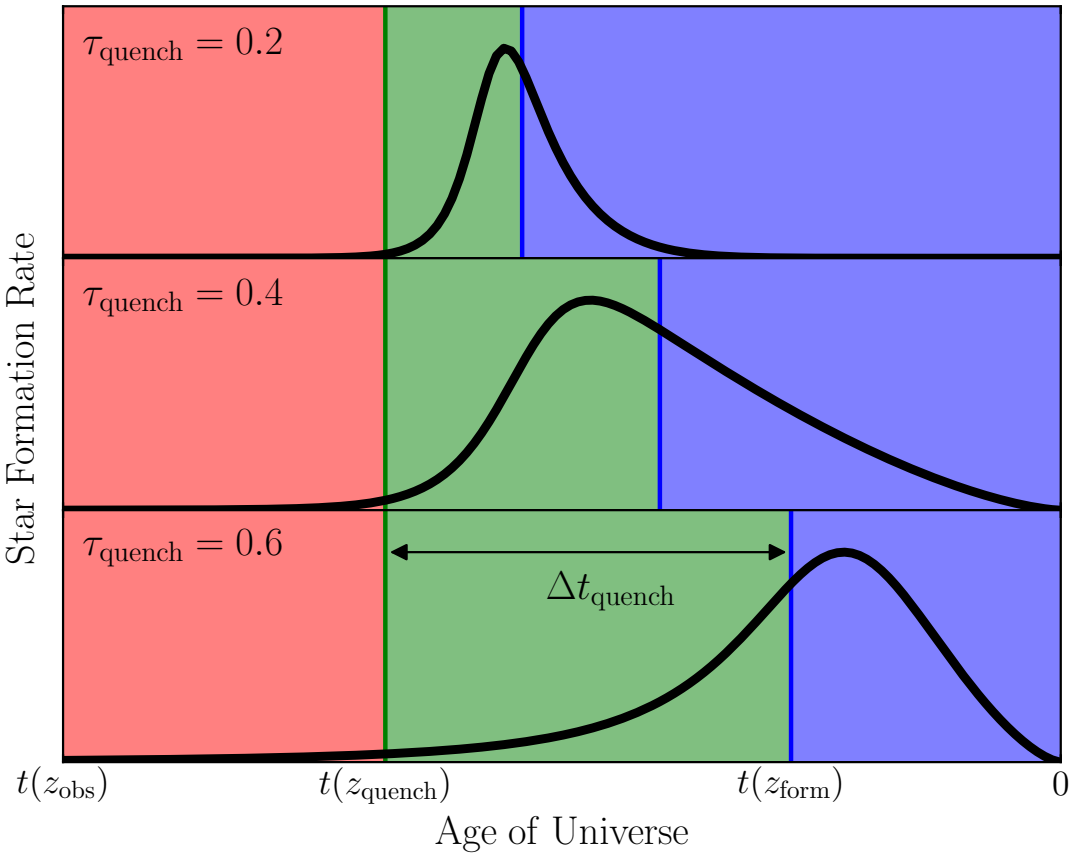

Figure 2.7 *Pictorial representation of the scheme for describing and comparing galaxy SFHs introduced in Section 2.4.2. The three example double-power-law SFHs shown all quench at the same time but have different shapes, leading to different values of* $\tau_{\rm quench}$*, which can be seen to trace both the speed of quenching, and, to a lesser extent, how quickly the SFH rises. The quenching timescale can be thought of as the time between peak star formation and the time of quenching.*

2.4.3 Generating mock observations of Mufasa galaxies

Mufasa (Davé et al. 2016) is a suite of cosmological hydrodynamic simulations that correctly reproduces many of the observed properties of the population of quenched galaxies at $z = 0$, including the slope of the red sequence and downsizing trend, as well as correct number densities for high-mass ($M_* \gtrsim 10^{10.7}$ M$_\odot$) quenched galaxies out to $z \sim 2$ (Davé et al. 2017). This makes Mufasa an ideal resource for testing the abilities of different models to recover realistic SFHs for massive quenched galaxies across a wide range of redshifts.

I will use the post-processed data-products from Mufasa described by Davé et al. (2017), which consist of a list of star particles for each galaxy with individual masses, ages, metallicities and A_V values. I use snapshots of the simulation taken between $z = 0.5$ and 2.5 at intervals of $\Delta z = 0.5$. To begin with, I select all galaxies from each of these snapshots with $M_* > 10^{10}$ M$_\odot$ in order to match the observed sample from UltraVISTA (see Section 2.5).

For each Mufasa galaxy I generate a spectrum using Bagpipes by modelling each star particle as a burst of star formation. I use the BC03 models and the Calzetti et al. (2000) attenuation curve, which has been shown to be a robust average attenuation curve for galaxies across a wide range of redshifts (e.g. Cullen et al. 2017, 2018; McLure et al. 2018b). I also fix $\log_{10}(U) = -3$ and $a_{\rm BC}$ to 10 Myr for the nebular-emission model, as these properties have only marginal impact on broad-band photometry for quiescent galaxies. For each galaxy I then calculate observed magnitudes in the twelve filters of the UltraVISTA catalogue described in Section 2.2.

I also compile the simulated SFHs of these objects in 100 Myr bins and calculate the normalised SFR (Equation 2.13) at the time of observation using the SFR in the most recent 100 Myr bin as the current ongoing SFR. I then select objects with normalised SFR < 0.1 as the quenched sample, as described in Section 2.4.2. For the SFHs of the quenched sample I then calculate the three indicators described in Section 2.4.2, M_*, $z_{\rm quench}$ and $z_{\rm form}$.

Once the sample had been selected, Gaussian noise was added to the mock photometry at the levels listed for the UltraVISTA Wide strips in Table 2.1. An error floor was added such that no photometry point corresponded to a detection of greater than 20σ, or, in the case of the IRAC channels, 10σ. Finally a scatter was added to the catalogued observed redshifts of the Mufasa galaxies at the

same level ($dz = 0.02$) as the photometric redshifts calculated for the UltraVISTA catalogue. At the end of this process I had built a catalogue of mock observations for 677 simulated quenched galaxies, tailored to match as closely as possible the sample I intended to select from the UltraVISTA data of Section 2.2.

2.4.4 Recovering star-formation histories for Mufasa quenched galaxies

In this section I present the results of fitting the mock observations produced in Section 2.4.3 with different SFH models. I primarily consider SFH recovery, and leave discussion of other physical parameters to future work. I evaluate the ability of different parameterisations to accurately recover the Mufasa SFHs using the three indicators introduced in Section 2.4.2, M_*, $z_{\rm quench}$ and $z_{\rm form}$.

For all of the fitted models I assume Calzetti dust with fixed $\epsilon = 3$, and fit A_V with a uniform prior between 0 and 4 magnitudes. For the nebular model I fix $a_{\rm BC} = 10$ Myr and $\log_{10}(U) = -3$. In principle all of these fixed parameters could be fitted with Bagpipes, however due to their marginal impact on observed photometry, particularly for quenched galaxies, and the computational expense of fitting large numbers of parameters they were instead fixed to representative values, consistent with results from the literature (see Section 2.3.1).

I first build a baseline model of the kind commonly used in the literature (e.g. Mortlock et al. 2017), by parameterising the SFH with a single exponentially declining component (Equation 2.10) and applying uniform priors to all SFH parameters. I set the limits of the priors to zero and 50 Myr less than $t(z_{\rm obs})$ for T_0, between 300 Myr and 10 Gyr for τ (following Wuyts et al. 2011 in both cases), between 1 and 10^{13} $\rm M_\odot$ in stellar mass formed, and between 0.2 and 2.5 $\rm Z_\odot$ in metallicity (Solar metallicity is assumed to take the value $\rm Z_\odot = 0.02$).

Once posterior distributions had been obtained for model parameters, I generated posterior SFHs and used these to calculate the posterior distributions of M_*, $z_{\rm quench}$ and $z_{\rm form}$ for each object. A representative comparison between Mufasa input and Bagpipes output from fitting the exponentially declining model to mock photometry is shown in the left-hand panels of Fig. 2.8. At the top the SFH is shown, in the middle the mass-assembly history, and at the bottom the history of the normalised SFR (see Equation 2.13). It can be seen that, for this object, this model over-predicts $M_{\rm formed}$, and hence M_*, and over-predicts both

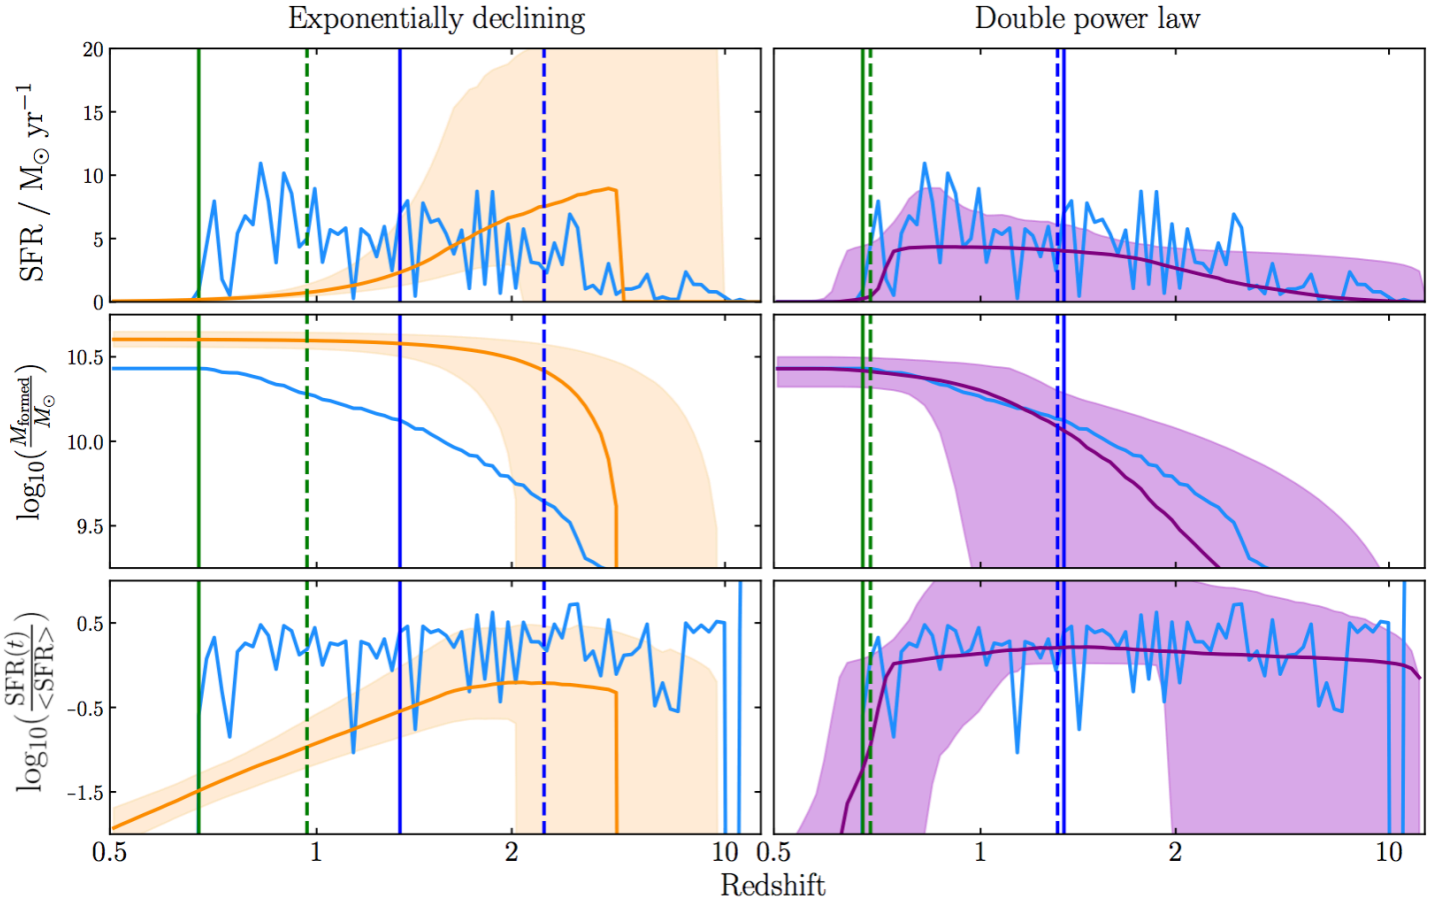

Figure 2.8 *SFH (top), mass-assembly history (middle) and normalised SFR history (bottom) for the same example object from* MUFASA *as shown in Fig. 2.4. The input histories from* MUFASA *are shown in blue with the posteriors shown in orange (left) and purple (right) for the exponentially declining and double-power-law SFH models of Section 2.4.4 respectively. The shaded regions show the 16$^{\rm th}$ – 84$^{\rm th}$ percentiles of the posterior. The true redshifts of quenching and of formation (as defined in Section 2.4.2) are overplotted as green and blue solid vertical lines respectively, with corresponding posterior median fitted values shown as dashed lines.*

Table 2.3 *Fixed and fitted parameters with their associated priors for the double-power-law SFH model used to fit the UltraVISTA catalogue. Priors listed as logarithmic are uniform in log-space. Observed redshift values were taken from the Mortlock et al. (2017) catalogue (see Section 2.5). Priors placed on parameters that are well constrained by the data (e.g. stellar mass) have little to no effect on the results obtained.*

Free Parameter	Prior	limits	Fixed Parameter	Value
A_V	Uniform	(0, 4)	$a_{\rm BC}$ / Myr	10
$\log_{10}(M_{\rm formed}$ / $\rm M_\odot)$	Uniform	(1, 13)	$\log_{10}(U)$	-3
Z / $Z_\odot$	Uniform	(0.2, 2.5)	ϵ	3
τ / Gyr	Uniform	$(0, t(z_{\rm obs}))$	$z_{\rm obs}$	$z_{\rm M17}$
α	Logarithmic	$(10^{-2}, 10^3)$		
β	Logarithmic	$(10^{-2}, 10^3)$		

$z_{\rm quench}$ and $z_{\rm form}$. There are several obvious discrepancies in the histories when viewed as a whole, with the overall shape of the SFH being poorly reproduced by the model, the mass assembly occurring much earlier than in MUFASA.

Fig. 2.9 shows a comparison between input values of of M_*, $z_{\rm quench}$ and $z_{\rm form}$ from MUFASA and the posterior median values obtained by fitting the exponentially declining model to the whole of the MUFASA quenched sample. It can be seen that biases are present in the estimates of the SFH parameters. This is most notable for the estimated stellar masses, which are over-predicted on average by 0.06 dex ($\sim$ 15 per cent), with the posterior median stellar masses of $\sim$ 80 per cent of objects being overestimates. Whilst this is a relatively small bias compared to other uncertainties that exist in the determination of stellar masses (e.g. Mobasher et al. 2015), it is nevertheless significant, and interesting given the prevalence of this method for estimating stellar masses in the literature.

Both $t(z_{\rm quench})$ and $t(z_{\rm form})$ are underestimated on average by $\simeq$ 0.4 Gyr. This is perhaps surprisingly good agreement, given that exponentially declining SFHs were not designed to reproduce these properties. However, as can be seen in the middle-left panel of Fig. 2.9, there is considerable systematic variation in the degree of underestimation of $t(z_{\rm quench})$ with the time between $t(z_{\rm obs})$ and $t(z_{\rm quench})$, most notably at $z_{\rm obs} = 0.5$, with the quenching times of the earliest-quenching objects actually being, on average, overestimated. This is a problem as it results in inconsistencies between quenching properties derived for different sub-samples of objects.

Motivated by these observations, I then tested other SFH parameterisations and combinations of priors in order to attempt to obtain better agreement with the

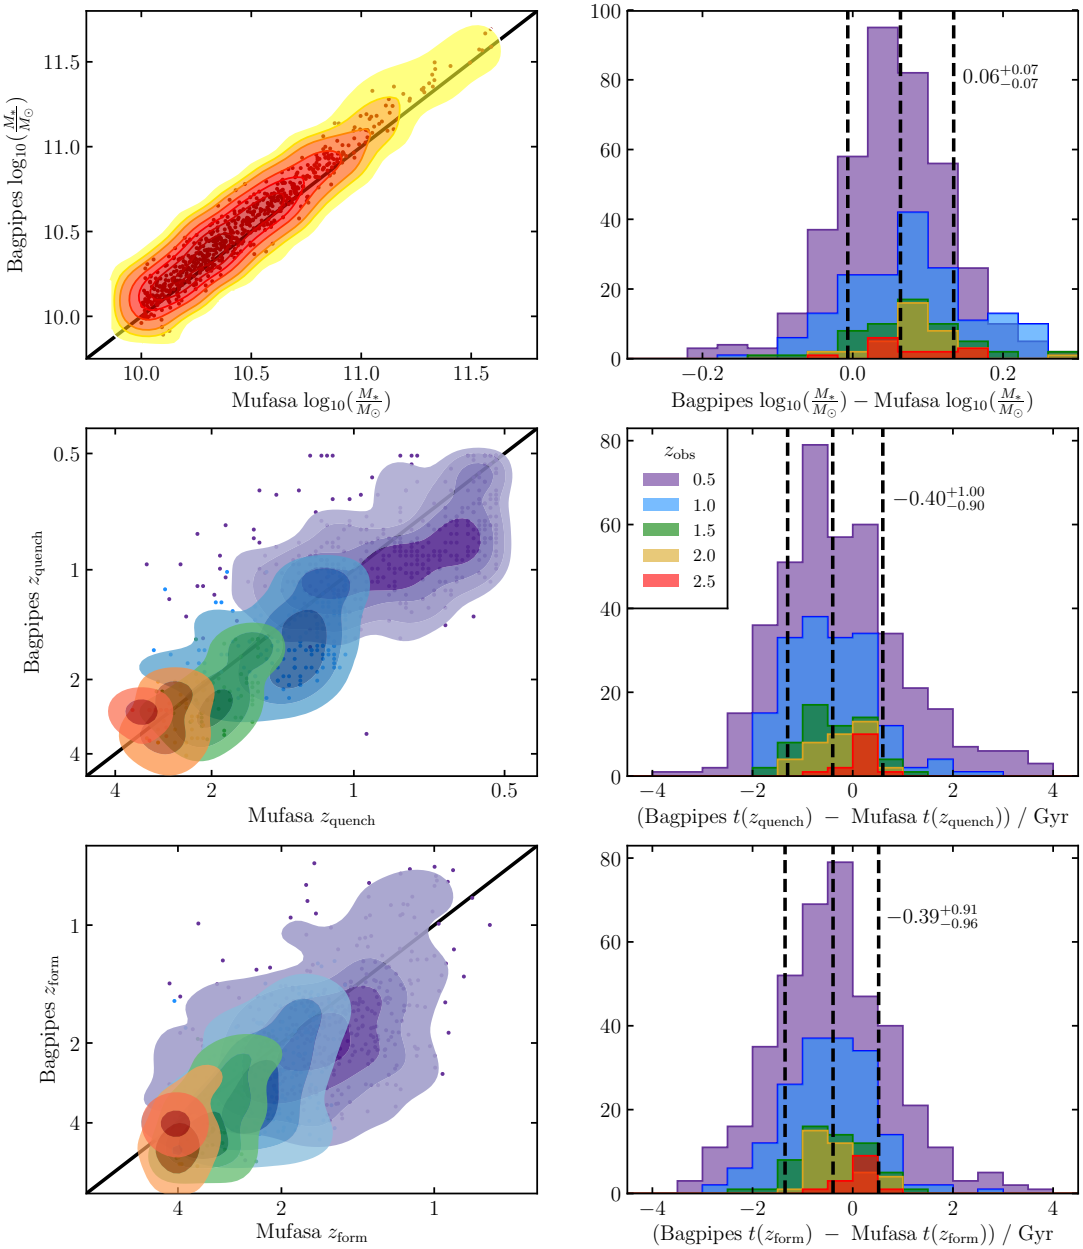

Figure 2.9 *Comparison between SFH properties of* Mufasa *galaxies and posterior median values derived by fitting their mock photometric observations with* Bagpipes *using the exponentially declining SFH parameterisation described in Section 2.4.4. The top-left panel shows stellar mass, middle-left panel shows redshifts of quenching and bottom-left panel shows redshifts of formation for the whole* Mufasa *quenched sample of 677 objects. The right-hand panels show histograms of the offset between fitted and true values of log-stellar mass (top), time of quenching (middle) and time of formation (bottom). The dashed lines on the right-hand panels indicate the 16th, 50th and 84th percentiles of the combined distribution for all observed redshifts. The stellar masses are overestimated by* ~ 15 *per cent on average, and both* $t(z_{\rm quench})$ *and* $t(z_{\rm form})$ *are underestimated on average by* $\simeq 0.4$ *Gyr.*

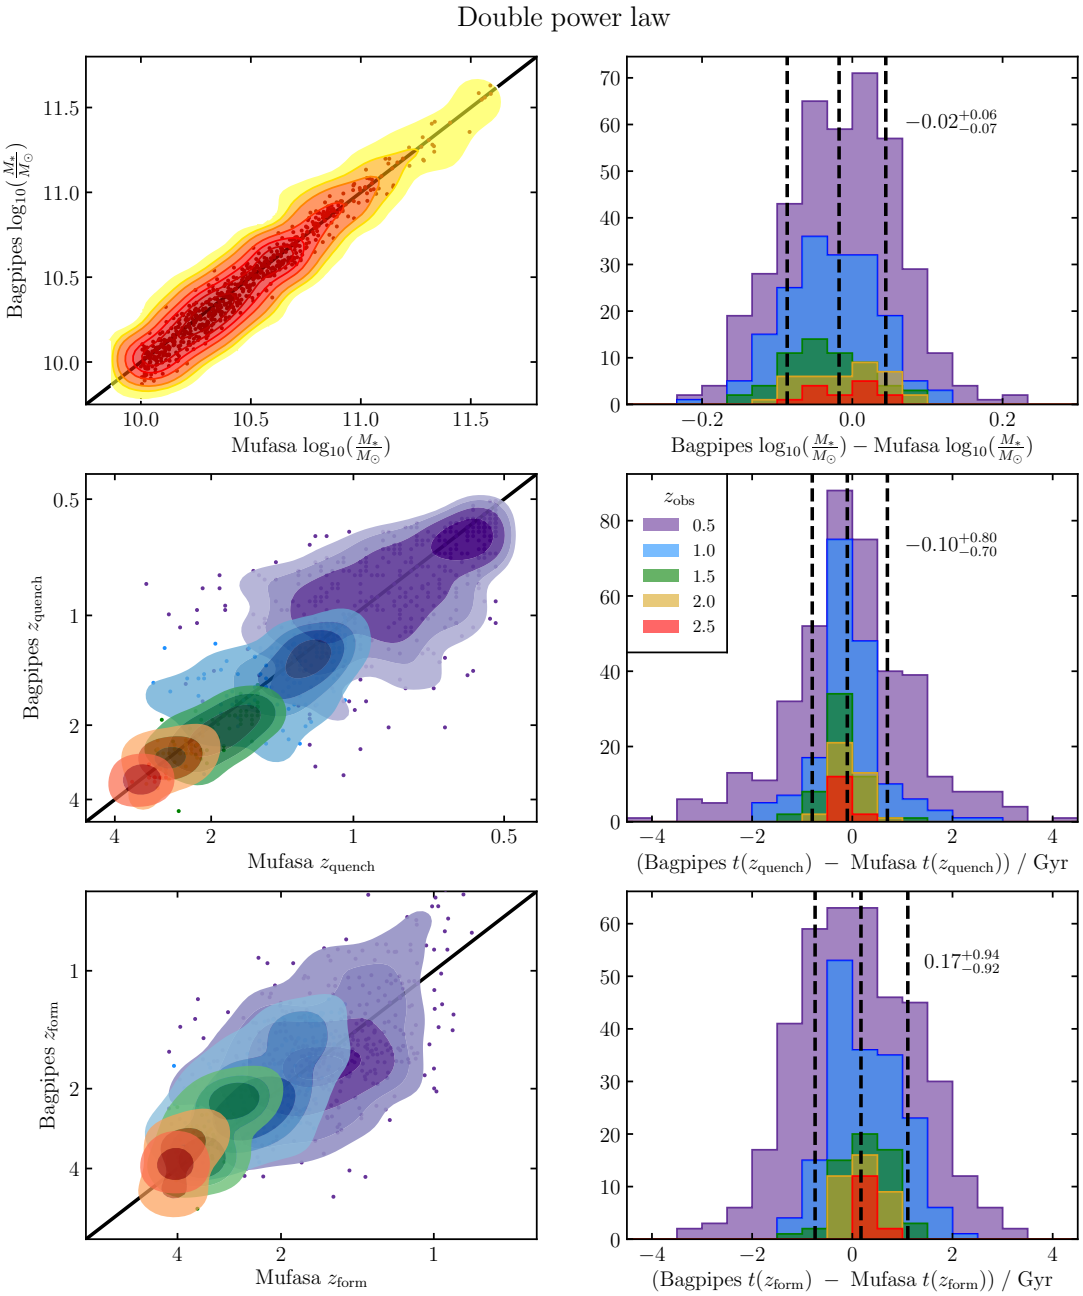

Figure 2.10 *Comparison between SFH properties of* Mufasa *galaxies and posterior median values derived by fitting their mock photometric observations with* Bagpipes *using the double-power-law SFH parameterisation described in Table 2.3. Plot details are as in Fig. 2.9. Considerably better agreement can be seen for all three parameters at each observed redshift when compared to the exponentially declining SFH results displayed in Fig. 2.9.*

true values of M_*, $z_{\rm quench}$ and $z_{\rm form}$ from MUFASA. I first investigated the effects of fitting the exponentially declining SFH parameterisation with different combinations of priors, but found no other combination that lead to significantly better agreement with MUFASA.

I then considered the double-power-law form of Equation 2.11. To incorporate the prior expectation, based on the hierarchical growth of galaxies observed in cosmological simulations, that most SFHs are extended almost all the way back to the beginning of cosmic time, I imposed a logarithmic prior on the rising slope, β. In the absence of constraining data, this produces an extended, gradual rise in the SFH up to the point at which the imprint of more recent star formation becomes obvious in the shape of the observed SED. I also applied a logarithmic prior to the falling slope, α, however this can be easily overcome by the strong imprint that recent, rapidly quenched star formation leaves on the observed SED.

A list of the parameters and priors for the double-power-law model are presented in Table 2.3. Fig. 2.10 shows the same comparison between input and fitted parameters for the double-power-law model as was shown for the exponentially declining model in Fig. 2.9. It can be seen that considerably better agreement is achieved for all three of the SFH parameters, with stellar masses now offset by only 0.02 dex. The most significant improvement is in the estimation of $t(z_{\rm quench})$, where the bias on the median has been cut to 100 Myr. More significantly, the contours on the middle-left panel are now well aligned with the diagonal, indicating no change to the bias in the estimates of $z_{\rm quench}$ with increasing time interval between $t(z_{\rm obs})$ and $t(z_{\rm quench})$. Additionally, the observed scatter in the estimates of $t(z_{\rm quench})$ have been cut by $\simeq 20$ per cent.

Fig. 2.4 shows example spectral and corner plots for a fit to one of the MUFASA objects using this model. This is the same object as shown in Fig. 2.8. The fitted SFH, mass assembly history and normalised SFR history for the double-power-law model are shown on the right-hand panels of Fig. 2.8. These can be seen to be in considerably better agreement with the input histories from MUFASA than for the exponentially declining model, with close agreement being achieved with the input histories at all times.

The better agreement achieved with input MUFASA SFH parameters using the double-power-law model led me to take this model forward for use in fitting the UltraVISTA sample in Section 2.5. Under the assumption that the SFHs of real massive quenched galaxies look similar to those from MUFASA, I have shown that

this model will allow me to obtain negligibly biased estimates of M_*, $z_{\rm quench}$ and $z_{\rm form}$ for the galaxies in the UltraVISTA sample.

2.5 UltraVISTA fitting and sample selection

In this section I describe the process of selecting the sample of 9289 massive quenched galaxies from the UltraVISTA catalogue introduced in Section 2.2. I begin by fitting the double-power-law model within Bagpipes to every object in the catalogue, except for those that were identified as stars or potential AGN by Mortlock et al. (2017). Parameters and priors for this model are listed in Table 2.3. Fitting the Bagpipes model whilst allowing observed redshift to vary produces more scatter in the recovered photometric redshifts than Mortlock et al. (2017) obtained by taking the median result of five photometric-redshift codes, therefore it was decided to fix the observed redshifts to these median values.

I then select objects with posterior median stellar masses, $M_* > 10^{10}$ M$_\odot$ and calculate the posterior distributions of the normalised SFR (as defined in Section 2.4.2) for these objects. I then select the quenched sample to contain galaxies that have posterior median normalised SFR values less than 0.1. At this stage I also exclude poorly fitted objects with minimum reduced chi-squared values of greater than 3. These objects are almost exclusively the result of individual datapoints affected by large systematic errors. They comprise ~ 1 per cent of the sample and are not strongly clustered in observed redshift or stellar mass.

Fig. 2.11 shows the final sample of quenched galaxies on the UVJ plane, as well as star-forming galaxies that match all of the other selection criteria except the cut in normalised SFR. A UVJ selection of the kind proposed by Williams et al. (2009) is also plotted. The normalised SFR selection and UVJ selection can be seen to select very similar samples of objects at all redshifts, with typically $\gtrsim 90$ per cent agreement between the two methods. The normalised SFR method can be seen to be more subtle in removing star-forming galaxies close to the UVJ boundary.

I split the sample into six bins in observed redshift between $z = 0.25 - 3.75$, chosen to span similar intervals in cosmic time. I then split each of these into five bins in stellar mass. The first four mass bins have widths of 0.3 dex from $\log_{10}(M_*/{\rm M}_\odot) = 10 - 11.2$, whilst the highest mass bin includes all objects with

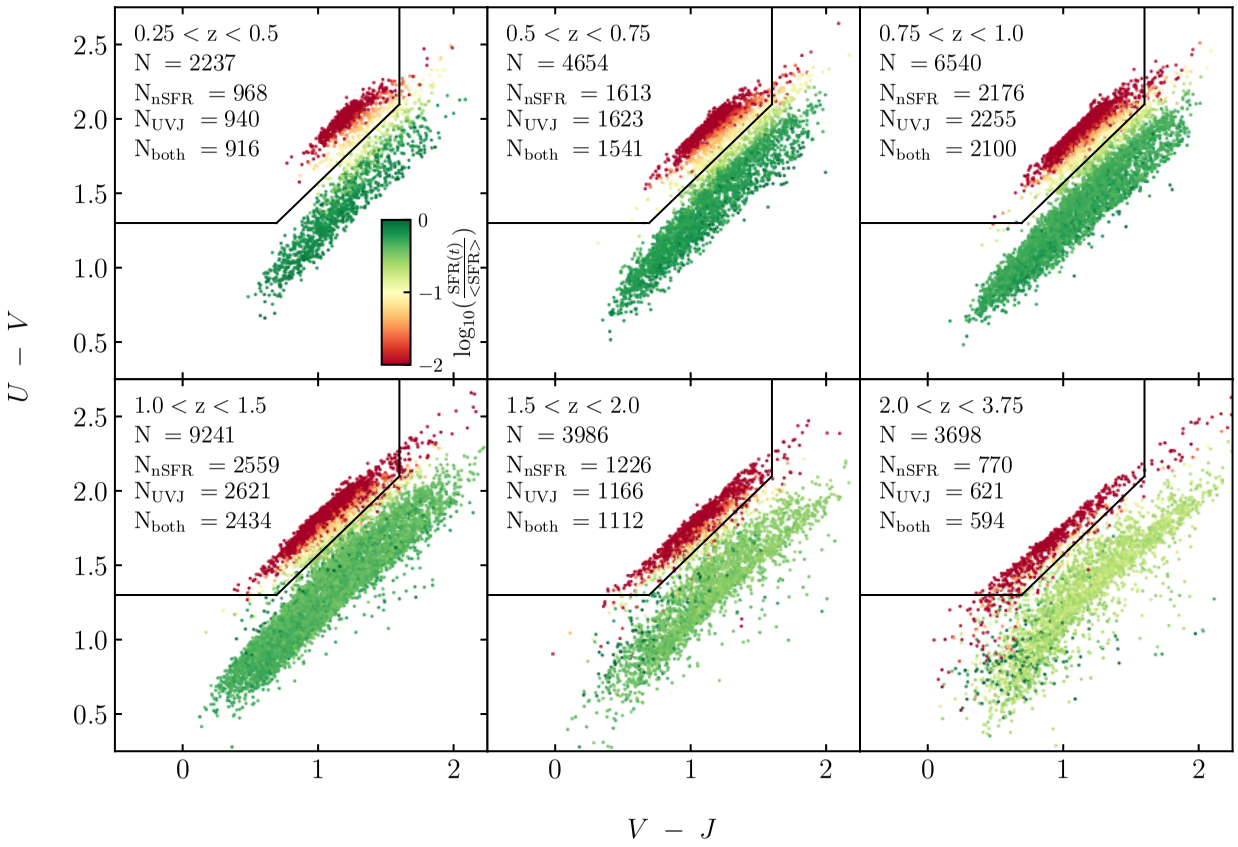

Figure 2.11 *UVJ selection diagram for the sample of galaxies from UltraVISTA described in Section 2.5. The population of star-forming galaxies that match all of the criteria except the cut in normalised SFR (nSFR; see Equation 2.13) are also shown. A strong trend in nSFR perpendicular to the star-forming sequence can be seen, as has been observed for sSFR (e.g. Fang et al. 2018). A UVJ selection of the kind proposed by Williams et al. (2009), with limits* $(U-V) > 1.3$, $(V-J) < 1.6$, $(U-V) > 0.88\times(V-J) + 0.69$ *is overplotted in black for reference, however this was not included in the selection criteria. The selection criterion for identifying quenched galaxies by* $\mathrm{nSFR} < 0.1$ *can be seen to be in good agreement with the UVJ selection shown in black at all observed redshifts.*

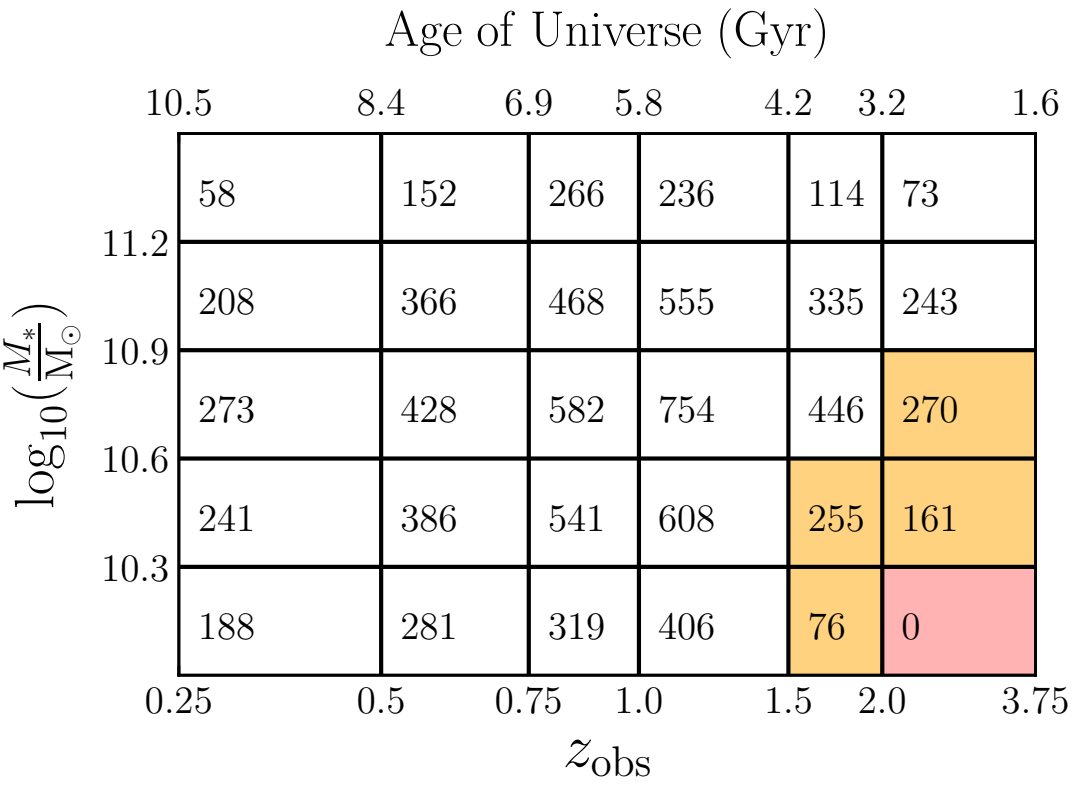

Figure 2.12 *Number of galaxies in each of the bins in stellar mass and observed redshift. The edges of the bins are shown in black. Bins shaded orange are those for which UltraVISTA wide is not mass complete, and hence objects from the wide area were removed from the sample. The bin shaded red falls below the mass-completeness limit of both the wide and deep regions, and hence was excluded entirely from the sample.*

$\log_{10}(M_*/\mathrm{M}_\odot) > 11.2$. Fig. 2.12 shows the distribution of the final UltraVISTA quenched sample across the bins in stellar mass and observed redshift.

The Mortlock et al. (2017) catalogue is mass complete down to $\log_{10}(M_*/\mathrm{M}_\odot) = 10$ over the UltraVISTA wide region at $z \lesssim 1.5$, and over the deep region at $z \lesssim 2$. I therefore exclude objects from the highest-redshift, lowest-mass bins if they fall below the mass-completeness limits of the UltraVISTA region they are drawn from. Bins are shaded orange in Fig. 2.12 for objects in the wide region having been excluded, or red for all objects having been excluded. These final criteria leave a sample of 9289 objects.

2.6 The star-formation histories of UltraVISTA quenched galaxies

In this section I present and discuss the results of my analysis of the SFHs of massive quenched galaxies from UltraVISTA. I analyse the SFHs using the scheme presented in Section 2.4.2, by calculating posterior SFHs for each object and using these to calculate posterior distributions for $t(z_{\mathrm{form}})$, $t(z_{\mathrm{quench}})$, $\Delta t_{\mathrm{quench}}$ and τ_{quench}. The goal of this analysis is to address Questions (i) and (ii) posed in Section 2.1. Hence, in Section 2.6.1, I consider Question (i): when did quenched galaxies form their stellar masses? Then, in Section 2.6.2, I consider Question (ii): how long did the process of quenching take?

As can be seen from Fig. 2.4, the SFHs of individual objects are relatively poorly constrained. At this point, the question of how best to present and draw conclusions from the richness of the posterior information I have obtained naturally arises. The optimal approach to analysis of this dataset would be to construct a Bayesian hierarchical model in which the values of z_{form} and z_{quench} for each object were assumed to be drawn from some overall sample distribution with stellar-mass and observed-redshift dependencies, which could then be constrained by the whole dataset. However the computational expense of implementing this (e.g. using Gibbs sampling) was judged to be prohibitive. Instead I have elected to represent the posterior distribution for each object by the median values of relevant parameters, and to take advantage the large sample size to make inferences about how these vary across the stellar-mass and observed-redshift bins plotted in Fig. 2.12.

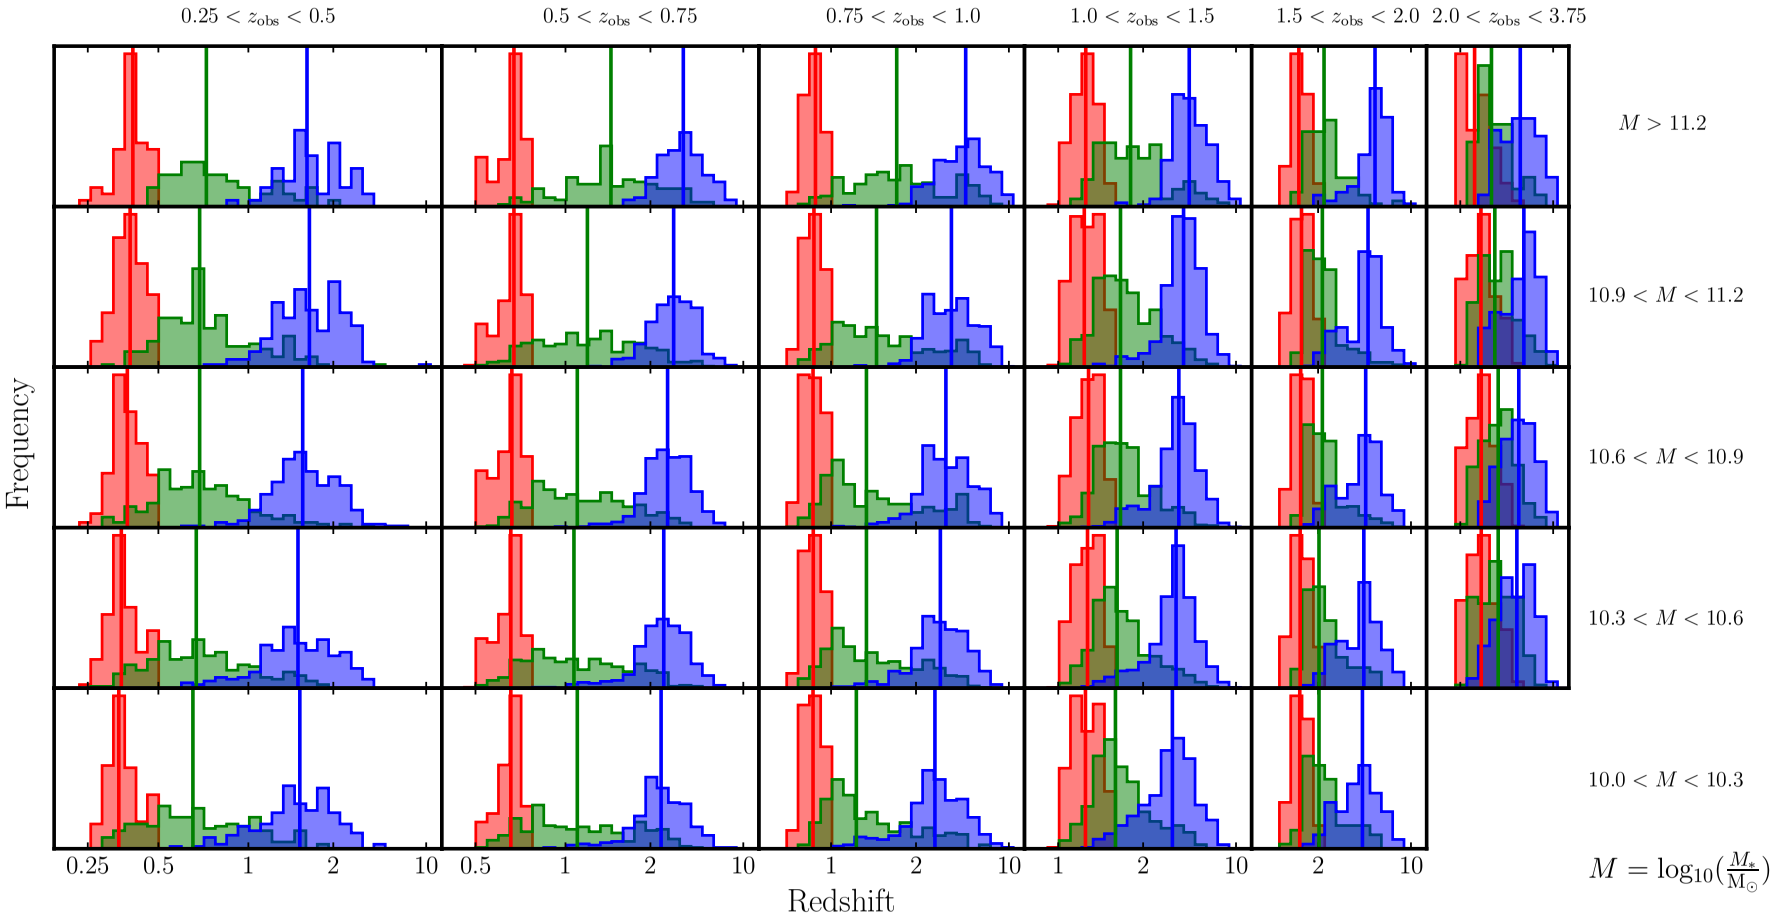

Figure 2.13 *Histograms of the redshifts of observation (red), posterior median redshifts of quenching (green) and posterior median redshifts of formation (blue) for the sample of UltraVISTA quenched galaxies, split into the stellar-mass and observed-redshift bins shown on Fig. 2.12. The medians of the distributions for each bin are denoted by solid vertical lines.*

2.6.1 When did quenched galaxies form their stellar masses?

Fig. 2.13 shows histograms of the posterior median values of $z_{\rm form}$ and $z_{\rm quench}$ for objects in the UltraVISTA sample, as well as the distributions of observed redshifts, split into the stellar-mass and observed-redshift bins shown in Fig. 2.12. The distributions of $z_{\rm form}$ and $z_{\rm quench}$ can be seen to be, on the whole, unimodal and roughly Gaussian in shape. It is therefore reasonable to represent these distributions by their median values, which are plotted as solid vertical lines in Fig. 2.13. The median values of $z_{\rm form}$ with their standard errors are plotted against median stellar mass in Fig. 2.14.

Two clear trends are visible in the distributions of $z_{\rm form}$ from Fig 2.13 and Fig. 2.14. Firstly there is the well-known downsizing trend, in which more massive galaxies at fixed observed redshift have higher median $z_{\rm form}$. This is discussed in Section 2.6.1. Secondly, there is a trend towards progressively lower median $z_{\rm form}$ for objects at lower observed redshifts. This is discussed in Section 2.6.1.

Downsizing: higher formation redshifts at higher masses

The downsizing trend observed, most clearly in Fig. 2.14, has also been observed in many other studies (see Section 2.1). Siudek et al. (2017) consider a sample of quenched galaxies at $0.4 < z < 1.0$ and show, in their fig. 11, redshifts of formation calculated from stacked spectra in bins of $\Delta z = 0.1$ across the same range of stellar masses I consider. Their redshifts of formation are considerably lower than those I show in Fig. 2.14 for corresponding redshift bins. However, as they note, their methodology of assuming the SFHs of their galaxies to be composed of only a single burst component means their redshifts of formation trace the time of the most recent episode of significant star formation. This means the formation redshifts they report are more analogous to my $z_{\rm quench}$ than $z_{\rm form}$. It can be seen from Fig. 2.13 that the median $z_{\rm quench}$ I derive for the $0.5 < z < 0.75$ and $0.75 < z < 1.0$ bins are consistent with the redshifts of formation shown on their fig. 11, ranging from $z \simeq 1.1$ at $M_* \sim 10^{10}$ M$_\odot$ to $z \simeq 1.5$ at the highest masses. It can also be seen from fig. 12 of Siudek et al. (2017) that these results are broadly consistent with analyses by Moresco et al. (2010), Jørgensen & Chiboucas (2013) and Onodera et al. (2015), whereas the results of Thomas et al. (2010) and Choi et al. (2014) confirm the continuation of the trends I observe in Fig. 2.14 down to $z \simeq 0$.

Pacifici et al. (2016) perform an analysis more similar to that presented in this chapter than any of the studies mentioned above, both in terms of sample selection (see Fig. 2.6) and SFH recovery method, with the median SFHs they calculate also found to be well fitted by a double-power-law. Their redshifts of peak star formation, $z_{\rm peak}$, are therefore reasonably analogous to my $z_{\rm form}$. In their lowest-redshift bin ($0.2 < z < 0.5$) they find a range of $z_{\rm peak}$ between $\simeq 0.85$ at $M_* \sim 10^{10}$ M$_\odot$ and $\simeq 1.3$ at the highest masses; this is lower than the range of $z \simeq 1.5$ to 1.7 I find at similar observed redshifts in Fig. 2.14. This discrepancy is still apparent, though less marked at higher observed redshifts, up until $z_{\rm obs} \simeq 1.3$, after which these results are in good agreement. This disagreement is a consequence of differences in how the two redshifts are measured, with my $z_{\rm form}$ being higher than $z_{\rm peak}$ for SFHs that decline more rapidly than they rise (e.g. the middle panel of Fig. 2.7). The majority of the UltraVISTA galaxies do follow this pattern (see Section 2.6.2) and so have $z_{\rm form}$ higher than $z_{\rm peak}$.

The downsizing trend I observe is relatively weak, with only $\simeq 0.5$ Gyr in $t(z_{\rm form})$ separating the lowest and highest-mass galaxies in the sample at all observed redshifts. Pacifici et al. (2016) tentatively report a stronger downsizing trend at lower observed redshifts. This is not obvious in my results, with the strength of the downsizing trend across the whole observed redshift range being similar to that which Pacifici et al. (2016) see at higher observed redshifts. However, in my lowest-redshift bin, the UltraVISTA number counts are significantly lower than at higher redshifts, increasing the statistical uncertainties on my results, as can be seen in Fig. 2.14.

In summary, the UltraVISTA SFHs exhibit a downsizing trend broadly consistent with results from the literature. This trend is not particularly strong and does not evolve significantly with observed redshift. The median $t(z_{\rm form})$ value evolves by $\simeq 0.5$ Gyr over a range of ~ 1.25 dex in stellar mass, meaning the trend weakens as a fraction of the age of the Universe with decreasing observed redshift.

Decreasing formation redshift with decreasing observed redshift

Along with the downsizing trend discussed in Section 2.6.1, another common observation from similar analyses has been a trend towards lower average redshifts of formation with decreasing observed redshift for quiescent galaxies of all masses. This effect is observed for both $z_{\rm form}$ and $z_{\rm quench}$ in Fig. 2.13, and is clearly seen for $z_{\rm form}$ in Fig. 2.14. The effect is strongest at low redshift, with the lowest

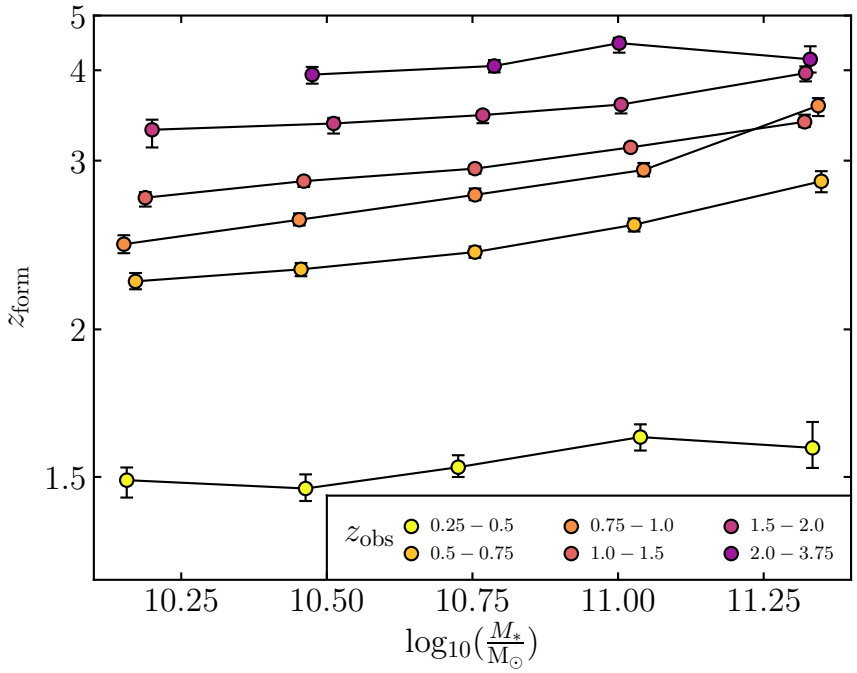

Figure 2.14 *Median $z_{\rm form}$ for each mass and redshift bin with their standard errors. Two clear trends are visible. Firstly, at fixed observed redshift, galaxies with higher stellar masses are seen to have higher median $z_{\rm form}$. This is the commonly observed downsizing trend. Secondly, the median redshift of formation decreases with observed redshift at all masses.*

observed-redshift bin being strongly offset from the others, and persists across the whole observed-redshift range of the UltraVISTA sample.

The simplest interpretation of this observation (as suggested by both Pacifici et al. 2016 and Siudek et al. 2017) relates to the continuing assembly of the red sequence throughout this observed-redshift range (e.g. Faber et al. 2007; Tomczak et al. 2014). This means that new and younger galaxies continue to quench throughout this time period, and so the median $z_{\rm form}$ and $z_{\rm quench}$ values at low redshift are averages over a diverse population that quenches at a wide range of redshifts. Conversely, the quenched galaxies in the higher observed-redshift bins are a biased sub-sample of these, as the quenched galaxy selection that has been applied to the sample necessarily excludes high redshift star-forming galaxies that will have quenched by lower redshifts.

The histograms plotted in Fig. 2.13 support this picture, with the distributions of $z_{\rm form}$ and $z_{\rm quench}$ observed to be far more extended in the lower-observed-redshift bins. The fact that the distributions of $z_{\rm quench}$ are extended all the way down to

the low-redshift ends of each observed-redshift bin is consistent with the assembly of the red sequence persisting throughout this observed-redshift range.

Another possible contributory factor to this trend of decreasing median $z_{\rm form}$ with decreasing observed redshift would be continuing evolution of the high-redshift quiescent population post-quenching, either through periods of rejuvenated star-formation, or merger events. This would mean that even if no galaxies joined the quiescent population across this observed redshift range, the median $z_{\rm form}$ I calculate at lower observed redshifts would still be lower than that which I observe at $z_{\rm obs} > 2$.Whilst my analysis confirms the well-known result that new galaxies join the red sequence across this observed redshift range (I see a strong increase in the number density of quiescent galaxies), it is also possible to use these results to gain insights into the roles of rejuvenated star-formation episodes and mergers.

A simple way of doing this is to compare the number density of quiescent galaxies at high redshift (e.g. $z_{\rm obs} > 1.5$) with the number density of extremely old galaxies at $0.25 < z_{\rm obs} < 0.5$ (those with $z_{\rm quench} > 1.5$). If all galaxies that are quenched at $z > 1.5$ undergo purely passive evolution from that time onwards, I would expect to observe the same number density of objects in the $0.25 < z_{\rm obs} < 0.5$ bin with $z_{\rm quench} > 1.5$ as I find for quiescent galaxies at observed redshifts above 1.5.

However, if galaxies that are quenched at $z > 1.5$ undergo further evolution after this time, I should observe fewer extremely old objects at low redshift, as newer stellar populations have been added to objects that were already quenched at $z > 1.5$ since that time, decreasing the $z_{\rm quench}$ measured for these objects at lower redshift.

I restrict this comparison to objects with $\log_{10}(M_*/{\rm M}_\odot) > 10.3$, so that I have a mass-complete sample across the whole observed redshift range. Of the 780 galaxies in this mass range with $z_{\rm obs} < 0.5$, just 27 have $z_{\rm quench} > 1.5$, a number density of $4.7 \pm 0.9 \times 10^{-5}$ Mpc^{-3}. By contrast, the number density for quiescent galaxies of these masses at $z_{\rm obs} > 1.5$ is $11.9 \pm 0.3 \times 10^{-5}$ Mpc^{-3}. This suggests that the majority (61±8 per cent) of massive galaxies that quench at $z_{\rm obs} > 1.5$ undergo significant further evolution post-quenching by $z = 0.5$. This is in agreement with recent results from IllustrisTNG (Nelson et al. 2018), who find significant mass growth of galaxies post-quenching. It should be noted that these uncertainties take into account only Poisson counting errors, and not the uncertainties in $z_{\rm quench}$ for individual objects.

2.6.2 How long did the process of quenching take?

A key piece of information necessary to constrain the physical mechanisms by which galaxies quench their star formation is the timescale over which quenching takes place. Accurate recovery of SFHs is the ideal method for studying quenching timescales, however, as discussed in Section 2.4.1, recovering physically realistic SFHs is challenging. Even if this can be achieved there are questions as to which parameters of the SFH are most representative of the quenching timescale, usually taken to mean the timescale over which the galaxy crosses the green valley, i.e. between an initial position on the star-forming main sequence and some threshold defining passivity.

Recently, a common approach has been to fit SFH models similar to the exponentially declining model of Equation 2.10 to spectral indices such as D4000 and the strength of the Hδ feature (e.g. Nogueira-Cavalcante et al. 2018), or UV-optical colours derived from broad-band photometry (e.g. Smethurst et al. 2018). The timescale over which the SFR decreases, τ is then used as a proxy for quenching timescale. However, when fitting this kind of SFH model, strong degeneracies between τ and T_0 mean that constraining τ is extremely challenging. Also, the effects of the age-metallicity degeneracy on D4000 and Hδ, and the age-metallicity-dust degeneracy on UV-optical colours must be properly treated, or the derived uncertainties in τ will be significantly underestimated.

In Section 2.4.2 I define an alternative scheme, in which the quenching timescale, $\Delta t_{\rm quench}$ is the time interval between the formation redshift, $z_{\rm form}$ (corresponding to the mass-weighted age of the galaxy) and the redshift at which the SFR of the galaxy falls below 10 per cent of its average value across the SFH of the galaxy, $z_{\rm quench}$. I finally define the normalised quenching timescale, $\tau_{\rm quench}$ to be $\Delta t_{\rm quench}$ as a fraction of the age of the Universe at $z_{\rm quench}$ (the fraction of the age of the Universe at the time of quenching that the galaxy takes to quench).

In Section 2.4.4, I confirm the ability of the double-power-law Bagpipes model to accurately recover both $z_{\rm form}$ and $z_{\rm quench}$. Hence, on average, the the correct values of $\tau_{\rm quench}$ should be recovered for the UltraVISTA objects. This allows information about the distribution of their SFH shapes to be obtained, allowing variations across the stellar-mass and observed-redshift bins shown in Fig. 2.12 to be considered.

Fig. 2.15 shows the distributions of posterior median values of $\tau_{\rm quench}$ for objects

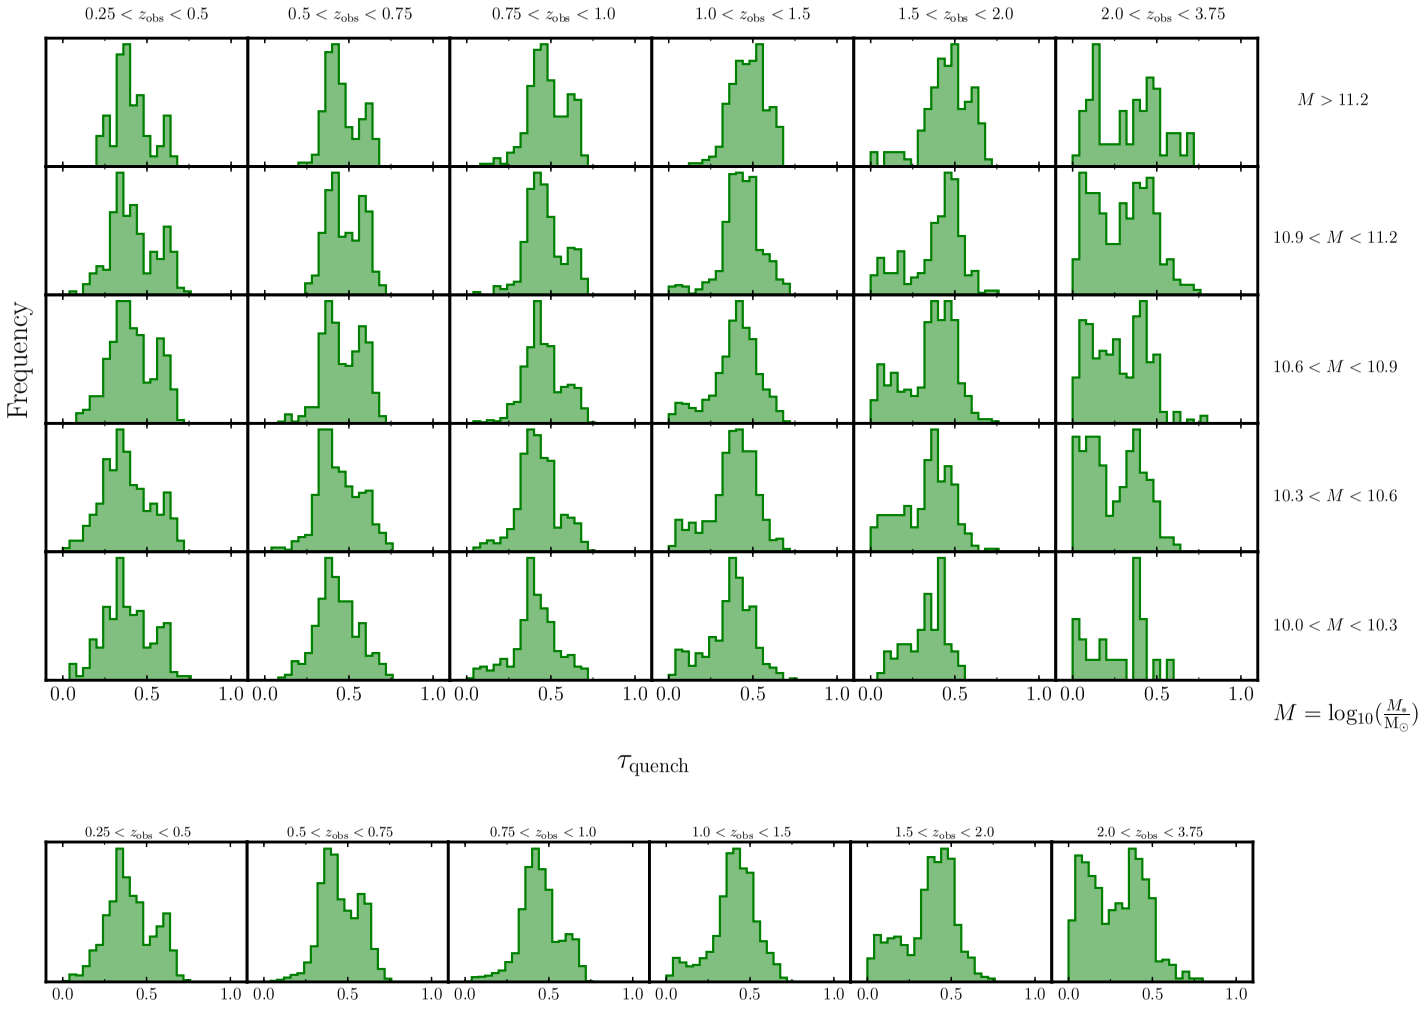

Figure 2.15 *Distributions of posterior median values of* $\tau_{\rm quench}$ *(quenching timescale as a fraction of the age of the Universe when the galaxy quenched) across the bins in stellar mass and observed redshift. The bottom row shows the results from the grid collapsed along the stellar-mass axis for improved statistics.*

in each of the stellar-mass and observed-redshift bins. Additionally, stacked histograms across all stellar-mass bins within a given observed-redshift bin are shown below the main grid. In Section 2.6.2 I discuss trends in these histograms across the stellar-mass and observed-redshift bins, and what these reveal about trends in the shapes of the SFHs. In Section 2.6.2 I perform a check to understand the effects of the priors imposed by the double-power-law SFH model on these trends. In Section 2.6.2 I consider what these results reveal about the quenching mechanisms that have influenced the SFHs of the galaxies in the UltraVISTA quenched sample.

Trends in SFH shape revealed by $\tau_{\rm quench}$

Firstly, in the highest observed-redshift bin, $2.0 < z_{\rm obs} < 3.75$, there is a strong peak in the distribution of $\tau_{\rm quench}$ centred on $\tau_{\rm quench} \simeq 0.1$ in each of the mass bins. This peak is also apparent, though reduced in strength, across the two redshift bins between $1.0 < z_{\rm obs} < 2.0$, and begins to take on a more obvious mass dependency, being observed most strongly at lower masses. A tail of small numbers of objects with $\tau_{\rm quench} < 0.2$ is apparent down to $z = 0.25$ at low masses. In contrast, virtually no objects are observed to fall into this region above $M_* = 10^{10.5}$ $\rm M_\odot$ and below $z_{\rm obs} = 1$. A $\tau_{\rm quench}$ of 0.1 corresponds to a quenching timescale of ~ 300 Myr at $z_{\rm quench} = 2$, or ~ 600 Myr at $z_{\rm quench} = 1$.

As demonstrated by the top panel of Figure 2.7, in order to reach such small values of $\tau_{\rm quench}$, SFHs must not only be rapidly quenched, they must also rise rapidly at late times, just before they quench. This shows that at $z_{\rm obs} \gtrsim 1$ there is a significant component of the quenched population with SFHs that both rise and fall rapidly, in a small fraction of the age of the Universe at their time of quenching. There is also a small component with this SFH shape at lower redshifts ($z \lesssim 1$) in the lower-mass bins.

Secondly, across all of the bins in stellar mass and observed redshift there is a strong peak centred around $\tau_{\rm quench} \simeq 0.4$. This accounts for the bulk of objects in the sample, and indicates an extended period over which star formation rises, followed by a relatively abrupt shutoff (e.g. the middle panel of Fig. 2.7). These SFHs have quenching timescales of just under half the age of the Universe at their time of quenching (e.g. ~ 1 Gyr at $z_{\rm quench} = 2$, or ~ 2 Gyr at $z_{\rm quench} = 1$), almost as rapid as the previous category of objects. However they are distinct from the previous category as they undergo a much more extended period of rising star

formation before they reach their peak.

Finally, a third peak begins to emerge at $z_{\rm obs} < 1.0$, centred around $\tau_{\rm quench} \simeq 0.6$. This population is more apparent at $0.5 < z_{\rm obs} < 1.0$ in the higher-mass bins, but by the lowest-redshift bin this peak is clearly visible even in the lowest-mass bin. These objects differ significantly from the previous two populations, in that their SFHs rise more quickly than they fall (e.g. the bottom panel of Fig. 2.7). These objects resemble the exponentially declining SFHs traditionally used in SED fitting, or the lognormal SFHs of Gladders et al. (2013), but represent a minority of the UltraVISTA sample. These objects have quenching timescales greater than half of the age of the Universe at $z_{\rm quench}$ (e.g. ~ 3 Gyr at $z_{\rm quench} = 1$).

These trends are more clearly visible in the bottom row of Fig. 2.15, where mass resolution has been sacrificed in order to obtain better statistics. The strongest peak at $\tau_{\rm quench} \sim 0.4$ is clearly dominant, whilst the presence and gradual decline of the population below $\tau_{\rm quench} = 0.2$ is clear at $z_{\rm obs} > 1$, and the presence and gradual rise of the peak centred on $\tau_{\rm quench} = 0.6$ is clear at $z_{\rm obs} < 1$.

Initially, Figure 2.15 appears to show surprisingly little variation in quenching timescales with stellar mass. However, it is important to remember that $\tau_{\rm quench}$ is normalised by the age of the Universe at $z_{\rm quench}$. As can be seen from Figure 2.13, more massive galaxies have higher median $z_{\rm quench}$ than less massive ones. This means that Figure 2.15 actually indicates more rapid quenching for more massive galaxies.

A check on the effects of priors

Considerable effort has already been made in Section 2.4 to understand the dependence of the results presented in this section on the model used to fit the observed photometry. Here, I perform one further test to directly check for potential biasing effects of the priors of the double-power-law SFH model on $\tau_{\rm quench}$. To do this I draw SFH parameters at random from the prior distributions listed in Table 2.3 and calculate $t(z_{\rm form})$, $t(z_{\rm quench})$ and $\tau_{\rm quench}$ for a statistical sample of SFHs.

The value of $\tau_{\rm quench}$ depends only on the shape of the SFH, not its normalisation, therefore the prior distribution on $\tau_{\rm quench}$ is independent of stellar mass. On the top row of Fig. 2.16, the prior distributions are plotted in purple. Only a very weak dependence on observed redshift exists, meaning the trends I observe across

the stellar-mass and observed-redshift bins cannot be a result of biasing effects of the chosen prior distributions.

It is still important however to understand how the overall shape of the $\tau_{\rm quench}$ distribution is affected by the choice of priors. As can be seen from Fig. 2.16, the prior distribution has strong, narrow peaks at $\tau_{\rm quench} = 0$ and 0.5. Upon experimentation it was found that these are due to allowing α and β for the double-power-law SFHs (see Equation 2.11) to reach very low and very high values respectively. Whilst these peaks in the prior do not correspond with any of the peaks discussed in Section 2.6.2, I still wished to understand how they might affect the results obtained.

In order to check that these peaks do not strongly affect the results obtained, I reduced the range of allowed α and β values from (0.01, 1000) to (0.1, 100), which was found through experimentation to produce a relatively flat prior on $\tau_{\rm quench}$. I then re-fitted the sample using this reduced range of allowed α and β values. This new prior distribution is shown in purple on the bottom row of Fig. 2.16, along with the results that were obtained by fitting with this new prior. It can be seen that the structure described in Section 2.6.2 is still clearly visible under this new, flattened prior.

Physical interpretation of trends in SFH shape

To briefly recap Section 2.6.2, I observe three peaks in the distributions of $\tau_{\rm quench}$ in Fig. 2.15, corresponding to three distinct SFH shapes. I refer to these three cases as (a), (b) and (c), which correspond roughly to the top, middle and bottom SFHs shown in Fig. 2.7 respectively. To summarise their properties, I have:

(a) A peak centred on $\tau_{\rm quench} \simeq 0.1$, corresponding to SFHs that rise and decline rapidly, in a small fraction of the age of the Universe at $z_{\rm quench}$. Strongly visible at $z_{\rm obs} > 1$, and also present in lower-mass bins at lower redshift. Physical quenching timescales of < 1 Gyr.

(b) A peak centred on $\tau_{\rm quench} \simeq 0.4$, corresponding to SFHs that decline more rapidly than they rise. Strongly visible in all of the bins of stellar mass and observed redshift, making up the majority of the UltraVISTA sample. Physical quenching timescales of $1 - 2$ Gyr.

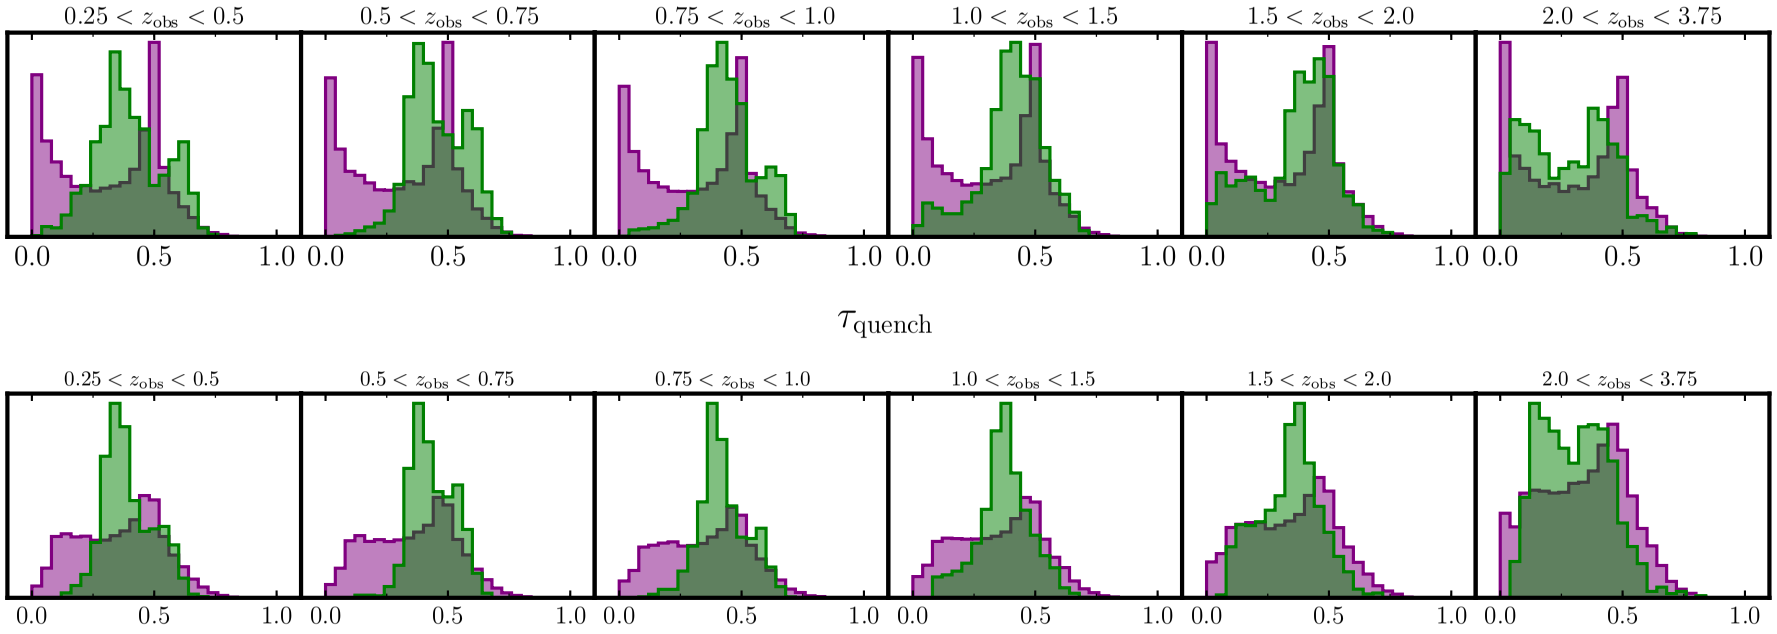

Figure 2.16 *Distributions of τ_{quench} in bins of observed redshift. The top row is the same as shown at the bottom of Fig. 2.15 with the prior distribution shown in purple. The bottom row shows the prior distribution and posterior median τ_{quench} values for the case in which the prior limits on α and β have been reduced to (0.1, 100) from their original ranges of (0.01, 1000). This change, whilst flattening the prior distribution of τ_{quench}, does not significantly affect the results obtained.*

(c) A peak centred on $\tau_{\rm quench} \simeq 0.6$ corresponding to SFHs that rise more rapidly than they decline. Only visible at $z_{\rm obs} < 1$ and most visible at higher masses, making up a further minority of theUltraVISTA sample. Physical quenching timescales of $\gtrsim 3$ Gyr.

As these SFH types are seen to be clearly distinct from each other in Fig. 2.15, it is reasonable to try to associate them with qualitatively different quenching processes. This is possible through comparing the physical quenching timescales inferred for each case to predictions for the behaviour of different physical processes.

Case (a) is associated with a process that rapidly quenches star formation, on timescales of < 1 Gyr. This process is most active at $z_{\rm obs} > 1$, and is also associated with a rapid rise in star-formation activity immediately preceding quenching. However, these SFHs need not necessarily be the result of a monolithic-collapse scenario. As discussed in Section 2.4.4, more recent star formation leaves a more significant impact on the observed spectrum. Therefore SFHs with a slow, shallow rise from the beginning of cosmic time, followed by a large, rapidly rising and quenching burst will be fitted with something like case (a) if a double-power-law form is used. This is because the more recent star-formation episode dominates the observed spectrum. More strongly constraining data would be needed to capture the detail of the early-time evolution of such SFHs.

The speed of quenching in case (a) is indicative of an ejective quenching process (see Section 2.1), as shutting off the supply of new gas would not be sufficient to cause the almost immediate cessation of star-formation activity. Quasar-mode feedback is one process that has the capability to cause extremely rapid quenching, and, if triggered by major-merger events, would also be expected to be much more common at high redshift, as the rate of major mergers is known to be a strong function of redshift (e.g. Lotz et al. 2011). The merger-triggered, quasar-mode feedback scenario would also explain the rapid rise of star-formation activity that is observed immediately before these galaxies quench.

Virtually no examples of this kind of SFH are observed in the lowest-redshift bins, except at the lowest masses, supporting the conclusion I reach in Section 2.6.1 that the majority of galaxies that are observed to be quenched at high redshift undergo further interactions that affect the shapes of their SFHs when they are observed later in cosmic time. This is consistent with the findings of

Gabor et al. (2011), who conclude that systems that quench due to AGN-driven outflows should begin to re-accrete gas and resume star formation within 1−2 Gyr.

The tidal/ram-pressure stripping experienced by galaxies falling into clusters is another process that can generate SFHs with this shape (e.g. Diemer et al. 2017). This could be responsible for the tail of objects at low masses at $z_{\rm obs} < 1$ with this kind of SFH, as the most massive galaxies are unlikely to become satellites, and mergers are less common later in cosmic time.

Case (b) is the dominant mode of quenching within the UltraVISTA sample, and is associated with a process that quenches star formation on intermediate timescales of $\sim$ 1−2 Gyr. The dominant mode of quenching in modern hydrodynamic simulations such as Mufasa (Davé et al. 2016) and IllustrisTNG (Nelson et al. 2018) is AGN-driven feedback in the low-accretion (jet) mode, the implementation of which has been shown to significantly improve agreement with the observed properties of the $z = 0$ red sequence (e.g. Gabor & Davé 2015). The typical timescale for quenching due to this process in IllustrisTNG is ~ 1.6 Gyr, which is consistent with what I observe in case (b). I therefore tentatively associate case (b) with AGN-driven feedback in the low-accretion mode.

Finally, case (c) is associated with some process that leads to a gradual decline in star formation over a period of $\gtrsim 3$ Gyr. Clearly this process must be preventative, rather than ejective, however it is the most challenging to associate with a known physical mechanism. As previously noted, this kind of SFH is similar in shape to an exponentially declining or lognormal function. This, and the fact that these objects are observed earliest ($z \simeq 1$) at the highest stellar masses, suggests a scenario similar to the model of Gladders et al. (2013). In this model, star formation in individual galaxies naturally dies down in the same way as the cosmic SFR density as the cosmic gas supply gradually dwindles. In the Gladders et al. (2013) model, the rate is determined, in part, by stellar mass, with the most massive galaxies reaching the end-point of their evolution the fastest. A similar scenario for the quenching of low-redshift galaxies was also presented by Schreiber et al. (2016).

Having tentatively identified the quenching processes at work, it is interesting to consider how their relative importances vary as a function of stellar mass. As was noted in Section 2.6.2, the distributions of $\tau_{\rm quench}$ are remarkably consistent across the bins in stellar mass, which translates into faster physical quenching

timescales for more massive galaxies, as these are found to be quenched at higher redshifts on average (Figure 2.12).

This finding is in agreement with previous work (e.g. Thomas et al. 2010; Pacifici et al. 2016). However the consistency of the distributions of $\tau_{\rm quench}$ across the stellar mass bins suggests that faster quenching at higher stellar masses is a further consequence of mass-accelerated evolution (or downsizing), rather than a consequence of a change in quenching mechanism across the range of masses probed ($M_* > 10^{10}\mathrm{M}_\odot$). The fact that case (c) galaxies (which seem to be associated with a more advanced stage in galaxy evolution) first begin to be seen at the highest masses lends further weight to this theory. The same argument applies to the fact that case (a) galaxies are observed down to lower observed redshifts only at lower masses. These less-massive galaxies appear to be following their more-massive counterparts along the same evolutionary track at a slower pace.

2.7 Conclusions

In this chapter I present the new galaxy spectral model generation framework and fitting tool Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation, or Bagpipes, which can be used to rapidly generate complex, physically realistic model galaxy spectra and fit these to arbitrary combinations of spectroscopic and photometric data using the MultiNest nested sampling algorithm.

I verify the ability of Bagpipes to recover the properties of the SFHs of quenched galaxies by fitting mock observations for a sample of simulated massive quenched galaxies with realistic SFHs, chemical-enrichment histories and dust properties from the Mufasa suite of cosmological hydrodynamic simulations. I investigate thoroughly the effects of different SFH models and priors on the biases in derived SFH properties. I find that the use of an exponentially declining SFH model produces a small but significant overestimation of the stellar masses of quenched galaxies, and an overestimation of the redshifts at which their stellar masses were assembled. A significant improvement is observed when fitting a double-power-law SFH model with the priors listed in Table 2.3.

I then use Bagpipes to perform a detailed analysis of the SFHs of a sample of

9289 quenched galaxies from the UltraVISTA Survey with stellar masses, $M_* > 10^{10}$ $M_\odot$ across an observed redshift range from $0.25 < z_{\rm obs} < 3.75$. I make the following observations based on my analysis:

1. I observe a clear downsizing trend in my results (Fig. 2.14), with the most massive galaxies being on average $\simeq 0.5$ Gyr older than those with $M_* \simeq 10^{10}$ $M_\odot$. This appears to remain relatively constant across the whole observed-redshift range.

2. A simple analysis of the redshift evolution of the number density of galaxies with very old stellar populations (Section 2.6.1) suggests that the majority of galaxies that quench at $z > 1.5$ will undergo significant further evolution through rejuvenated star-formation or merger events by $z = 0.5$.

3. The SFHs of the UltraVISTA quenched galaxies fall into three distinct classes with different shapes (Fig. 2.15). I refer to these as cases (a), (b) and (c), which roughly correspond to the top, middle and bottom panels of Fig. 2.7 respectively. I propose that these different cases correspond to different quenching mechanisms (Section 2.6.2).

4. Galaxies with SFHs of type (a) are common at $z_{\rm obs} > 1$. They experience a rapid rise followed by a rapid decline in SFR over a small fraction of the age of the Universe at the redshift of their quenching, and have rapid quenching timescales of < 1 Gyr. I tentatively identify these as the result of (possibly merger-triggered) quasar-mode AGN feedback.

5. Galaxies with SFHs of type (b) make up the bulk of the UltraVISTA quenched sample across all redshifts and masses. They experience a slower rise and more rapid decline in SFR, with intermediate quenching timescales of $1-2$ Gyr. These SFHs match predictions from simulations for quenching by low-accretion (jet) mode AGN feedback (e.g. Davé et al. 2017; Nelson et al. 2018).

6. Galaxies with SFHs of type (c) begin to enter the UltraVISTA sample at $z_{\rm obs} < 1$. Their SFHs rise more rapidly than they fall, and have long quenching timescales of $\gtrsim 3$ Gyr. These are not obviously associated with a specific physical process, and I speculate that they may be the first galaxies

in which star formation 'naturally' dies down with the decrease of the overall cosmic gas supply (e.g. see Schreiber et al. 2016).

7. In common with other studies, I observe more rapid quenching at higher stellar masses. However, in contrast to some previous studies, I see no clear evidence of a change in the relative importances of different quenching mechanisms with stellar mass (at masses greater than $M_* = 10^{10}$). My results suggest that this trend is an extension of the well-known mass-accelerated evolution (or downsizing) trend, which appears to affect the quenching timescale as well as the redshift of formation.

To summarise, my results suggest that the red sequence at $M_* > 10^{10}$ $M_\odot$ has been built up since $z \sim 4$ by a number of different quenching processes that make different relative contributions at different epochs. At $z \gtrsim 1$, (potentially merger-triggered) quasar-mode AGN feedback plays a significant role in the rapid but short-lived quenching of galaxies. Throughout cosmic time, AGN feedback in the low-accretion mode is the dominant quenching mechanism, being responsible for the bulk of the build-up of the red sequence. However galaxies that quench by this processes at $z \gtrsim 1$ typically either experience some periods of rejuvenated star formation at later times or merge with younger stellar populations, which further reduces the number of extremely early-quenching objects observed at low redshift. Finally, at $z \lesssim 1$, a third mode of 'natural' quenching begins to act, shutting down star-formation much more slowly as the cosmic gas supply declines. It is also possible that at lower masses I see rapid quenching due to tidal/ram-pressure stripping of satellite galaxies. This is qualitatively similar to the model of Schawinski et al. (2014). A joint analysis combining morphology and environment with detailed SFH recovery would provide an additional test of the scheme I propose.

In this chapter I have considered a large sample of galaxies with high quality photometric data. These data are nevertheless weakly constraining on the SFHs of individual objects (e.g. Fig. 2.4). I have relied instead upon the large UltraVISTA sample to obtain information about the distribution of the SFH properties of quenched galaxies across bins in stellar mass and observed redshift. This was motivated, in part, by a desire to establish a baseline for how well the SFHs of quenched galaxies could be constrained using photometric observations alone.

Recently, several authors (e.g. Pacifici et al. 2012; Thomas et al. 2017) have

confirmed that significantly better constraints on physical parameters can be obtained by fitting combinations of rest-frame ultraviolet spectroscopy and multi-wavelength photometry. In Chapter 4 I will extend my analysis to fitting spectroscopy from VANDELS (McLure et al. 2018a; Pentericci et al. 2018) with Bagpipes, in order to obtain more detailed information on the properties of individual SFHs, to build up a more detailed understanding of the quenching processes that affect these galaxies.

Chapter 3

How to measure star-formation histories with parametric models

The material in this chapter was originally published in Carnall et al. (2019a). A companion paper, Leja et al. (2018), deals with non-parametric models.

3.1 Introduction

One of the most important processes driving the evolution of galaxies is star formation. This means that the stellar masses and star-formation rates (SFRs) of galaxies are two of their most fundamental properties. Measurements of these quantities underpin many of the most important results in the study of galaxy formation, such as the redshift evolution of the galaxy stellar-mass function (GSMF; e.g. Tomczak et al. 2014), cosmic star-formation-rate density (SFRD; e.g. Madau & Dickinson 2014) and the galaxy star-forming main sequence (SFMS; e.g. Speagle et al. 2014).

To measure these quantities we rely on modelling and fitting the observed spectral energy distributions (SEDs) of galaxies (see Conroy 2013), using methods ranging from monochromatic SFR indicators and single-colour mass-to-light relationships to full spectral fitting.

Models used to fit galaxy SEDs normally include a star-formation history (SFH). The fitted SFH is then used to derive the SFR and stellar mass (as distinct

from the total stellar mass formed; see Section 3.3), as well as other quantities of interest such as the specific SFR (sSFR) and mass-weighted age. This means that the priors on these quantities are not set explicitly, but instead are set implicitly by the priors applied to the SFH. Because of this, SFH priors affect results obtained for many of the most fundamental galaxy properties.

As the choice of prior is subjective, it is generally desirable for the conclusions reached to be as insensitive to the prior assumptions made as possible. For the prior to be of minimal importance, two conditions must be met. Firstly, the model being fitted must be capable of accurately describing the data-generating processes. Secondly, the data must be strongly constraining on the values of all model parameters. The problem of fitting SFHs to galaxy SEDs is particularly challenging because, in general, neither of these conditions are met.

In this case the data-generating processes encompass the entire physics of galaxy formation and evolution, and are thus extremely complex. It has been shown through simulations of galaxy formation (e.g. Davé et al. 2016; Nelson et al. 2018) that these processes can give rise to a huge diversity of complex galaxy SFHs, presenting a significant modelling challenge.

Progress has been made in addressing this challenge, however there is still much debate in the literature as to the best approach. Three different families of models are in common usage, all of which impose substantially different priors on SFHs, both in terms of the range of allowed shapes and the relative prior weights assigned to different allowed shapes. Firstly, parameterised SFH models, which are the subject of this chapter (see Section 3.2), secondly non-parametric models (e.g. Cid Fernandes et al. 2005; Ocvirk et al. 2006; Tojeiro et al. 2007; Cappellari 2017; Leja et al. 2017; Chauke et al. 2018), and thirdly models drawn directly from simulations (e.g. Brammer et al. 2008; Pacifici et al. 2012). At the core of this debate is a trade-off between flexibility and computational tractability, with more flexible models generally being more computationally intensive to fit.

Even if a SFH model can be defined that is both computationally tractable and flexible enough to encompass the inherent complexity of galaxy SFHs, the priors assumed for its parameters will still be important unless all of the model parameters are well constrained by the data. However, as was extensively demonstrated by Ocvirk et al. (2006), the problem of inferring SFHs from galaxy SEDs is poorly conditioned, meaning even small perturbations of the data can lead to large perturbations of the inferred SFH. The main underlying reason for

this is the strong evolution of the mass-to-light ratios of stellar populations with age, meaning that the early-time evolution of the SFH has little effect on the observed SED at later times if star-formation is ongoing.

The situation therefore is one in which inferences made about key galaxy physical parameters are highly sensitive to the SFH prior (e.g. see Wuyts et al. 2009, 2011; Lee et al. 2009; McLure et al. 2011; Pforr et al. 2012; Mobasher et al. 2015; Salmon et al. 2015; Iyer & Gawiser 2017). SFH priors are therefore a subject of prime importance for all SED fitting analyses.

In this context it is important to critically evaluate SFH priors to see whether the priors they impose on parameters of interest are well-motivated. Two significant risks are the over-interpretation of data if observational uncertainties are not carefully propagated, and the unintentional imposition of informative priors, which can lead to strong posterior constraints even in the absence of strongly constraining data. It is possible that issues of this nature contribute to the known tensions between the observed GSMF, SFRD and SFMS (e.g. Madau & Dickinson 2014; Leja et al. 2015; Pacifici et al. 2015).

Recently, significant advances in statistical and computational techniques (e.g. Skilling 2006; Feroz et al. 2009, 2013; Goodman & Weare 2010; Acquaviva et al. 2011; Foreman-Mackey et al. 2013) have made it possible to rapidly fit complex galaxy SED models to data within a fully Bayesian framework, including full control of the applied priors. A new generation of galaxy SED fitting tools has been built to exploit this, such as Beagle (Chevallard & Charlot 2016), Prospector (Leja et al. 2017; Johnson et al. in preparation). These codes allow detailed analyses of the effects of SFH priors, and direct comparisons between results obtained under different assumptions.

In this chapter, I use Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation (Bagpipes; see Chapter 2) to conduct an investigation into the priors imposed by four commonly used parametric SFH models: exponentially declining, delayed exponentially declining, lognormal and double power law.

In Section 3.2 I introduce the parametric SFH models. In Section 3.3 I consider the effects of these models in the limit of weakly constraining data by discussing the priors they impose on physical parameters, in particular sSFR and mass-weighted age/formation time. In Section 3.5, I consider the biases introduced in the case of fitting these models to high signal-to-noise ratio (SNR) broad-band

photometric data. This is achieved by constructing and fitting a simple mock galaxy catalogue, designed to span a wide range of scenarios for the formation of galaxies. I also discuss how well these mock data can discriminate between different SFH models.

Finally, in Section 3.6, I perform a long-discussed consistency check (e.g. Heavens et al. 2004) by using these four models to fit a volume-complete sample of local galaxies and comparing the redshift evolution of the cosmic SFRD inferred from their SFHs to the relationship obtained by measuring SFRs across cosmic time by Madau & Dickinson (2014). For this analysis I use a sample in the redshift interval $0.05 < z < 0.08$ with high-quality photometric data and spectroscopic redshifts from the Galaxy and Mass Assembly (GAMA) Survey (Driver et al. 2009, 2016; Baldry et al. 2018).

Throughout this chapter, all times, t, are measured forwards from the beginning of the Universe such that $t(z)$ is the age of the Universe at redshift z. I also assume a Kroupa (2001) initial mass function (IMF).

3.2 Parametric star-formation-history models

Parametric models approximate galaxy SFHs using simple functional forms, typically involving $2-3$ shape parameters plus a normalisation. These are, in some respects, the least flexible option for fitting SFHs, imposing strong prior limitations on the range of allowable shapes. However their relative speed and simplicity of fitting means that they are widely used. It has also been demonstrated that complex SFHs from simulations can be reasonably well described by parametric models (e.g. Simha et al. 2014, Diemer et al. 2017).

The four widely used parametric SFH models I consider (and some recent examples of their use in the literature) are: exponentially declining (Mortlock et al. 2017; Wu et al. 2018a; McLure et al. 2018a), delayed exponentially declining (Ciesla et al. 2017; Chevallard et al. 2019), lognormal (Diemer et al. 2017; Cohn 2018) and double power law (see Chapter 2) For brevity I will refer to them as tau, delayed, lognormal and DPL models respectively.

The models are introduced in the following four subsections; the prior probability densities I assume for their parameters are reported in Table 3.1. I have chosen a simple set of prior probability densities in each case to act as a basis for

Table 3.1 *Parameters and prior distributions for each SFH model. Logarithmic priors are uniform in log_{10} of the parameter.*

Model	Parameter	Symbol / Unit	Range	Prior
Tau/	Start time	T_0 / Gyr	$(0, t_{\rm obs} - 0.1)$	uniform
Delayed	Timescale	τ / Gyr	(0.3, 10)	uniform
Lognormal	Peak time	$t_{\rm max}$ / Gyr	(0.1, 15)	uniform
	FWHM	$T_{\rm FWHM}$ / Gyr	(0.1, 20)	uniform
DPL	Falling slope	α	(0.1, 1000)	logarithmic
	Rising slope	β	(0.1, 1000)	logarithmic
	Turnover	τ / Gyr	$(0.1, t_{\rm obs})$	uniform
Global	Normalisation	$M_{\rm formed}$ / $\rm M_\odot$	$(1, 10^{13})$	logarithmic

comparison. The effects of changing these will be discussed in Section 3.3.2. In all cases I normalise the SFH by the total stellar mass formed, $M_{\rm formed}$ at the time of observation, $t_{\rm obs}$. I assign a logarithmic prior to $M_{\rm formed}$, which is the minimally informative prior when the uncertainty spans several orders of magnitude (e.g. Simpson et al. 2017).

3.2.1 Exponentially declining

Exponentially declining SFHs (tau models) are probably the most commonly applied SFH model. They assume that star-formation jumps from zero to its maximum value at some time T_0, after which star-formation declines exponentially with some timescale τ,

$$\mathrm{SFR}(t) \propto \begin{cases} \exp\left(-\frac{t - T_0}{\tau}\right) & t > T_0 \\ 0 & t < T_0. \end{cases} \tag{3.1}$$

Tau models are often used as a fiducial model against which others can be compared, however they have been shown to become less appropriate at higher redshifts (e.g. Reddy et al. 2012) as they cannot reproduce rising SFHs. They have also been shown to produce biased estimates of stellar mass, SFR and mass-weighted age when used to fit mock observations of simulated galaxies (e.g. see Chapter 2; Simha et al. 2014; Pacifici et al. 2015). The priors listed in Table 3.1

for this model are adapted from those used by Wuyts et al. (2011).

3.2.2 Delayed exponentially declining

A simple extension of exponentially declining SFHs are delayed exponentially declining SFHs (delayed models). Multiplying the tau-model SFR by the time since T_0 removes both the discontinuity in SFR at T_0 and the condition that star formation can only decline after that point. This results in a more flexible and more physical model, which can reproduce rising SFHs if τ is large. The SFR is now described by

$$\mathrm{SFR}(t) \propto \begin{cases} (t - T_0)\,\exp\left(-\dfrac{t - T_0}{\tau}\right) & t > T_0 \\ 0 & t < T_0. \end{cases} \tag{3.2}$$

For the delayed model I apply the same prior probability densities for τ and T_0 as for the tau model.

3.2.3 Lognormal

Lognormal models for individual galaxy SFHs were postulated by Gladders et al. (2013), based, in part, on the evolution of the cosmic SFRD being well fitted by the lognormal function. The SFR is described by

$$SFR(t) \propto \frac{1}{t}\,\exp\left(-\frac{\left(\ln(t) - T_0\right)^2}{2\tau^2}\right) \tag{3.3}$$

where τ and T_0 are free parameters. Because these parameters do not have intuitive interpretations (e.g. star-formation does not peak at $t = e^{T_0}$) I follow Diemer et al. (2017) in re-parameterising in terms of t_{max}, the time at which star-formation peaks, and T_{FWHM}, the full-width at half maximum of the SFH. I constrain t_{max} to be less than 15 Gyr after the beginning of the Universe and T_{FWHM} to be less than 20 Gyr in order to limit the prior volume containing models with very large t_{max} and T_{FWHM} that have almost identical shapes at times earlier than the $z = 0$ age of the Universe.

3.2.4 Double power law

The double power law (DPL) function introduces another free parameter in order to separate the rising and declining phases of the SFH, which are modelled by two separate power-law slopes. This function has been shown to provide a good description of the redshift evolution of the cosmic SFRD (Behroozi et al. 2013b; Gladders et al. 2013), as well as producing good fits to SFHs from simulations in Chapter 2 (see also Pacifici et al. 2016; Diemer et al. 2017). The functional form is

$$\mathrm{SFR}(t) \propto \left[\left(\frac{t}{\tau}\right)^{\alpha} + \left(\frac{t}{\tau}\right)^{-\beta}\right]^{-1} \tag{3.4}$$

where α is the falling slope, β is the rising slope and τ is related to (but not the same as) the peak time. The priors reported in Table 3.1 are based on those used in Chapter 2.

3.3 Priors on physical parameters

As described in Section 3.1, in most galaxy SED fitting analyses, the priors on the parameters of interest (e.g. stellar mass, star-formation rate, mass-weighted age) are not set explicitly, instead being set implicitly by the priors applied to the SFH. In this section I use Bagpipes to sample the prior distributions listed in Table 3.1 for the models described in Section 3.2 and report and discuss the priors imposed on parameters of interest.

For each draw from the prior, I obtain the stellar mass, M_* (which is the total mass in stars and remnants at the time of observation), by integrating the SFH multiplied by the mass-return fraction as a function of stellar age. This is distinct from the normalisation of the SFH, M_{formed} described in Section 3.2, which is the total stellar mass formed. I also average over the most recent 100 Myr of the SFH to obtain an estimate of the current SFR, which I denote SFR_{100}.

The dependence of M_* on the shape of the SFH is relatively weak, meaning that the prior on M_* is largely independent of the parametric form used. The M_* prior instead closely mirrors the prior placed on M_{formed}, typically with an offset of ~ 0.2 dex (the return fraction). I therefore do not discuss stellar masses in this

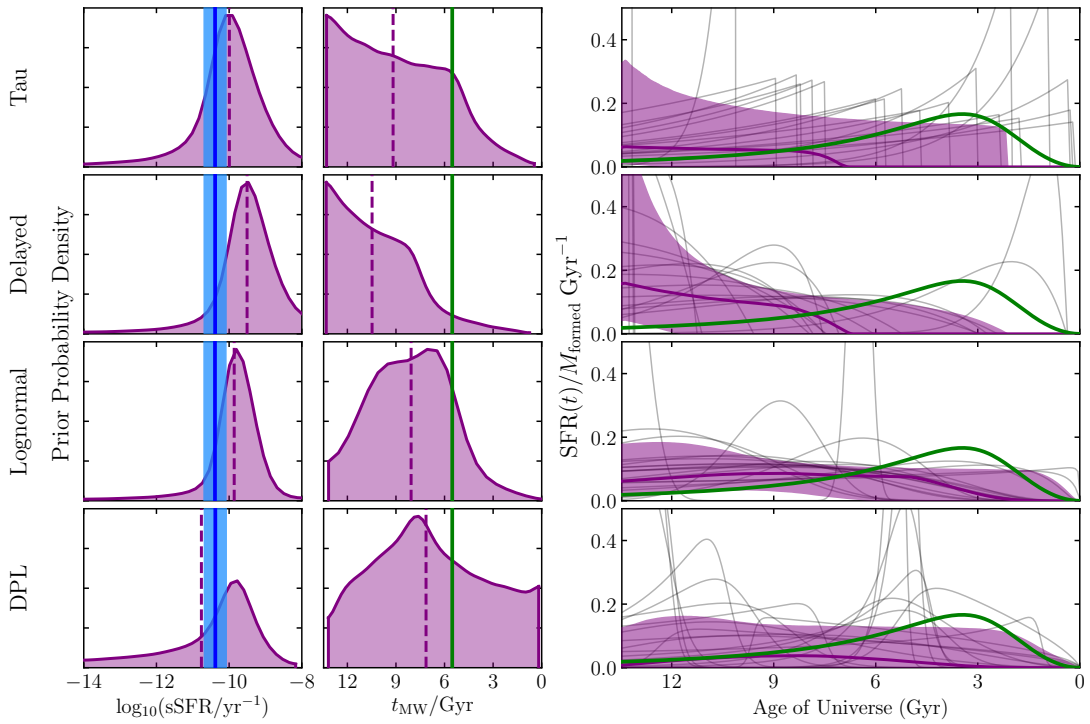

Figure 3.1 *Prior probability densities for sSFR, mass-weighted formation time and SFH shape imposed by each of the four parametric SFH models (see Section 3.2 and Table 3.1). To the left, the prior medians are shown as dashed purple lines. The $z = 0$ SFMS of Speagle et al. (2014) for a representative stellar mass of $M_* = 10^{10.5}\mathrm{M}_\odot$ (close to the characteristic mass of the GSMF) is shown on the far-left panels as a blue solid line. The blue shaded region shows a scatter of 0.3 dex. The mass-weighted formation time for the stars in the Universe, derived from the Madau & Dickinson (2014) SFRD curve, is shown as a solid green line on the centre-left panels. To the right, the solid purple line is the prior median and the shaded region shows the 16^{th}–84^{th} percentiles. A sample of draws from each prior is shown in grey. The Madau & Dickinson (2014) SFRD curve is shown in green.*

section, however it should be noted that the SFH prior can still bias measurements of M_* if the model cannot reproduce the true SFH shape (see Section 3.5).

The SFR prior has significant dependencies on both the SFH model and the stellar mass. In order to isolate the effect of the SFH model, I will consider the prior on the specific star-formation rate, sSFR, which I calculate by $\mathrm{sSFR} = \mathrm{SFR}_{100}/M_*$.

The second physical parameter I consider in this section is the mass-weighted formation time, t_{MW}. This corresponds to the more commonly used mass-weighted age, but is measured forwards from the beginning of the Universe to maintain consistency between objects at different observed redshifts. I calculate

this by

$$t_{\rm MW} = \frac{\int_0^{t_{\rm obs}} t\ {\rm SFR}(t)\ dt}{\int_0^{t_{\rm obs}} {\rm SFR}(t)\ dt}. \tag{3.5}$$

This parameter gives an indication of the epoch at which the stellar masses of galaxies were assembled, which should ultimately agree with the measured redshift evolution of the cosmic SFRD (see Section 3.6).

In Section 3.3.1 I report the priors imposed on sSFR and mass-weighted formation time by our fiducial models at $z = 0$. Then, in Section 3.3.2, I consider the effects of varying the prior probability densities for individual model parameters. Finally, in Section 3.3.3, I explore the effects of moving to higher observed redshifts.

3.3.1 Fiducial models at redshift zero

I first sample the fiducial prior distributions shown in Table 3.1 at a fixed redshift of $z = 0$. Fig. 3.1 shows the priors imposed on sSFR, $t_{\rm MW}$ and SFH shape by the four parametric models.

The far-left panels show the priors imposed on sSFR, with the consensus $z = 0$ SFMS reported by Speagle et al. (2014) shown as a solid blue line (the shaded region represents a scatter of 0.3 dex). It can be seen that the sSFR priors imposed by all four models are strongly peaked around sSFR $\sim 10^{-10}$ yr^{-1}, with tails out to lower sSFRs. A limit of sSFR $\leq 10^{-8}$ yr^{-1} is imposed by the use of a SFR timescale of 100 Myr.

Given that the SFRs of galaxies are poorly constrained by the observed SED (see Section 3.1), the fact I observe different sSFR priors for different models is one of the main reasons that different locations are observed for the SFMS in different studies (e.g. Speagle et al. 2014; Pacifici et al. 2015). However, the strongly peaked (and hence informative) nature of these prior distributions also suggests another possibility: that galaxy SFRs measured assuming these parametric SFH models could be driven towards a narrow range of sSFR values by the SFH prior. This could cause a SFMS to be observed even when the data does not have the necessary constraining power to infer reliable SFRs, or artificially tighten an otherwise less-well-defined SFMS.

The centre-left panels show the priors imposed on mass-weighted formation time, with the mass-weighted formation time for the stellar population of the Universe at $z = 0$, as calculated from the Madau & Dickinson (2014) SFRD curve, shown by a solid green line. It can be seen that all of the models considered favour larger t_{MW} than Madau & Dickinson (2014), with the tau and delayed models in particular favouring young galaxy stellar populations. The lognormal and DPL models produce broader priors on t_{MW}, but still favour later times than the epoch at which it is known that galaxies assembled the majority of their stellar masses. The right panels show the priors imposed on SFH shape by dividing out the dependency on M_{formed}. The Madau & Dickinson (2014) curve is also shown, clearly demonstrating that these fiducial priors favour later stellar-mass assembly.

In response to issues of this nature, there has been a recent rapid increase in the diversity and complexity of parametric SFH models in use in the literature (e.g. Ciesla et al. 2017; Glazebrook et al. 2017; Merlin et al. 2018; Schreiber et al. 2018). This increasing diversity in applied methodologies necessitates a firm basis for comparisons between different methods.

The method presented in this section (sampling SFHs from the prior probability distributions on model parameters, then using these to construct the priors on the parameters of interest) is applicable to any kind of SFH model. I therefore suggest that these tests should always be conducted when a new model is used, in order to understand the priors being imposed, and to facilitate comparisons with other methods.

This method illustrates the power of the fully Bayesian methodology being pioneered in SED fitting analyses. To give an example of how this could be used to refine a SFH model, a more physical prior could be motivated by the ansatz that the shape prior should resemble the overall cosmic SFRD (e.g. Gladders et al. 2013). In this case we would wish to find a prior that rises rapidly to a peak around $z = 2$ on the right panels of Fig. 3.1 and then declines more slowly, with t_{MW} peaked around the green line on the centre-left panels (see Section 3.4). Another option would be to design a prior centred on the SFH of a galaxy that follows the SFMS throughout its evolution (e.g. Ciesla et al. 2017). If no compelling argument can be made for a prior belief that the distribution of galaxy sSFRs at $z = 0$ is strongly peaked at $\sim 10^{-10}$ yr^{-1}, then a less informative prior on sSFR would also be desirable.

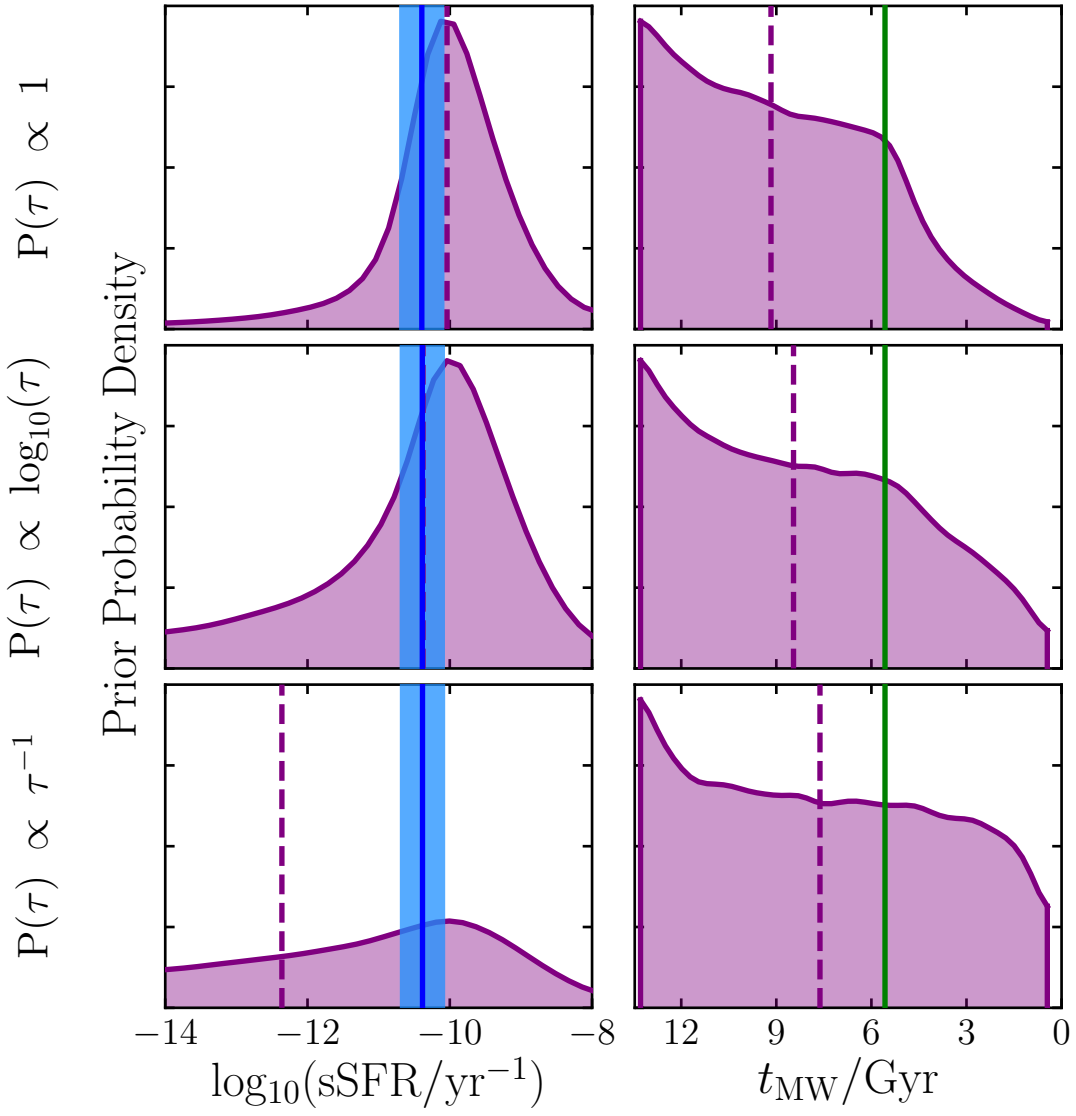

Figure 3.2 *Priors imposed on sSFR and mass-weighted formation time by the exponentially declining SFH parameterisation (Equation 3.1) under the assumption of different prior probability densities for the τ parameter. Other priors are the same as listed in Table 3.1. Vertical lines are as in Fig. 3.1.*

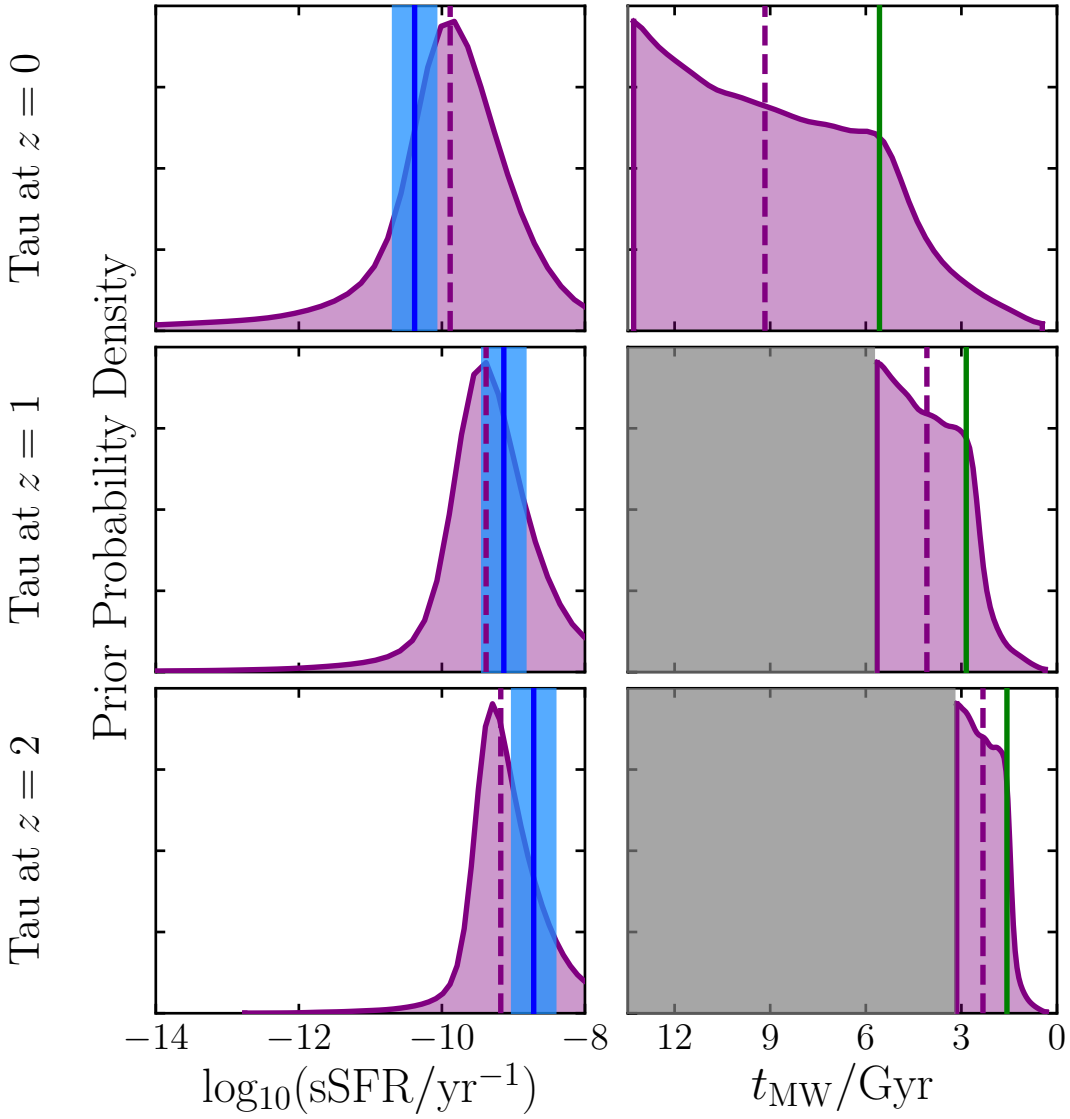

Figure 3.3 *Priors imposed on sSFR and mass-weighted formation time by the exponentially declining SFH parameterisation (Equation 3.1) with the priors shown in Table 3.1 at different observed redshifts (different T_0 priors). Vertical lines are as described in the Fig. 3.1 caption at the redshifts shown.*

3.3.2 The effects of changing the prior probability densities assumed for model parameters

The discussion of Section 3.3.1 refers only to one set of fiducial prior probability densities for the parameters of the four SFH models, chosen to be typical of applications of these models in the literature. Changing any of these will affect the priors shown in Fig. 3.1, with a variety of different combinations possible.

Whilst Bagpipes can be used to obtain priors on physical parameters for any of these combinations, I here restrict the discussion to two cases of particular interest involving the prior probability density assumed for the τ parameter of the exponentially declining SFH model. These provide representative examples from which the magnitude of the effects of such changes can be seen. Another example will be provided in Section 3.3.3, where I will consider changing the observed redshift. This is effectively a specific case of changing the prior on the T_0 parameter. An expanded discussion of the prior probability densities for the lognormal and DPL models is also provided in Section 3.4.

An alternative parameterisation of the tau model to that given in Equation 3.1 is $\mathrm{SFR}(t) \propto e^{-\gamma(t-T_0)}$, where $\gamma = \tau^{-1}$ (e.g. Pacifici et al. 2015; Salim et al. 2016). A uniform prior applied to γ corresponds to a prior probability density on τ of $P(\tau) \propto \tau^{-1}$, as opposed to a uniform prior on τ. The use of a logarithmic prior on τ is also often discussed. I therefore test two alternative prior probability densities for τ, $P(\tau) \propto \log_{10}(\tau)$ and $P(\tau) \propto \tau^{-1}$ whilst maintaining the ranges and priors for T_0 and M_{formed} quoted in Table 3.1.

Fig. 3.2 shows the priors on sSFR and mass-weighted formation time for the two alternative τ priors, as well as the original uniform prior from Fig. 3.1 for reference. It can be seen that both of these alternative priors broaden the prior on sSFR towards lower values, with the distribution under the assumption of $P(\tau) \propto \tau^{-1}$ being relatively flat. Additionally, both of these priors extend the distribution of mass-weighted formation times further towards earlier times, again with the assumption of $P(\tau) \propto \tau^{-1}$ producing a relatively flat prior on t_{MW}.

I conclude, by a comparison of Figures 3.1 and 3.2, that changing the prior probability densities assumed for SFH model parameters can produce significant changes in the priors imposed on the parameters of interest, of the same order of magnitude as changing between different SFH parameterisations. I also conclude that when using the exponentially declining SFH model, a prior probability

density for the τ parameter of of $P(\tau) \propto \tau^{-1}$ produces a less informative prior on parameters of interest than the more common uniform τ prior, $P(\tau) \propto 1$.

3.3.3 The effects of changing the observed redshift

I have so far restricted the discussion to the priors imposed at $z = 0$. It is also interesting to consider if and how these priors change as a function of observed redshift. In particular, given the conclusion in Section 3.3.1 that the priors imposed could bias results from galaxy SED fitting in favour of a tight SFMS within a narrow range in sSFR, it is interesting to consider whether the redshift evolution of the prior on sSFR matches the redshift evolution of the SFMS.

Fig. 3.3 shows the priors imposed on sSFR and mass-weighted formation time by the fiducial tau model as described in Table 3.1 at a range of observed redshifts, chosen to span the cosmic time interval between the epoch of peak star-formation and the present day. Increasing the redshift effectively changes the prior on the T_0 parameter by reducing $t_{\rm obs}$, the upper limit. This discussion is therefore simply a specific case of changing prior probability densities, as discussed more generally in Section 3.3.2.

It can be seen that the prior on mass-weighted formation time retains the same shape and retains a bias towards later formation than the $t_{\rm MW}$ values calculated from the Madau & Dickinson (2014) SFRD curve.

The sSFR priors imposed by this model can be seen to be a strong function of redshift, with the prior evolving towards sharper peaks at higher sSFRs with increasing redshift. This evolution is in the same sense as the SFMS, although the evolution is weaker than for the SFMS of Speagle et al. (2014) at fixed stellar mass.

It is informative to compare the results shown in Fig. 3.3 with those of Ciesla et al. (2017). The left hand panel of their fig. 8 shows that, even at $z = 0$, the peak of the sSFR prior for their tau model falls below the SFMS at the fiducial stellar mass of $M_* = 10^{10.5}{\rm M}_\odot$. This contrasts with the top two panels of Fig. 3.3, where the prior peaks above the SFMS. This difference is a consequence of different lower limits placed on the τ parameter: Ciesla et al. (2017) use 1 Myr, as opposed to the limit of 300 Myr used here.

I conclude that the use of standard parametric SFH models both has the potential

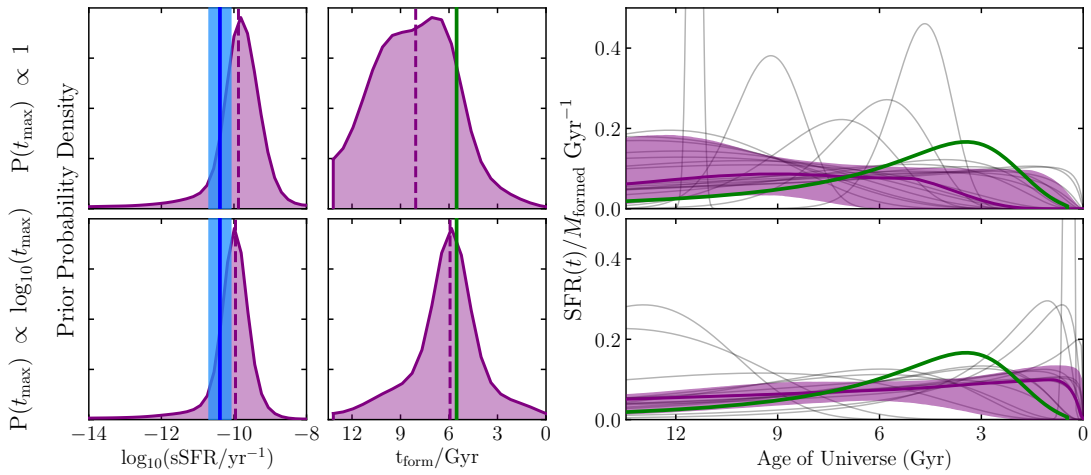

Figure 3.4 *The effects of changing the prior on the $t_{\rm max}$ parameter of the lognormal SFH model. All details are as in Figure 3.1.*

to push galaxies towards a narrow range of sSFR values, and that this range of sSFR values evolves with observed redshift in the same sense as the SFMS. Clearly, therefore, an important consideration when designing galaxy SED fitting analyses is to avoid unintentionally biasing results towards recovery of a tight SFMS by the use of a strong prior on sSFR where the observational data being fitted may not be constraining enough to reliably infer SFRs.

3.4 Effects of different priors for the lognormal and double power law models

As discussed in Section 3.3.1, the recent increase in the diversity of SFH models used in the literature precludes a side-by-side comparison of all available options. Instead I advocate that authors who use new models should present the results of tests such as those performed in Section 3.3, in order to understand the priors that are being imposed on the physical parameters of interest. As the lognormal and DPL models I consider in this thesis are relatively novel, I here further elaborate on the results of Section 3.3 for these models. The aim is to provide a discussion that helps authors wishing to use these models to select appropriate prior probability densities for model parameters.

3.4.1 The lognormal model

A number of variations might reasonably be considered on the priors reported for this model in Table 3.1. For example, Diemer et al. (2017) introduce a relatively complex set of priors to penalise extremely large $t_{\rm max}$ and $T_{\rm FWHM}$ values. As the uncertainties on $t_{\rm max}$ and $T_{\rm FWHM}$ span several orders of magnitude, it is reasonable to consider assigning them priors that are uniform in logarithmic, rather than linear space. This was tested in the case of $T_{\rm FWHM}$ and found to result in extremely narrow, bursty SFHs, which, when fitted to data, typically adopted $T_{\rm FWHM}$ values consistent with the lower limit on the prior. As this is not thought to be a physically realistic shape for the SFHs of most galaxies, this option was discounted. Instead, in Figure 3.4 I demonstrate the effects of imposing a prior of $\mathrm{P}(t_{\rm max}) \propto \log_{10}(t_{\rm max})$ between the limits quoted in Table 3.1, as opposed to the original uniform prior shown on Fig. 3.1. This can be seen to significantly change the shape prior, encoding a preference for earlier formation, which brings the prior on $t_{\rm form}$ into better agreement with Madau & Dickinson (2014). However, this change also significantly narrows the priors on SFH shape and $t_{\rm form}$, reducing the ability of this prior to describe the diversity of possible SFH shapes.

3.4.2 The double power law model

The priors for the DPL model listed in Table 3.1 were arrived at following extensive experimentation with different options, detailed in Chapter 2. The shape of the SFH changes roughly uniformly with the logarithm of α and β between the limits of 0.1 and 1000. Outside of this range, further increase or decrease of these parameters results in no further change to the SFH shape, as the slopes are essentially horizontal for values < 0.1, or vertical for values > 1000. Using uniform priors on α and β leads to the same effect as was discussed for the $T_{\rm FWHM}$ parameter of the lognormal model in Section 3.4.1, with extremely bursty SFHs being assigned the bulk of the prior mass.

Given the success achieved in emphasising earlier formation in Fig. 3.4, I perform a similar test to Section 3.4.1 by changing the prior on the τ parameter from the initial uniform prior to $\mathrm{P}(\tau) \propto \log_{10}(\tau)$. The results of this test are shown in the middle row of Fig. 3.5. It can be seen that this change overweights models that form the majority of their stars at very early times, producing a prior preference

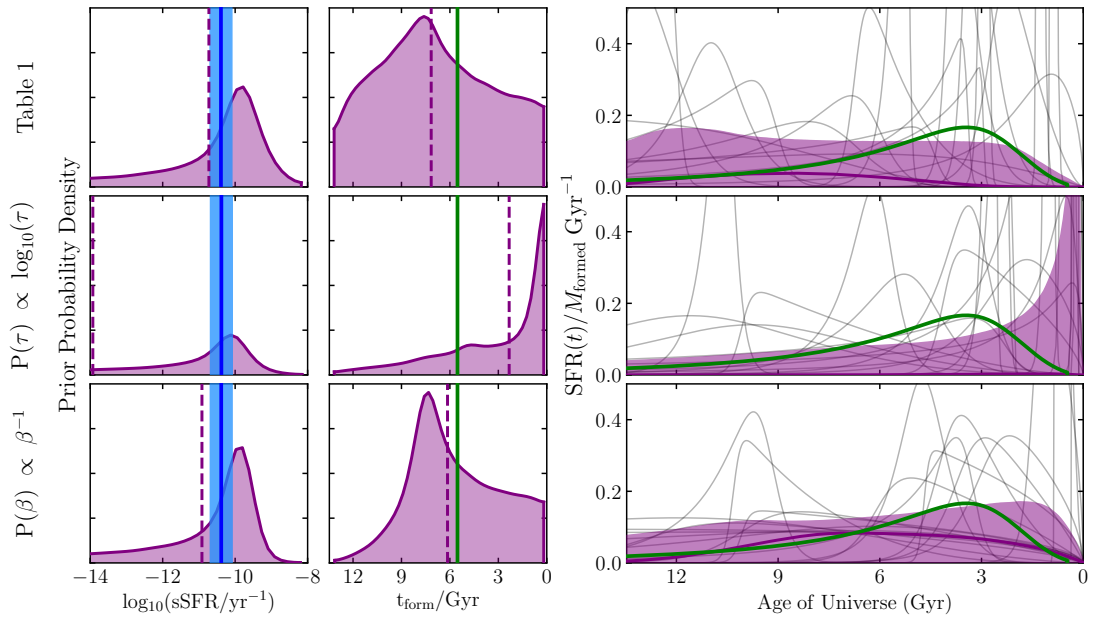

Figure 3.5 *The effects of changing the priors on the τ and β parameters of the DPL SFH model. All details are as in Figure 3.1.*

for older stellar populations than the Madau & Dickinson (2014) curve. I therefore suggest retaining a uniform prior on τ for the DPL model.

The main advantage of the DPL model is the ability to decouple the early and late time SFHs, with separate parameters controlling the slopes of each. With this in mind, and considering the results of Fig. 3.1, I assess the possibility of changing the prior on the rising slope from $P(\beta) \propto \log_{10}(\beta)$ to $P(\beta) \propto 1/\beta$. This change lends more prior weight to flatter rising slopes, thus causing star-formation to be more extended back towards earlier times before it peaks. These results are shown in the bottom row of Fig. 3.5. This is probably the most promising of the alternatives I have considered to the priors listed in Table 3.1, with the average of the $t_{\rm form}$ prior in good agreement with Madau & Dickinson (2014). However the average SFH shape can still be seen to be quite different from Madau & Dickinson (2014), with the sSFR prior also retaining a fairly strong peak.

As discussed in Section 3.5.4 and demonstrated in this section, the process of tuning a parametric SFH model involves a large amount of trial and error, with no clear physical link between the priors chosen and the physics the SFH model represents. I therefore suggest that non-parametric SFH models are a better choice, as they allow physical information to be included in the prior in a more direct way. For further information and examples see Leja et al. (2018).

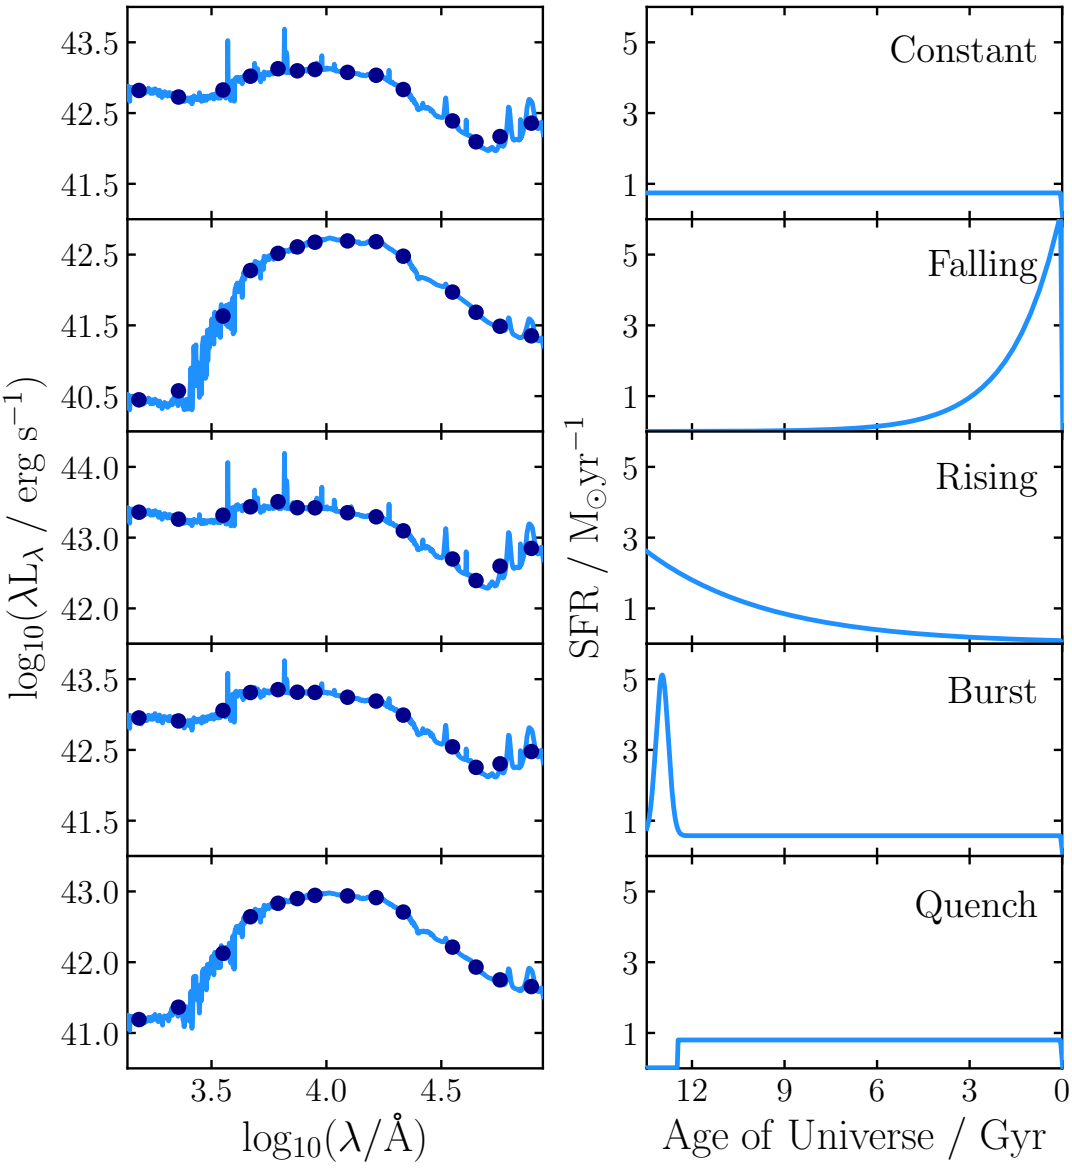

Figure 3.6 *SEDs (left) and SFHs (right) for each of the mock galaxies introduced in Section 3.5.1. Photometric fluxes derived from these spectra are shown as circles on the left for each of the following filters: GALEX FUV/NUV, SDSS ugriz, 2MASS JHKs and Spitzer/IRAC channels* 1 – 4.

3.5 Testing parametric models with mock observations

In Section 3.3 I have discussed how, in the absence of strongly constraining data, the priors imposed by the SFH model on physical parameters can affect the results obtained by assigning greater prior weights to SFHs of certain shapes. However, even in the high-SNR regime, the limited range of shapes that a given parametric SFH model is capable of reproducing can introduce biases into physical parameter estimates. To understand these biases it is necessary to examine some test cases involving high-SNR data.

In this section I test the abilities of the parametric SFH models introduced in Section 3.2 to recover galaxy physical parameters from mock high-SNR broad-band photometric data. In Section 3.5.1 I describe the mock catalogue, and the generation and fitting of mock photometry with Bagpipes. In Section 3.5.2 I discuss the biases introduced by the use of parametric SFH models to fit these mock data. In Section 3.5.3 I compare the quality of the fits obtained using different parametric SFH models. Finally, in Section 3.5.4 I consider the implications of these results for the use of broad-band photometric observations as a tool for learning about galaxy SFHs.

3.5.1 Generating and fitting mock data

To facilitate these tests, a set of five mock galaxies at $z = 0$ with different SFHs was generated. The SFHs are:

- **Constant**: equal SFR from $t = 0$ to $t_{\rm obs}$.
- **Falling**: exponential decline (see Equation 3.1) with $T_0 = 0$ and $\tau = t_{\rm obs}/10 = 1.4$ Gyr.
- **Rising**: $\mathrm{SFR}(t) \propto e^{\frac{t}{\tau}}$ with $\tau = t_{\rm obs}/4 = 3.4$ Gyr.
- **Recent burst**: constant SFR from $t = 0$ to $t_{\rm obs}$ making up 80% of $M_{\rm formed}$ and a Gaussian burst centred 500 Myr before $t_{\rm obs}$ with width $\sigma = 200$ Myr making up the other 20% of $M_{\rm formed}$.

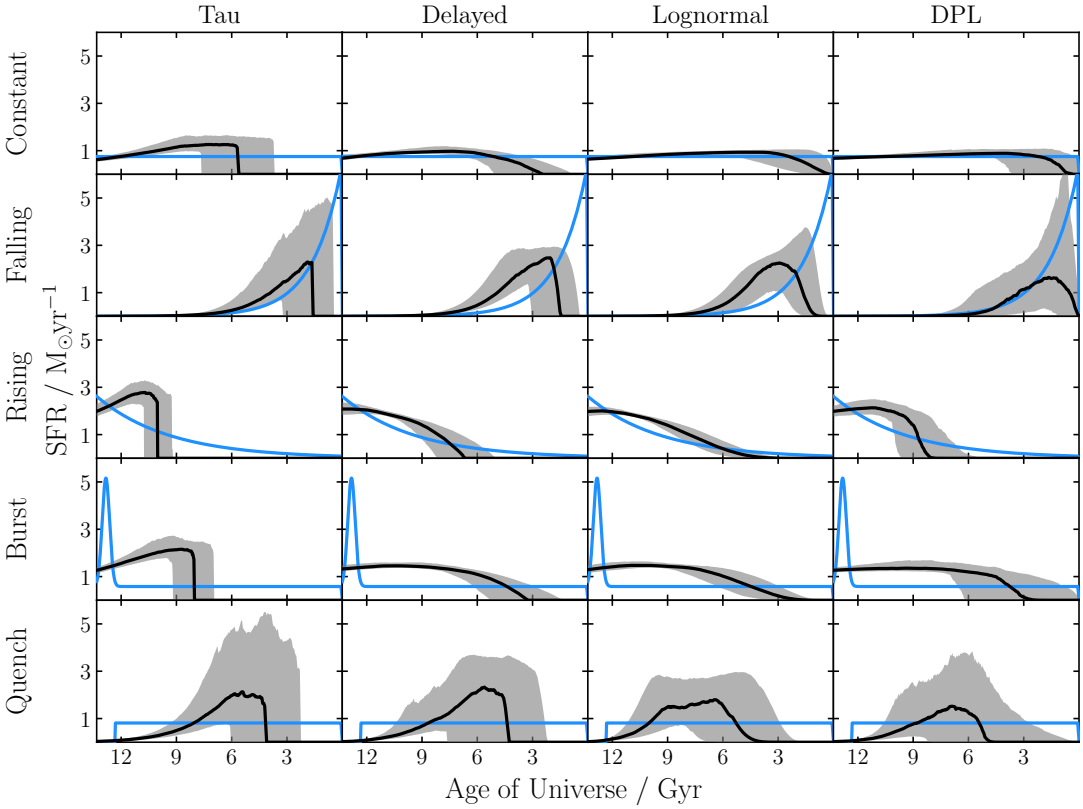

Figure 3.7 *Recovery of SFHs from the mock data described in Section 3.5.1 using the four parametric SFH models described in Section 3.2 and Table 3.1. Different rows show different mock galaxies, different columns show different fitted SFH models. Blue lines are the input mock SFHs, black lines show the posterior medians and the shaded regions show the 16^{th}–84^{th} percentiles.*

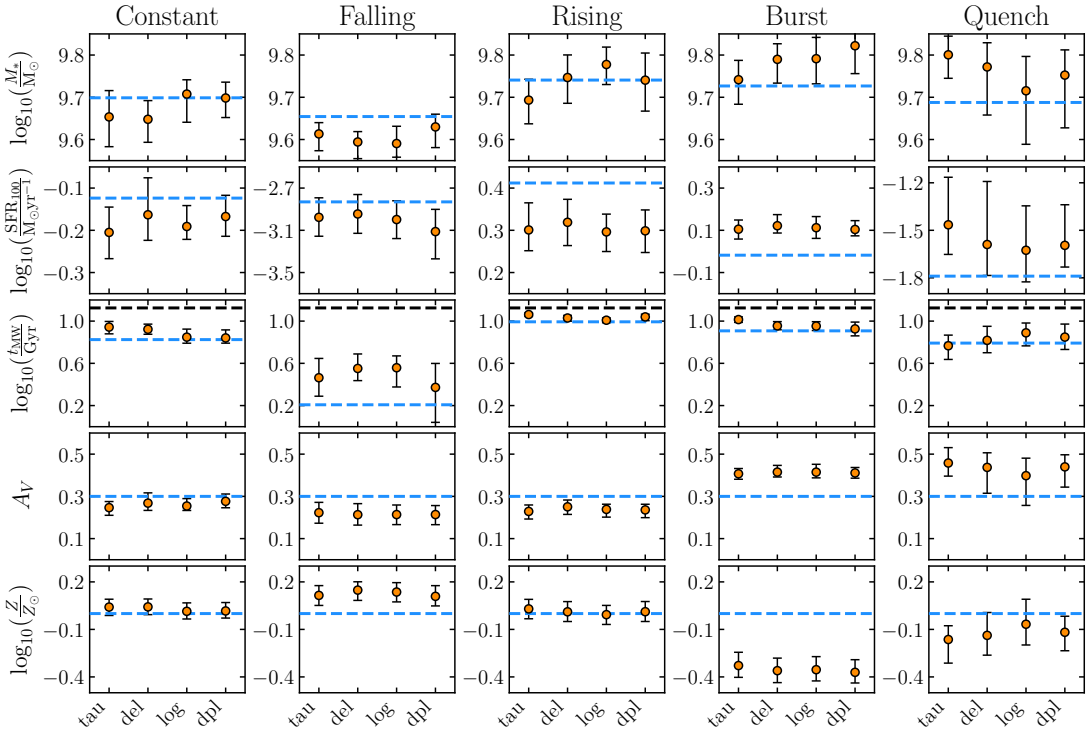

Figure 3.8 *Physical parameter recovery from mock data (see Section 3.5.1) using the four SFH models introduced in Section 3.2 and Table 3.1. Posterior median values are shown as orange circles, while the errorbars show the 16^{th}–84^{th} percentiles of the posteriors. Blue dashed lines show the true values, black dashed lines show $t_{\rm obs}$. The observed biases result from the parametric SFH models being unable to reproduce the correct SFH shapes (see Fig. 3.7) and assigning different prior weights to different SFH shapes.*

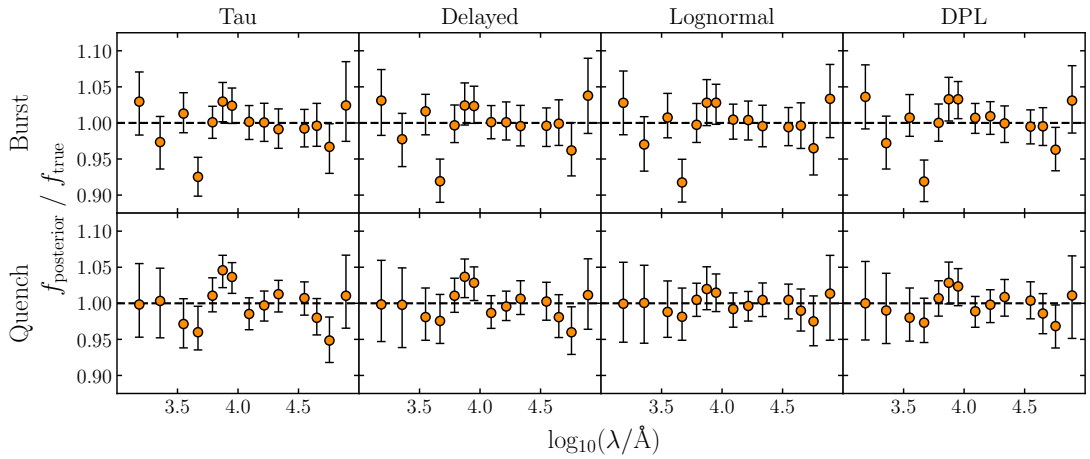

Figure 3.9 *Posterior predictions for photometric observations compared to the input values. The two rows correspond to the recent burst and rapid quench mocks (results for the other three mocks are very similar). The four columns correspond to the four parametric SFH models. Posterior medians are shown as orange circles, error bars show the 16^{th}–84^{th} posterior percentiles.*

- **Sudden quench**: constant star formation from $t = 0$ to $t_{\rm obs} - 1$ Gyr, then constant star formation at 2% of the original level from that point to $t_{\rm obs}$.

Mock photometry was generated with Bagpipes using the methods described in Chapter 2 for each of the following photometric filters: *GALEX FUV/NUV*, SDSS *ugriz*, 2MASS *JHKs* and *Spitzer*/IRAC channels $1-4$. I assign each photometric flux an uncertainty corresponding to a high SNR of 25. However, I do not perturb the model fluxes by these uncertainties, so as to isolate the effects of the SFH parameterisation. The mock SEDs and SFHs are shown in Fig. 3.6.

Each model was assigned a total stellar mass formed, $M_{\rm formed} = 10^{10}$ M$_\odot$ and a metallicity of $Z = 0.02$ (which I take to be Solar metallicity, $Z_\odot$). An ionization parameter of $\log_{10}(U) = -3$ was assumed for nebular emission. Dust attenuation with $A_V = 0.3$ mag was applied using the Calzetti et al. (2000) attenuation curve. Attenuation was assumed to be doubled for stars formed in the last 10 Myr and for nebular emission. For this project, Bagpipes was upgraded to use the dust emission models of Draine & Li (2007). I assume values of 2 for $Q_{\rm pah}$, the percentage of dust mass in polycyclic aromatic hydrocarbons, 1 for $U_{\rm min}$, the minimum starlight intensity to which the dust is exposed, and 0.01 for $\gamma_{\rm e}$, the fraction of the incident starlight at $U_{\rm min}$.

This mock catalogue was chosen to span a wide range of scenarios for the

formation of galaxies. However it should be noted that, by constructing this catalogue, I am expressing prior beliefs about the SFHs of real galaxies. The results of Sections 3.5.2 and 3.5.3 are necessarily dependent on these prior beliefs (see Section 3.5.4).

Each of the four parametric models introduced in Section 3.2 was fitted to each set of mock photometry assuming the priors given in Table 3.1 at a fixed redshift of $z = 0$. The other free parameters are A_V for the Calzetti et al. (2000) dust attenuation law, to which I assign a uniform prior between 0 and 4 mag, metallicity, Z, to which I assign a logarithmic prior over the range $0.2 < Z/Z_\odot < 2$, and the dust emission parameters, to which I assign the same priors as Leja et al. (2017), except that I do not extrapolate $Q_{\rm pah}$ beyond a maximum value of 4.58.

3.5.2 Recovery of SFHs and physical parameters

Fig. 3.7 shows the posterior SFHs obtained by the process described in Section 3.5.1 compared to the input mock SFHs. Fig. 3.8 shows the posterior constraints obtained for a number of physical parameters compared to their true values. It can be seen that the quality of the recovered posterior SFHs and physical parameters is a strong function of the shape of the input mock SFH and a weaker, but still significant, function of the parametric SFH model that is fitted. This finding is supported by Fig. 3.8, where it can be seen that often all four models return consistent values for physical parameters that are in strong tension with the true value.

The SFHs of the first three "simple" mocks (constant, falling, rising) are relatively well recovered by all of the SFH models, with the exception of the tau model, which predictably struggles to recover the constant and rising SFHs. It is also interesting to highlight the case of the tau model fitting the falling mock, as in this case the true SFH exists within the prior. It can be seen from Fig. 3.7 that, whilst the tau model recovers the shape of the input SFH well, the recovery is not perfect, as might be expected. This is due to the non-Gaussian shape of the posterior distribution, in this case for the T_0 and τ parameters.

This is a common consequence of poorly constraining data, and means that the most probable (maximum a posteriori) and posterior median parameter estimates may be offset. In this case, because the model fluxes have not been perturbed, the most probable model is the input model. However, when dealing with

real observational uncertainties and non-Gaussian posteriors, the maximum a posteriori parameter estimates are often poorly representative of the posterior.

For the "simple" mocks, the stellar mass, SFR, dust attenuation and metallicity posteriors can be seen from Fig. 3.8 to typically fall within $1-2\sigma$ of the true values, with biases of the order of ~ 0.1 dex for stellar mass and ~ 0.2 dex for SFR. The posterior mass-weighted formation times are more strongly in tension with the input values (up to $\sim 4\sigma$), and are typically overestimates of the true values. These biases are due to all four models, to some extent, failing to reproduce the early-time evolution of these mock SFHs.

By contrast, all four SFH models struggle to reproduce the more complex shapes of the recent burst and sudden quench models. More significant biases in the recovery of physical parameters can be seen in these cases, up to 0.3 dex in SFR, 0.2 mag in dust attenuation, and 0.3 dex in metallicity. For the burst mock in particular, the posterior probability distributions strongly exclude the true values for all parameters except $\boldsymbol{M_*}$. The posterior SFHs shown in Fig. 3.7 appear unable to simultaneously reproduce both the overall shapes of these mock SFHs and their rapid variability at late times, with the shapes fitted representing a compromise between the two. The strong biases in the recovered dust and metallicity values are a consequence of the age-metallicity-dust degeneracy, which allows the mock photometry to be well matched even whilst fitting radically different SFHs.

Based upon these results, I conclude that the choice of parametric SFH model has the potential to significantly affect the physical parameter estimates obtained when fitting high-SNR broad-band photometric observations, in particular by at least 0.1, 0.3 and 0.2 dex for stellar mass, SFR, mass-weighted formation time respectively. The finding of ~ 0.1 dex variations in stellar-mass measurements is in good agreement with similar analyses in the literature (e.g. Pacifici et al. 2015; Mobasher et al. 2015; Iyer & Gawiser 2017). For SFRs however, both Pacifici et al. (2015) and Iyer & Gawiser (2017) find a wider range of variations, suggesting this limited library of mocks does not encompass the most pathological cases.

However, often all four models return consistent posterior estimates for physical parameters that are strongly biased from the input values. This suggests that the true shape of the SFH is the most important factor in whether parametric models can return unbiased physical parameter values, with even a simple mock catalogue containing SFHs that are not well described by any of the parametric models. The parametric models are least reliable for galaxies that experience

recent, rapid changes in their SFRs. This limitation can be improved upon by the use of non-parametric models (Leja et al. 2018).

3.5.3 Comparisons between SFH models

Given the biases discussed in Section 3.5.2, it is interesting to consider whether there is a case, based on the results of this section, for using one of these parametric models in preference to others. In the ideal case, I would be able to identify one parametric model that reliably returns smaller biases in physical parameter estimates than the others and recommend this for general use. Failing this, I would hope to be able to identify the model that returns the least biased physical parameter estimates for a specific object (without knowledge of the true values) by assessing the relative quality of the fits. It should again be stressed, as in Section 3.5.1, that any conclusions as to which parametric SFH models are more preferred depend on the prior beliefs about real galaxy SFHs I expressed when building the mock galaxy catalogue used in this section.

Considering the relative biases returned by the use of different models, as discussed in Section 3.5.2, often the different SFH models return consistent parameter estimates, all of which are biased with respect to the input values. In cases where the relative biases differ, the tau model generally returns more highly biased physical parameter estimates, however there is little distinction between the biases obtained when using the other three parametric models. I therefore conclude that none of the delayed, lognormal or DPL models can be said to reliably produce less-biased physical parameter estimates than any other when fitting high-SNR photometric data.

I therefore move on to consider which models produce the best fits to the mock data. Fig. 3.9 shows the residuals between the posterior predictions for photometry and the input mock data. Results are shown for each of the four SFH models fitted to the recent burst and sudden quench mocks (results for the other three mocks are very similar). It can be seen that all of the models produce posterior distributions that appear to be acceptable fits to the mock photometry. Furthermore, it is impossible to distinguish by eye between the posteriors obtained by fitting each of the different models.

In order to quantify this I consider the Bayesian evidence for each model fitted to each set of mock photometry. Evidence is analogous in a Bayesian framework

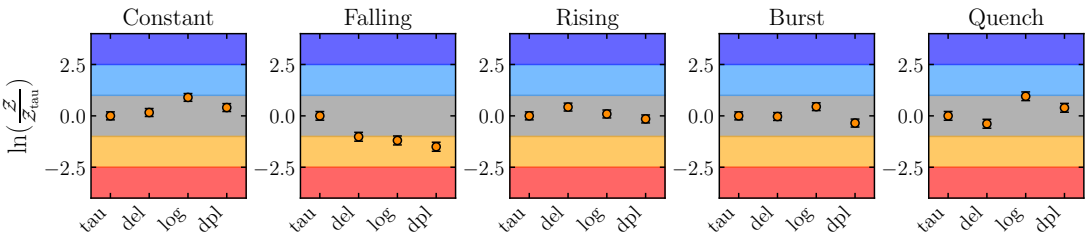

Figure 3.10 *Bayesian evidence plots (e.g. Trotta 2008) for each SFH model fitted to each mock, relative to the tau model. The coloured stripes denote (from top to bottom): a strong preference for (dark blue), a weak preference for (light blue), no preference for (grey), a weak preference against (orange) and a strong preference against (red) the model compared to the tau model. It can be seen that for none of these mocks is there a strong preference for any one of the SFH models above the others.*

to the more commonly used minimum reduced chi-squared value, and is often used to discriminate between models (e.g. Trotta 2008; Salmon et al. 2016). Fig. 3.10 shows the evidence for each model fitted to each set of mock photometry relative to the tau model. The coloured stripes represent (from top to bottom) strong evidence for, weak evidence for, no evidence for, weak evidence against and strong evidence against one model compared to another.

It can be seen from Fig. 3.10 that none of the parametric models is strongly favoured or disfavoured compared to the tau model for any of the mocks. Even for the falling mock, where the input model is within the tau-model prior, only a very weak preference is visible. This means that it is not possible, in general, to identify which SFH model produces the least biased physical parameter estimates for individual objects by assessing the "goodness of fit" to broad-band photometry. A specific case was recently demonstrated by Belli et al. (2019), who show that it is not possible to distinguish between parametric SFH models for a sample of quiescent galaxies.

3.5.4 Broad-band photometry as a tool for understanding galaxy SFHs

In Section 3.5.3 I have demonstrated that, even though I observe from Fig. 3.8 that the tau model produces more highly biased physical parameter estimates for the mock catalogue, this is not associated with a worse quality of fit. This means that, for the case in which we have no knowledge of the true physical

parameter values, we have no basis for deciding which parametric SFH model produces less-biased physical parameter estimates.

The conclusion that the tau model produces more highly biased physical parameter estimates in Section 3.5.3 is a consequence of the prior beliefs about real galaxy SFHs expressed when constructing the mock catalogue in Section 3.5.1. If all of the mocks in the catalogue were similar to the falling mock, we would conclude that the tau model describes galaxy SFHs equally as well or better than the other three models.

This leads to an important conclusion: high-SNR broad-band photometry cannot, in general, constrain prior beliefs about which parametric SFH models are most appropriate for describing the SFHs of real galaxies. Hypotheses such as galaxy SFHs being well described by the lognormal function (Gladders et al. 2013) cannot be proven or disproven using broad-band photometric observations. The physical parameter inferences made are therefore necessarily dependent on prior beliefs, to at least the levels reported in Section 3.5.2.

There are three possible responses to this conclusion. Firstly, one can perform a more sophisticated version of the analysis presented in Sections 3.5.1 and 3.5.2, by generating a catalogue of mock data that represents a set of prior beliefs about galaxy formation (e.g. Buat et al. 2014; Ciesla et al. 2015), then fitting those mocks with different parametric SFH models to decide which is the most appropriate for representing those prior beliefs.

A popular choice is to use catalogues of mock observations drawn from simulations of galaxy formation, as in Chapter 2 (see also Pacifici et al. 2015; Diemer et al. 2017). However, the lack of flexibility in parametric SFH models means that the process of selecting an appropriate model usually involves a significant amount of trial and error, with no clear physical link between the chosen model and the prior beliefs expressed, as discussed in Section 3.4. Additionally, the use of prior beliefs drawn from simulations renders comparisons between observational results and simulation outputs of questionable value.

This brings us to the second option: the use of non-parametric SFH models. Because of the increased flexibility of these models, it is far easier to encode more specific prior beliefs about the shapes of galaxy SFHs into the model, with or without the use of mock observations drawn from simulations. This option is the subject of a companion paper (Leja et al. 2018), and is further discussed in Section 3.7.

Table 3.2 *A comparison of $z \sim 0.05$ SFRD and SMD estimates. The first four rows show estimates using the different SFH models used in this chapter to fit the GAMA sample (see Section 3.6.1). I also show several results from the literature, converted to the Kroupa (2001) IMF where necessary. The uncertainties quoted do not include systematic effects.*

Model	$\log_{10}\left(\frac{\text{SMD}}{\text{M}_\odot\ \text{Mpc}^{-3}}\right)$	$\log_{10}\left(\frac{\text{SFRD}}{\text{M}_\odot\ \text{yr}^{-1}\ \text{Mpc}^{-3}}\right)$
Tau	$8.40^{+0.02}_{-0.03}$	$-2.041^{+0.002}_{-0.001}$
Delayed	$8.39^{+0.02}_{-0.02}$	$-2.033^{+0.001}_{-0.001}$
Lognormal	$8.37^{+0.02}_{-0.02}$	$-2.082^{+0.002}_{-0.001}$
DPL	$8.38^{+0.03}_{-0.03}$	$-2.088^{+0.001}_{-0.002}$
Madau & Dickinson (2014)	8.55	-1.97
Wright et al. (2017)	$8.35^{+0.06}_{-0.07}$	–
Driver et al. (2018)	$8.30^{+0.01}_{-0.01}$	$-1.95^{+0.00}_{-0.00}$

The final option is the analysis of higher quality observational data, such as high-SNR continuum spectroscopy, which has been shown to be more strongly constraining on galaxy SFHs (e.g. Gallazzi et al. 2005, 2008; Ocvirk et al. 2006; Pacifici et al. 2012; Thomas et al. 2017). These kinds of analyses are extremely promising, as they have the potential to demonstrate a clear preference in favour of one model for the SFHs of galaxies over others, providing deeper insights into the physics driving the assembly of stellar mass in galaxies.

3.6 Testing parametric models with observational data

In Sections 3.3 and 3.5 I have demonstrated the effects of different parametric SFH models on results obtained through SED fitting for galaxy stellar masses, SFRs and mass-weighted formation times. As noted in Section 3.1, all measurements of these quantities are dependent, to some extent, on the SFH model. However, it is possible to minimise the impact of the SFH by looking at SFR indicators that are sensitive to star-formation on very short timescales, over which it is safe

to assume that star-formation is constant (e.g. Kennicutt & Evans 2012). By making these observations for representative samples of galaxies across cosmic time, it is possible to infer the redshift evolution of the cosmic SFRD and hence stellar-mass density (SMD) independently of individual galaxy SFHs (e.g. Madau & Dickinson 2014).

In this section, I fit a low-redshift, volume-complete sample of galaxies using the four parametric SFH models discussed in this chapter. I then consider whether the results obtained through SED fitting for the redshift evolution of the cosmic SFRD and SMD are consistent with these "SFH-free" estimates. This provides a valuable additional perspective on the physical motivation for the priors imposed by the use of parametric SFH models on stellar masses, SFRs and mass-weighted formation times (e.g. Heavens et al. 2004; Ocvirk et al. 2006; Gallazzi et al. 2008; Wuyts et al. 2011). I will also compare these results to similar SED fitting analyses in the literature as a check on my Bagpipes fitting methodology.

In Section 3.6.1 I discuss the fitting of the GAMA observational sample. In Section 3.6.2 I consider my stellar-mass measurements by calculating the inferred cosmic SMD. In Section 3.6.3 I consider my SFR measurements by calculating the inferred cosmic SFRD and comparing to independent, Hα-derived SFRs. Finally, in Section 3.6.4 I consider my mass-weighted formation time measurements by considering the shape of the inferred SFRD evolution.

3.6.1 GAMA data and fitting methodology

In this section, I use data from the Galaxy and Mass Assembly Survey (GAMA; Driver et al. 2009, 2016; Baldry et al. 2018) DR3 release. In particular, the 21-band aperture-matched catalogue for the equatorial G09, G12 and G15 fields generated using the Lambdar code by Wright et al. (2016). The catalogue includes far-UV to far-IR photometry from *GALEX*, SDSS, VISTA, *WISE* and *Herschel* over a 180 deg^2 area.

I define my sample by using the stellar mass estimates of Taylor et al. (2011) to select objects with $M_* > 10^9$ M$_\odot$. GAMA is mass-complete down to 10^9 M$_\odot$ at $z \lesssim 0.08$ (Lange et al. 2015; Driver et al. 2016), and I therefore select objects with GAMA spectroscopic redshifts in the range $0.05 < z < 0.08$. This results in a sample of 6134 objects. Whilst this is, in reality, a mass-complete rather than volume-complete sample, $\sim 95\%$ of the stellar mass in the Universe at $z = 0$ is

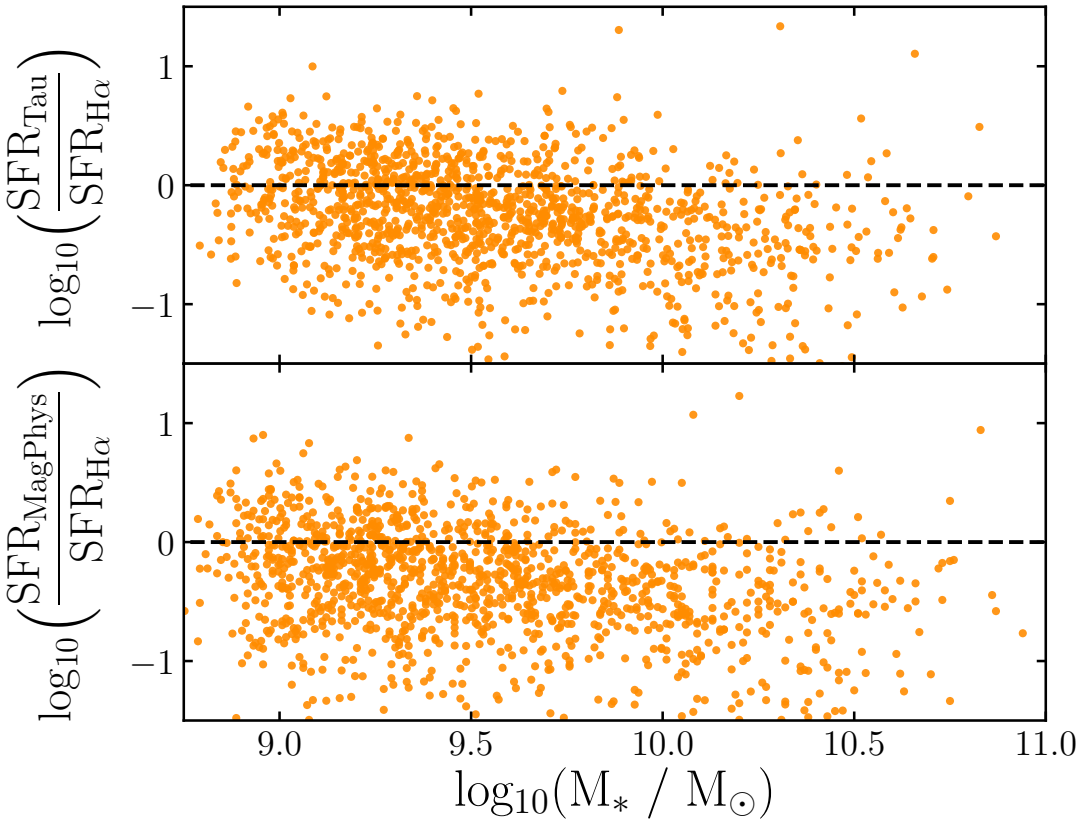

Figure 3.11 *Comparisons between SFRs derived from SED fitting and from dust-corrected Hα fluxes using the Kennicutt & Evans (2012) relation. The top panel shows SFRs obtained in this chapter using the tau SFH model; the other models produce similar results. The bottom panel shows SFR estimates from* MagPhys*, released as part of GAMA DR3.*

in galaxies above this mass limit (Tomczak et al. 2014). I therefore approximate this sample as volume-complete.

Each object is fitted with Bagpipes using each of the four parametric SFH models described in Section 3.2, under the same assumptions as described in Section 3.5.1. To calculate the inferred cosmic SFRD evolution, for each SFH model I extract 100 posterior draws for the SFH of each object. These are then summed across objects and divided by the comoving volume from which the sample is drawn. I use these curves to calculate the inferred SMD and SFRD at $z \sim 0.05$ for each SFH model. These values are reported in Table 3.2.

3.6.2 Inferred stellar masses

I first consider the stellar masses inferred for this sample. I evaluate these stellar masses by comparing the cosmic SMD results for each model, shown in Table 3.2, to results from the literature. My results can be seen to be consistent with those of Wright et al. (2017), derived from the same catalogue using the MagPhys code (da Cunha et al. 2008). A similar analysis by Driver et al. (2018) yields a slightly lower value, however all of these results are consistent to well within the ~ 0.3 dex systematic uncertainties normally assumed for stellar mass measurements (e.g. Mobasher et al. 2015). I thus conclude that my stellar mass estimates are consistent with similar analyses in the literature.

A "SFH-free" estimate of the cosmic SMD can also be obtained by integrating the Madau & Dickinson (2014) SFRD curve, then multiplying by $(1 - R)$, where $R = 0.27$ is the mass-return fraction. A ~ 0.2 dex offset has been widely observed between this estimate and those from SED fitting analyses (e.g. Leja et al. 2015). My results also reflect this tension, and although this offset is within the systematic uncertainties, the question remains as to which systematic effect is responsible. The results of Section 3.5.2 suggest that the use of parametric SFH models can lead to systematic offsets in stellar-mass measurements at levels of ~ 0.1 dex, with the main driver of these offsets being the shape of the true SFH. It is possible therefore that the majority of this offset is due to the biasing effects of parametric SFH models. However this would only be true in the scenario that all galaxy SFHs have true shapes that cause parametric models to underestimate their stellar masses (e.g. similar to the falling mock in Section 3.5).

3.6.3 Inferred star-formation rates

I now consider the star-formation rates I infer for the sample. As can be seen from Table 3.2, my SFRD results at $z \sim 0.05$ fall ~ 0.1 dex lower than both the SED fitting analysis of Driver et al. (2018) and the SFRD curve of Madau & Dickinson (2014). However the offset is, again, well within the systematic uncertainties of ~ 0.5 dex normally assumed for SFR measurements (e.g. Pacifici et al. 2015). It is interesting to note that, whereas the SMD measurements reported in Table 3.2 using different parametric models are consistent, the SFRD measurements are strongly inconsistent with each other. This suggests that SFR measurements are more strongly biased by SFH priors than stellar mass measurements.

It is also possible to assess the quality of the SFRs I derive on an individual basis by comparing my SFRs to those inferred from Hα fluxes. GAMA DR3 includes calibrated measurements of Hα and Hβ fluxes from GAMA and Sloan Digital Sky Survey (SDSS) spectra. I begin by selecting the 1373 objects of the 6134 in the sample that have SNR $>$ 5 in both Hα and Hβ. I then correct the Hα fluxes for dust attenuation using the measured Balmer decrements, following the process outlined in section 3 of Domínguez et al. (2013). I finally convert the dust-corrected Hα fluxes to SFRs using the calibration of Kennicutt & Evans (2012).

Fig. 3.11 shows my tau-model SFRs compared to those derived from Hα as a function of the inferred stellar masses. SED-derived SFRs, calculated using MagPhys, were also released as part of GAMA DR3. These are also shown compared to Hα on Fig. 3.11. It can be seen that both sets of results agree well with Hα at lower masses, however SFRs at progressively higher masses are increasingly under-predicted with respect to Hα.

It should be noted that the scatter observed in Fig. 3.11 is partially due to variations in galaxy SFRs on the very short timescales to which Hα is sensitive. The parametric models I consider are not capable of reproducing variations on such short timescales, however several approaches have been demonstrated that allow parametric models to reproduce this behaviour, such as the addition of bursts of star-formation (e.g. da Cunha et al. 2008), and resampling the average SFR over the last 10 Myr from a separate distribution (e.g. Pacifici et al. 2016).

The change observed in the mean SFR offset with stellar mass is consistent with the result from Section 3.5.2 that biases in SFRs inferred using parametric SFH models are a strong function of the true SFH shape. Lower-mass galaxies are known to form their stars later in cosmic history, and are likely to have SFHs consistent with the constant and/or rising mocks, for which I recover either unbiased or slightly underestimated SFRs.

Conversely, the significant underestimation of SFRs with respect to Hα at higher masses in Fig. 3.11 is not consistent with the SFR offsets observed for any of the mocks in Figure 3.8. As in Section 3.5.2, this again suggests that this mock catalogue does not encompass the most pathological cases. However it should also be noted that the dustier nature of higher mass galaxies (e.g. Garn & Best 2010) makes Hα a less-reliable SFR indicator.

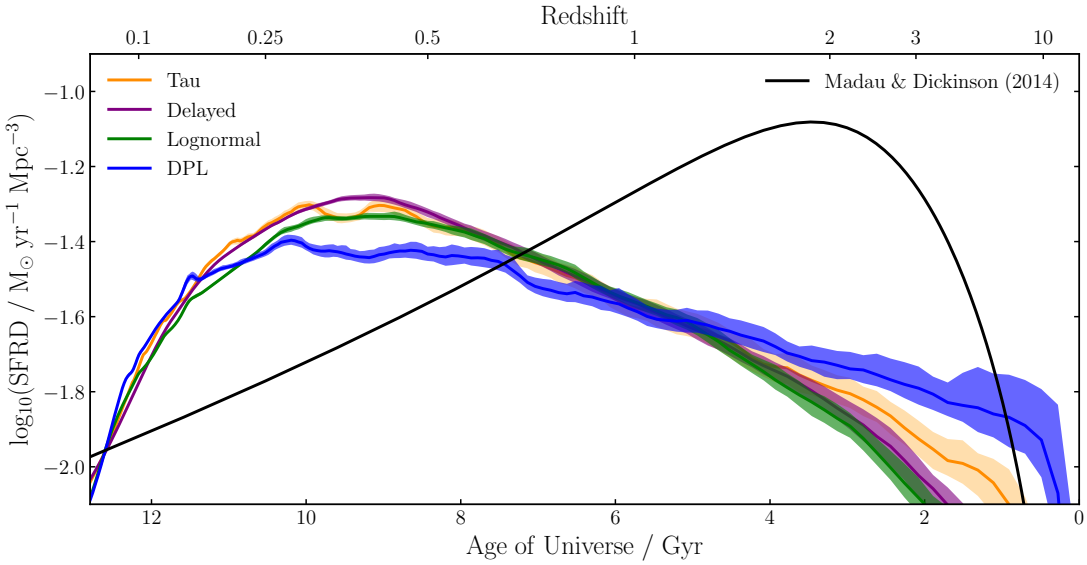

Figure 3.12 *Redshift evolution of the cosmic SFRD as derived from the GAMA sample under the assumption of different parametric SFH models. The solid lines show the posterior medians and the shaded regions show the 16^{th}–84^{th} percentiles. The result obtained by Madau & Dickinson (2014) by measuring the SFRs of galaxies across cosmic time is shown in black. The shape of the SFRD curve can be seen to be poorly reproduced, with mass assembly occurring later in cosmic history.*

3.6.4 Inferred mass-weighted formation times: the shape of the inferred SFRD evolution

I finally consider the mass-weighted formation times inferred for the GAMA sample, by considering the implied redshift evolution of the cosmic SFRD. These results are shown in Fig. 3.12, along with the Madau & Dickinson (2014) result. Whilst I have demonstrated in Sections 3.6.2 and 3.6.3 that the stellar masses and SFRs I infer for the GAMA sample at $z \sim 0.05$ are broadly consistent with Madau & Dickinson (2014), it can be seen that the cosmic SFRD evolution inferred is very different. The four SFRD curves peak significantly later in cosmic history, at $z \sim 0.4$, when the Universe is $9-10$ Gyr old. This is in marked contrast to the Madau & Dickinson (2014) curve, for which the SFRD peaks at $z \sim 2$, when the Universe is $2-3$ Gyr old.

Measuring SFRs for individual high-redshift galaxies is a more direct measurement of the redshift evolution of the cosmic SFRD, and can be regarded, for the purposes of this discussion, as ground truth. Fig. 3.12 therefore suggests that the four parametric models considered significantly overestimate the mass-weighted

formation times for galaxies in the GAMA sample, by as much as ~ 5 Gyr on average.

A similar overestimation of mass-weighted formation times (underestimation of mass-weighted ages) has been previously observed in SED fitting analyses using both parametric and non-parametric SFHs (e.g. Gallazzi et al. 2008, Wuyts et al. 2011, Leitner 2012). An analysis by Panter et al. (2003) and Heavens et al. (2004) extracted SFHs for almost 100,000 galaxies from SDSS spectra using the MOPED code, which employs an 11-parameter non-parametric SFH model. They found that the cosmic SFRD peaks at $z \sim 0.5$, when the Universe was $8 - 9$ Gyr old, similar to the results obtained with the parametric SFH models considered in this chapter. However, an updated analysis by Panter et al. (2007) found that star-formation peaked in their highest-redshift bin at $z > 1.87$, demonstrating the significant impact of SED-modelling assumptions on mass-weighted formation time inferences.

In Fig. 3.1, I showed that the priors imposed by the parametric SFH models favour later galaxy formation than Madau & Dickinson (2014). In fact, the prior medians shown on the centre-left panels of Fig. 3.1 correspond closely to the peaks in cosmic SFRD shown in Fig. 3.12. It seems likely, therefore, that the strong posterior constraints obtained for cosmic SFRD evolution, favouring a late peak in cosmic SFRD, are a consequence of the strong priors imposed by the parametric SFH models, rather than the photometric data fitted. This conclusion is supported by the fact that the DPL model exhibits both a prior preference for older ages than the other models in Fig. 3.1, and earlier mass assembly in Fig. 3.12.

A key goal for contemporary SED fitting analyses is to move beyond the acquisition of galaxy physical parameters at the redshift of observation to reliably infer mass-assembly histories at earlier times. This ability would be a valuable tool for explicitly linking galaxy populations at different epochs (e.g. McLure et al. 2018a; Pentericci et al. 2018). Whilst it has long been known that tau models do not produce accurate mass-assembly histories (e.g. Wuyts et al. 2011; Reddy et al. 2012; Pforr et al. 2012; Buat et al. 2014), these results suggest that newer parameterisations, such as the lognormal and DPL models, do not significantly improve our ability to obtain realistic mass-assembly histories from the observed SED. Whilst it is possible to calibrate these models to obtain unbiased results in certain circumstances (e.g. see Chapter 2), in order to obtain realistic mass-assembly histories for representative samples of galaxies, further consideration

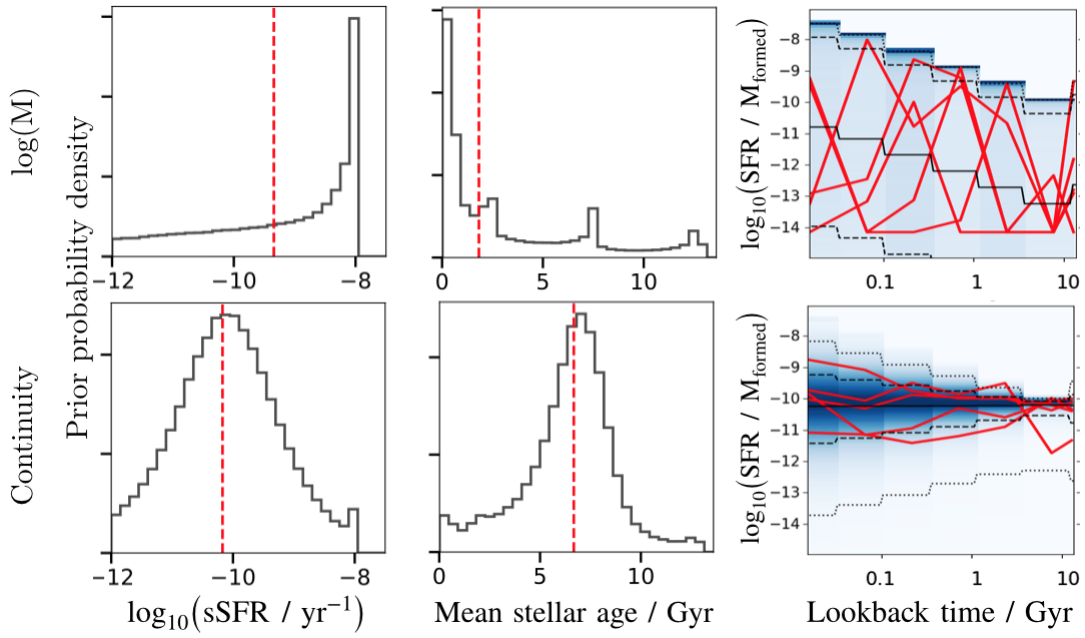

Figure 3.13 *The priors imposed on sSFR, mean stellar age and SFH shape by the two non-parametric models discussed in Section 3.7. In the left and centre panels, the dashed red lines indicate the prior medians. In the right panels, the red lines show individual draws from the prior, whilst the blue shading indicates the prior density. The solid black lines show the prior median, whereas the dashed black lines show the 16^{th}–84^{th} percentiles of the prior.*

should be given to non-parametric and simulation-derived SFH models.

3.7 Non-parametric SFH models

The focus of this chapter is on parametric models for galaxy SFHs, however in this section I give a brief introduction to some of the possibilities that are available when using non-parametric models. The discussion presented here is based upon the work presented in Leja et al. (2018).

The simplest form of non-parametric SFH simply divides up the lifetime of the Universe into several time bins, then fits the amount of star-formation that took place in each bin. Whilst this approach does not explicitly assign parameters to the shape of the SFH, there are still several choices to be made, namely the number and spacing of time bins, the parameter to be fitted within each bin (e.g. SFR, sSFR, total mass formed) and the prior that is to be placed on this parameter.

Simple versions of this model have been employed recently in the literature (e.g. Chauke et al. 2018, Morishita et al. 2018, Belli et al. 2019) however little consideration has been given to the impact of this prior on derived parameters. In this section I will consider a specific case, in which the total stellar mass formed in a series of bins is fitted, with a logarithmic prior placed on the mass formed in each bin (examples to date in the literature typically use linear priors). The time bins I will employ are:

$$\begin{aligned} 0 < t < 30 \text{ Myr} \\ 30 < t < 100 \text{ Myr} \\ 100 < t < 330 \text{ Myr} \\ 330 \text{ Myr} < t < 1.1 \text{ Gyr} \\ 1.1 < t < 3.6 \text{ Gyr} \\ 3.6 < t < 11.7 \text{ Gyr} \\ 11.7 < t < 13.7 \text{ Gyr}. \end{aligned} \tag{3.6}$$

The priors imposed by this model on sSFR, mean stellar age and SFH shape are shown in the top row of Fig. 3.13. It can be seen that this model favours extremely bursty SFHs, with the majority of the stellar mass typically being formed in a single time bin. This leads to sharp peaks in the mean stellar age prior, and a high probability of very young ages and high sSFRs, due to the two bins that fall entirely within the most recent 100 Myr. From this I conclude that the use of this prior in effect expresses a strong belief that galaxy SFHs are intrinsically very bursty, with regular changes of up to several orders of magnitude in sSFR.

In order to construct a contrasting case, I now consider a different kind of prior: one favouring continuity between different time bins rather than rapid fluctuations. This is achieved by fitting for $\Delta\log(\text{SFR})$ between adjacent time bins, using a prior that explicitly down-weights sharp transitions in SFR. I adopt a Student's-t distribution as the prior on $x = \log(\text{SFR}_n/\text{SFR}_{n+1})$, such that

$$P(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\Gamma\left(\frac{\nu}{2}\right)}\left(1 + \frac{(x/\sigma)^2}{\nu}\right)^{-\frac{\nu+1}{2}} \tag{3.7}$$

where Γ is the Gamma function, n is the bin index, running over the bins given in Equation 3.6, and ν and σ are constants. I will adopt $\nu = 2$ and $\sigma = 0.3$,

with both the choice of parameterisation and values of these constants motivated by a comparison with the Illustris hydrodynamic simulation (see appendix B of Leja et al. 2018). This process demonstrates one way in which specific prior information, potentially motivated by simulation results, can be encoded in non-parametric SFH models.

The bottom row of Fig. 3.13 shows the priors imposed by the use of this continuity SFH model. This model can be shown to favour sustained constant SFRs, leading to steady stellar mass growth across cosmic time. As a consequence of this, the mean stellar age prior peaks at approximately half of the current age of the Universe. The prior on sSFR is peaked in a similar way to the parametric models considered in Section 3.3, as both this model and the parametric models considered earlier disfavour strong variation in sSFRs, preferring sustained star-formation over long time periods.

3.8 Conclusion

In this chapter I have carried out an investigation of the effects of four parametric SFH models (exponentially declining, delayed exponentially declining, lognormal and double power law) on galaxy stellar masses, SFRs and mass-weighted formation times. I have considered:

- The priors imposed on physical parameters by the use of each parametric SFH model in Section 3.3.
- The biases introduced when fitting mock high-SNR broad-band photometric data in Section 3.5.
- The consistency of SFHs inferred for a volume-complete, low-redshift sample of galaxies from GAMA with the cosmic SFRD evolution reported by Madau & Dickinson (2014) in Section 3.6.

In Fig. 3.1 I demonstrate that each of these parametric models imposes relatively similar, strongly peaked priors on sSFR, which could act to tighten and shift the SFMS, depending on the details of the modelling assumptions used. All four SFH models also impose a prior preference for stellar-mass assembly at later times (younger stellar ages) than is observed to be the case through measuring galaxy

SFRs at high redshift. Fig. 3.2 demonstrates that changing the prior probability densities on model parameters can change the priors on physical parameters at least as significantly as changing the parametric SFH model adopted. In particular, a uniform prior on $1/\tau$ for the tau model is less informative on galaxy sSFRs than a uniform prior on τ.

By fitting a mock catalogue of high-SNR broad-band photometry, I have shown in Fig. 3.8 that inferred stellar masses, SFRs and mass-weighted formation times/ages are prior-dependent at levels of at least 0.1, 0.3 and 0.2 dex respectively. However, the dominant factor that determines how well the true values of these parameters can be recovered is the shape of the true SFH, rather than the parametric model being fitted. The four parametric models are all significantly limited in their ability to reproduce SFHs with strong, recent variations in SFR, and will consequently return strongly biased parameters when fitting galaxies with these SFHs.

Under the assumption that the mock catalogue constructed in Section 3.5 is representative of real galaxy SFHs, I have shown that tau models produce more strongly biased physical parameter estimates than the other three models. However, in Fig. 3.10, I demonstrate that high-SNR broad-band photometry cannot discriminate between prior beliefs about which parametric SFH models are most appropriate for describing real galaxy SFHs. This means that carefully considered, physically motivated priors are a necessary component of any SED fitting analysis.

Finally I have fitted a volume-complete sample of galaxies at $0.05 < z < 0.08$ with high-quality photometric data from the GAMA Survey. I demonstrate in Table 3.2 and Fig. 3.11 that the Bagpipes stellar-mass and SFR measurements at $z \sim 0.05$ are consistent both with other SED fitting analyses from the literature and the SFRD curve of Madau & Dickinson (2014).

However, in Fig. 3.12 I demonstrate that the mass-weighted formation times inferred are significantly overestimated (mass-weighted ages are underestimated), as the ensemble of the fitted SFHs predicts a much later peak in cosmic SFRD than Madau & Dickinson (2014). This analysis suggests that the cosmic SFRD peaked at $z \sim 0.4$, approximately 6 Gyr later than is directly observed. A comparison of Fig. 3.1 with Fig. 3.12 suggests that this result is a consequence of the poorly-motivated priors imposed by these parametric SFH models.

These analyses demonstrate the challenges involved in using parametric SFH

models as tools for understanding the history of galaxy stellar-mass assembly. Non-parametric SFH models are a promising alternative to the parametric forms discussed in this chapter. As discussed in Section 3.7, such models both provide greater flexibility than parametric models, and allow prior beliefs to be incorporated in a more direct way. In a companion paper (Leja et al. 2018), the advantages and disadvantages of such models are considered in detail, by subjecting them to similar tests to those I have employed in this chapter.

The observation of samples of galaxies at different points in cosmic history and subsequent attempts to connect them are powerful tools for understanding galaxy evolution (e.g. Wild et al. 2016; Belli et al. 2019). However, despite the aim of understanding populations of galaxies, all current SED fitting analyses treat individual galaxies as statistically independent from each other. A more powerful approach would be to simultaneously model and fit whole populations of galaxies through the use of a hierarchical Bayesian model.

This could be used, for example, to enforce continuity between galaxy populations at different redshifts by treating the redshift evolution of the cosmic SFRD as a prior distribution. The hyper-parameters of this prior could then be jointly constrained by samples of galaxies across a range of observed redshifts, using both instantaneous SFRs and SFHs. Under this scheme we would be able to self-consistently model the redshift evolution of the GSMF, SFMS and cosmic SFRD, as well as obtaining better constraints on the SFHs of individual galaxies through the use of well-motivated population priors (Bayesian shrinkage).

Chapter 4

The star-formation histories of VANDELS massive quiescent galaxies at $1.0 < z < 1.3$

The material in this chapter was originally published in Carnall et al. (2019b).

4.1 Introduction

Understanding the origin of the colour bimodality in the local galaxy population remains one of the most important goals of extragalactic astronomy. The emerging picture is one in which feedback of energy from baryonic processes plays a central role in the quenching of star-formation activity, with supernovae and active galactic nuclei (AGN) thought to dominate at lower and higher masses respectively. A range of other factors are also thought to play important roles, for example mergers and environmental effects such as ram-pressure stripping. However, we still lack a detailed physical description of these and other relevant processes.

The inherent complexity of galaxy formation physics means that large-scale numerical simulations are required to connect theoretical models with observable properties. The extreme computational expense of such simulations precludes, for the present, a Bayesian statistical approach to parameter estimation and model selection in this context. Instead, predictions from individual simulations are

made for a number of observable properties, such as the distributions of galaxy colours, stellar masses and star-formation rates (SFRs), which can be compared with observational results. These comparisons provide qualitative insights as to how the underlying physical model should be refined (e.g. Davé et al. 2017; Trayford et al. 2017; Nelson et al. 2018; Donnari et al. 2019; Cochrane & Best 2018; Baes et al. 2019).

However, numerical simulations have now reached a level of complexity such that a range of models with varying physical ingredients can produce good approximations of the standard set of well-constrained observable properties. The challenge for new observational studies therefore is to provide precise measurements of a wider range of physical parameters that are highly constraining on galaxy formation models, such as stellar ages, star-formation histories (SFHs), stellar and nebular metallicities and levels of dust attenuation. Crucially, these studies must be performed on large and representative samples of galaxies. They must also be extended to high redshift, in order to constrain the evolution of galaxy properties across cosmic time.

An area of particular theoretical interest in recent years has been the AGN-feedback processes that quench star-formation in the most massive galaxies (e.g. Davé et al. 2016, 2019; Weinberger et al. 2017; Pillepich et al. 2018; Nelson et al. 2019). These processes should leave strong imprints, not just on the bright end of the galaxy luminosity function, but also on the star-formation histories of massive quiescent galaxies (e.g. Croton et al. 2006). Additional constraints are also available from the physical properties of galaxies transitioning between the star-forming and quiescent populations. If strong constraints can be placed on these more-subtle indicators, it will be possible to begin ruling out models that are capable of matching simpler observables.

A huge literature exists on the stellar ages of massive quiescent galaxies (e.g. see Chapter 2; Heavens et al. 2000, 2004; Panter et al. 2003, 2007; Cimatti et al. 2004, 2008; Gallazzi et al. 2005, 2014; Daddi et al. 2005; Onodera et al. 2012, 2015; Jørgensen & Chiboucas 2013; Whitaker et al. 2013; Choi et al. 2014, 2019; Fumagalli et al. 2016; Pacifici et al. 2016; Citro et al. 2016; Siudek et al. 2017; Belli et al. 2019; Estrada-Carpenter et al. 2019). However, these measurements are challenging to make for several reasons.

As stellar populations age their luminosities fall rapidly, meaning that the evidence of earlier star-formation episodes can easily be lost in the glare of

younger stellar populations. Additionally, galaxy spectral energy distributions (SEDs) suffer from strong, interrelated degeneracies between different physical properties, such as the age-metallicity-dust degeneracy (e.g. Papovich et al. 2001; Lee et al. 2007; Conroy 2013). Because of these issues, photometric data often fail to strongly constrain galaxy physical parameters, meaning the applied priors can significantly impact the results obtained, as discussed in Chapter 3. Finally, significant systematic uncertainties exist in the empirical models used to interpret observational data (e.g. Han & Han 2019).

The situation has been improved in the local Universe by the advent of large systematic surveys designed to obtain high signal-to-noise ratio (SNR) continuum spectroscopy. These data are more-strongly constraining on subtle galaxy physical parameters (e.g. Pacifici et al. 2012; Thomas et al. 2017), however, until recently such data has been extremely scarce at higher redshifts. Additionally, interpreting spectroscopic data in a way that makes full use of the available information is challenging, both in terms of the complexity of the required models and the computational expense of fitting these models to data (see Section 4.4.1). The use of simplifying assumptions to reduce the complexity of the problem, for example by fixing nuisance parameters to fiducial values, typically leads to biases in derived physical parameter values and to underestimated uncertainties (e.g. Pacifici et al. 2015; Iyer & Gawiser 2017).

Despite these challenges, a consensus has emerged around several important results. Firstly, at fixed observed redshift, less-massive galaxies are found to have younger stellar populations than their more-massive counterparts (e.g. Gallazzi et al. 2005, 2014; Pacifici et al. 2016). This is often referred to as downsizing, or mass-accelerated evolution. Secondly, at fixed stellar mass, a trend towards lower average formation redshift is found with decreasing observed redshift. A combination of factors contribute to this, including new galaxies joining the red sequence (e.g. Brammer et al. 2011; Muzzin et al. 2013; Tomczak et al. 2014), mergers (e.g. Khochfar & Silk 2009; Khochfar et al. 2011; Emsellem et al. 2011), and periods of rejuvenated star-formation activity (e.g. Belli et al. 2017).

Finally, there is considerable evidence for at least two distinct quenching mechanisms with different timescales, which change in relative importance with observed redshift (e.g. see Chapter 2; Schawinski et al. 2014; Schreiber et al. 2016; Wild et al. 2016; Maltby et al. 2018; Belli et al. 2019). Rapid quenching, often associated with post-starburst galaxies (e.g. Wild et al. 2009), is thought to dominate at high redshift ($z \gtrsim 1$), whereas slower quenching, associated

with green-valley galaxies, is thought to dominate at lower redshifts ($z \lesssim 1$). Whilst these three fundamental results have gained broad acceptance, precise quantitative measurements are still lacking, in particular at high redshift.

Within the last year, two new, large, high-redshift spectroscopic surveys have been completed: Lega-C (van der Wel et al. 2016; Straatman et al. 2018) and VANDELS (McLure et al. 2018a; Pentericci et al. 2018). These surveys have greatly expanded the availability of high-SNR continuum spectroscopy within the first eight billion years of cosmic history, providing new opportunities for placing strong constraints on subtle galaxy physical parameters (e.g. Cullen et al. 2017, in prep.; Wu et al. 2018b,c; Chauke et al. 2018).

In parallel, a new generation of spectral modelling and fitting tools has been developed (e.g. Chevallard & Charlot 2016; Leja et al. 2017; Johnson et al. in prep). These codes include complex, flexible physical models of the kind necessary to reproduce the properties of observed spectroscopic data, and make use of modern computational and statistical methods to fit these models to data within a fully Bayesian framework. This allows the recovery of full posterior distributions for physical parameters, meaning realistic uncertainties can be obtained, including an understanding of complex, multi-parameter degeneracies.

In this chapter I present the first analysis of a sample of extremely deep rest-frame near-ultraviolet spectra from VANDELS. The targets are 75 UVJ-selected galaxies with stellar masses of $\log_{10}(M_*/\mathrm{M}_\odot) > 10.3$ at observed redshifts of $1.0 < z < 1.3$. I analyse these spectra, in parallel with multi-wavelength photometry, using Bayesian Analysis of Galaxies for Physical Inference and Parameter EStimation (Bagpipes; Chapter 2). The Bagpipes code is used to fit a complex physical plus systematic uncertainties model, allowing strong yet realistic constraints to be placed on the physical parameters of these galaxies. In particular I will discuss their SFHs, quantifying the downsizing trend at $z \sim 1$ and comparing my results to predictions from modern cosmological simulations. I will also consider the properties of the post-starburst and green-valley galaxies in the sample, in an attempt to understand the evolutionary pathways of galaxies towards the red sequence.

The structure of this chapter is as follows. In Section 4.2, I introduce VANDELS, and give details of the selection of the sample. Then, in Section 4.3, I give details of the physical model I construct within Bagpipes to describe the targets. In Section 4.4 I discuss spectral fitting approaches, and introduce my spectroscopic

plus photometric fitting methodology. WI present my results in Section 4.5, discuss these results in Section 4.6, and present my conclusions in Section 4.7. All times, t, are measured forwards from the beginning of the Universe. For posterior distributions I quote 50^{th} percentile values and 16^{th}–84^{th} percentile ranges.

4.2 VANDELS data and sample selection

VANDELS (McLure et al. 2018a; Pentericci et al. 2018) is a large, recently completed ESO Public Spectroscopic Survey using the VIMOS instrument on the VLT at Paranal Observatory. The survey targeted 2106 high-redshift galaxies in the UKIDSS Ultra-Deep Survey (UDS) field and Chandra Deep Field South (CDFS). Whilst 87 per cent of the VANDELS targets are star-forming galaxies at $z > 2.4$, the final 13 per cent are massive, UVJ-selected passive galaxies at $1.0 < z < 2.5$. In this section I describe the VANDELS data, as well as the selection of the mass-complete sample of 75 objects from the VANDELS DR2 public release[1] considered in this chapter.

4.2.1 Photometric catalogues and parent sample

The VANDELS photometric catalogues and sample selection procedure are both described in full in McLure et al. (2018a). Here I present a brief summary of the key points relevant to this chapter. Both sets of VANDELS pointings in UDS and CDFS are centred on the CANDELS fields (Grogin et al. 2011; Koekemoer et al. 2011). Because the VIMOS field of view is larger than the areas imaged by CANDELS, the VANDELS photometric catalogues were supplemented with a variety of ground-based public imaging data. Because of this, the VANDELS sample is drawn from four different photometric catalogues, each spanning a UV-NIR observed wavelength range from ~ 0.3–5 μm.

Each of these catalogues was subjected to an extensive SED fitting campaign to construct derived-parameter catalogues, including robust photometric redshifts, z_{phot}, stellar masses and rest-frame magnitudes. The initial VANDELS sample selection was performed using a 2015 version of these catalogues, which has since been supplemented by deeper data. The photometric data used in this chapter,

[1] `https://www.eso.org/sci/publications/announcements/sciann17139.html`

as well as the stellar masses used in Section 4.2.3, come from the v1.0 internal catalogues, the final versions of which will be made public as part of the final VANDELS data release.

The VANDELS passive sample was selected by the following process. Firstly, objects were required to have $1.0 < z_{\rm phot} < 2.5$. Objects were then selected to have observed H-band magnitudes of $H < 22.5$, corresponding to stellar masses of $\log(M_*/\mathrm{M}_\odot) \gtrsim 10$. Next, objects were selected by rest-frame UVJ colours (e.g. Williams et al. 2009; Whitaker et al. 2011). In order to ensure that all targets would be detected with sufficient SNR in the VIMOS spectra, a final selection criterion was applied, requiring an observed i-band magnitude of $i < 25$. This slightly biases the full VANDELS passive sample against the faintest and reddest objects. This process, summarised below, results in a parent sample of 812 objects.

- $1.0 < z_{\rm phot} < 2.5$
- $H < 22.5$
- $U - V > 0.88(V - J) + 0.49$
- $U - V > 1.2$
- $V - J < 1.6$
- $i < 25$.

4.2.2 VANDELS spectroscopic observations

The VANDELS observations are described in full in Pentericci et al. (2018), and so I again provide here only a brief summary of the relevant points. Of the 812 objects selected by the process detailed in Section 4.2.1, a random sample of 268 were assigned slits and observed. All observations were conducted using the MR grism, providing $R \simeq 600$ spectroscopy spanning an observed wavelength range from $\lambda = 4800{-}10000$Å. Objects were observed for either 20, 40 or 80 hours, depending on their i-band magnitudes, to obtain SNRs of 15–20 per resolution element (∼10Å) in the i-band. Spectroscopic redshifts, $z_{\rm spec}$, were measured and verified manually by the VANDELS team.

A known issue with the VANDELS spectra is a systematic drop in flux at the blue end ($\lambda \lesssim 5600$Å). This region typically has a low SNR in the passive spectra, but an empirical correction was derived and implemented based on the bluer star-forming sample (see section 4.1 of Pentericci et al. 2018). This median correction is applied to all of the VANDELS spectra, however, object-to-object variations persist at levels of up to ~ 30 per cent. This calibration uncertainty is fitted as part of the model fitted in this chapter, as discussed in Section 4.4.

Because of the rapid build-up of the red sequence across the target redshift range, the passive sample is heavily weighted towards lower redshifts, with 88 per cent having $z_{\rm phot} < 1.5$. This means that, for the vast majority of the sample, these spectra contain a full suite of rest-frame UV-optical absorption features, including $\rm Mg_{UV}$, Ca H and K, the 2640Å, 2900Å and 4000Å breaks and Hδ and higher order Balmer lines, as well as the [O II] 3727Å emission line.

4.2.3 The $1.0 < z < 1.3$ mass-complete sample

As discussed in Section 4.2.1, the VANDELS passive sample is not mass-complete across the whole redshift range from $1.0 < z < 2.5$. Furthermore, the full suite of rest-frame UV spectral features described in Section 4.2.2 is only available at the lower end of this redshift range. Based on these considerations, it was decided to impose additional redshift and stellar-mass limits on the sample to define a mass-complete sample for which all of these features are available.

I first apply a limit of $z_{\rm spec} < 1.3$, such that the 4000Å break falls blue-ward of the strong sky-line contamination long-ward of 9250Å. I then return to the v1.0 photometry and derived parameter catalogues and re-apply the initial VANDELS passive sample selection criteria, excepting the i-band magnitude limit (see Section 4.2.1). For this reduced redshift range, 98 per cent of objects with stellar masses of $\log(M_*/\rm M_\odot) > 10.3$ meet the i-band limit imposed in Section 4.2.1. I therefore impose this mass limit, meaning my final sample is a random draw from a 98 per cent mass-complete sample. I finally require that objects have received > 90 per cent of their final exposure time in VANDELS DR2, and have a spectroscopic redshift quality flag of 3 or 4 (corresponding to > 95 per cent probability of being correct), resulting in a sample of 75 objects.

Table 4.1 *Parameters and priors for the model I fit to the VANDELS data. The model and fitting methodology are described in Sections 4.3 and 4.4. For Gaussian priors, μ is the mean and σ the standard deviation of the prior distribution. The upper limit on the τ parameter, $t_{\rm obs}$, is the age of the Universe at $z_{\rm spec}$. Logarithmic priors are uniform in log base ten of the parameter. The SFH model used is a double-power law; the calibration model is a second order Chebyshev polynomial. The form of the Gaussian process noise model is given in Equation 4.6.*

Component	Parameter	Symbol / Unit	Range	Prior	Hyperparameters	
Global	Redshift	z	$z_{\rm spec} \pm 0.015$	Gaussian	$\mu = z_{\rm spec}$	$\sigma = 0.005$
	Velocity Dispersion	$\sigma_{\rm vel}$ / km s^{-1}	(40, 400)	logarithmic		
SFH	Stellar mass formed	M_* / M$_\odot$	(1, 10^{13})	logarithmic		
	Metallicity	Z / Z$_\odot$	(0.01, 2.5)	logarithmic		
	Falling slope	α	(0.1, 1000)	logarithmic		
	Rising slope	β	(0.1, 1000)	logarithmic		
	Peak time	τ / Gyr	(0.1, $t_{\rm obs}$)	uniform		
Dust	Attenuation at 5500Å	A_V / mag	(0, 8)	uniform		
	Power-law slope	n	(0.3, 1.5)	Gaussian	$\mu = 0.7$	$\sigma = 0.3$
Calibration	Zero order	P_0	(0.5, 1.5)	Gaussian	$\mu = 1$	$\sigma = 0.25$
	First order	P_1	(−0.5, 0.5)	Gaussian	$\mu = 0$	$\sigma = 0.25$
	Second order	P_2	(−0.5, 0.5)	Gaussian	$\mu = 0$	$\sigma = 0.25$
Noise	White noise scaling	a	(0.1, 10)	logarithmic		
	Correlated noise amplitude	b / $f_{\rm max}$	(0.0001, 1)	logarithmic		
	Correlation length	l / $\Delta\lambda$	(0.01, 1)	logarithmic		

4.3 Physical model

This section describes the physical model I construct within Bagpipes to describe the VANDELS targets. The process by which Bagpipes generates spectral models is described in full in Chapter 2. A summary of the parameters and priors of the model is provided in Table 4.1. The fitting of this model to the data, as well as the systematic uncertainties model fitted, is described in Section 4.4.

For all objects, I vary the observed redshift within a narrow range centred on the spectroscopic redshift (measured by the VANDELS team, as described in Section 4.2.2). I impose a Gaussian prior centred on $z_{\rm spec}$, with standard deviation, $\sigma = 0.05$, and allow deviations of up to 3σ in either direction. Velocity dispersion is modelled within Bagpipes by convolution of the spectral model with a Gaussian kernel in velocity space. I apply a logarithmic prior to velocity dispersion, $\sigma_{\rm vel}$, between 40 and 400 km s^{-1}.

4.3.1 Stellar population model

In this chapter, I use the default Bagpipes stellar population models, which are the 2016 updated version of the Bruzual & Charlot (2003) models[2], using the MILES stellar spectral library (Falcón-Barroso et al. 2011). I model the chemical-enrichment histories of the VANDELS galaxies with a delta-function, assuming that all stars within the galaxy have the same metal content with scaled-Solar abundances. This single metallicity is varied with a logarithmic prior between $-2 < \log_{10}(Z_*/Z_\odot) < 0.4$ (I define $Z_\odot = 0.02$).

I parameterise the SFHs of the VANDELS galaxies using the double-power-law model described in Chapter 2. The SFH is described by

$$\mathrm{SFR}(t) \propto \left[\left(\frac{t}{\tau}\right)^{\alpha} + \left(\frac{t}{\tau}\right)^{-\beta}\right]^{-1} \tag{4.1}$$

where α is the falling slope, β is the rising slope and τ is related to the peak time. In Chapter 2 I demonstrate that this model produces unbiased estimates of the redshifts of formation and quenching for quiescent galaxies from the Mufasa simulation across a wide redshift range.

[2] https://www.bruzual.org/~gbruzual/bc03/Updated_version_2016

4.3.2 Dust attenuation model

Recent studies at high redshift have favoured an average attenuation curve slope of $n \simeq 0.7$, where $A_\lambda \propto \lambda^{-n}$ (e.g. Cullen et al. 2017, 2018; McLure et al. 2018b), similar to that found by Calzetti et al. (2000) for local galaxies. However, several studies suggest significant object-by-object variation (e.g. Kriek & Conroy 2013; Narayanan et al. 2018). I therefore model dust attenuation with the modified Charlot & Fall (2000) model described in Section 2.3.1. I place a Gaussian prior on n, with mean, $\mu = 0.7$ and standard deviation, $\sigma = 0.3$. I set permissive lower and upper limits of 0.3 and 1.5 respectively. I likewise allow a wide range of V-band attenuations, A_V, for stellar continuum emission, adopting a uniform prior from $0 < A_V < 8$. I adopt a fixed value of $\epsilon = 2$ for the ratio of attenuation between stellar birth clouds and the wider interstellar medium (ISM). I also adopt a value of 10 Myr for $t_{\rm BC}$, the lifetime of stellar birth clouds, meaning that A_V is doubled for emission from stars formed in the last 10 Myr.

4.3.3 Nebular emission model

The VANDELS UVJ selection is designed to identify galaxies with low levels of ongoing star-formation, and hence nebular emission. However, it is important for the fitted model to be capable of reproducing the spectra of dusty star-forming galaxies that can contaminate UVJ-selected samples. I therefore implement the nebular emission model described in Section 2.3.1, with a fixed ionization parameter of $\log_{10}(U) = -3$. Light from stars formed more recently than the lifetime I assume for stellar birth clouds ($t_{\rm BC} = 10$ Myr) is processed through the Bagpipes nebular model. The resulting nebular continuum and line emission is attenuated by twice the ISM A_V, as described in Section 4.3.2.

4.4 Combining spectroscopic and photometric data

This section describes the fitting of the physical model described in Section 4.3 to the combined datasets described in Section 4.2. I begin in Section 4.4.1 by reviewing the literature on galaxy spectral fitting. Then, in Section 4.4.2, I show that a simplistic approach that does not allow for systematic uncertainties fails

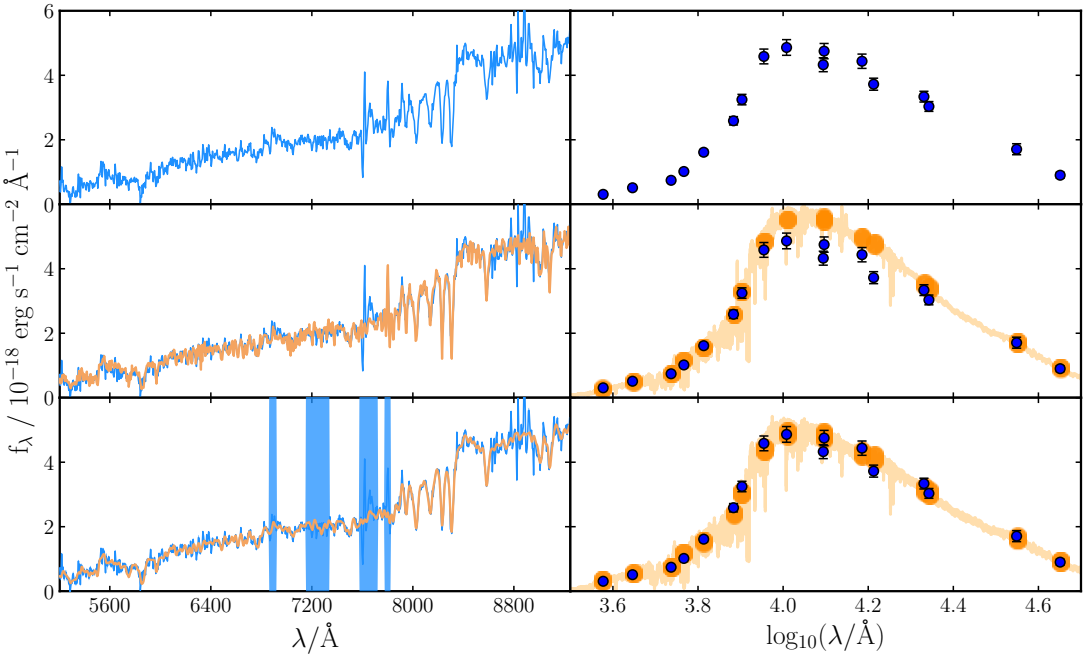

Figure 4.1 *An example object from the VANDELS sample fitted using different methods. Data are shown in blue, with spectroscopic data to the left and photometric data to the right. The blue shaded regions were masked in the final fits (see Section 4.4.4). Posterior distributions are shown in orange (the* 16^{th}–84^{th} *percentiles) for the simplistic model of Section 4.4.2 in the middle panels and for the joint physical plus systematics model of Section 4.4.3 in the bottom panels. An expanded version of the bottom-left panel is shown in Fig. 4.2.*

to describe the VANDELS data. In Section 4.4.3 I construct a model for these systematics, and in Section 4.4.4 I fit the combined physical plus systematics model to the joint datasets.

4.4.1 Historical approaches to spectral fitting

Historically, galaxy spectral fitting techniques have been applied to photometric data (e.g. Faber 1972), and this has remained a popular approach for several reasons. Photometric data are widely available across a wide range of wavelengths, and both random and systematic uncertainties are relatively simple to characterise (e.g. Mortlock et al. 2017; McLure et al. 2018a). It is also far simpler to construct models for the broad-band colours of galaxies than for detailed spectral features (e.g. Bell & de Jong 2001; Bell et al. 2003), and this has been shown to be sufficient for estimating basic physical properties such as stellar masses (e.g. Mobasher et al. 2015).

Such analyses typically assume that the uncertainties on photometric fluxes are well-determined, Gaussian distributed and independent. In this case, the scatter of observed fluxes f_i about their true values follows a chi-squared distribution (where i runs over a number, N_{bands}, of photometric bandpasses). A physical model that is a function of some parameters, $\boldsymbol{\Theta}$, can therefore be fitted to these observations using the log-likelihood function

$$\ln(\mathcal{L}_{\mathrm{phot}}) = K - 0.5 \sum_{i}^{\mathrm{N_{bands}}} \left(\frac{f_i - m_i(\boldsymbol{\Theta})}{\sigma_i} \right)^2 \tag{4.2}$$

where K is a constant, $m_i(\boldsymbol{\Theta})$ is the model prediction for the observed flux f_i, and σ_i is the corresponding uncertainty.

Significant failures of the above assumptions are fairly simple to identify by assessing the quality of fit, typically using the minimum reduced chi-squared value, and are rare enough that the affected objects can simply be excluded from the analysis. More-subtle failures can be modelled, for example by applying variable zero-point offsets to each band (e.g. Brammer et al. 2008), or by asserting that the uncertainty be greater than some fixed fraction of the observed flux (typically 5 per cent; e.g. Muzzin et al. 2013; Belli et al. 2019). This prevents uncertainties from being underestimated in the high-SNR regime, where the precision of the photometric calibration dominates the error budget.

However, as described in Section 4.1, photometric data are limited in their ability to constrain more-subtle galaxy physical parameters. It has been shown that spectroscopic observations have the potential to improve this situation, however difficulties in accurate spectrophotometric calibration (e.g. van der Wel et al. 2016) and the construction of precise, high resolution spectral models make these analyses challenging in practice. Historically therefore, analyses of spectroscopic data have been limited to individual spectral features (commonly Lick indices), such that results are independent of spectrophotometric calibration (e.g. Faber et al. 1985; Gorgas et al. 1993; Worthey et al. 1994). Whilst these analyses have produced many extremely valuable results, they do not make use of the full information content of spectroscopic data (e.g. Conroy et al. 2018).

More recently, with the advent of large, well calibrated spectroscopic surveys such as the Sloan Digital Sky Survey (SDSS; York et al. 2000), attention has shifted towards full-spectral-fitting methods, which attempt to model and fit the whole information content of spectroscopic data (e.g. Heavens et al. 2000, 2004; Panter

et al. 2003, 2007; Cid Fernandes et al. 2005; Ocvirk et al. 2006; Tojeiro et al. 2007). The simplest approach to full spectral fitting is to make the same assumption of well-determined, independent Gaussian uncertainties on each spectral pixel flux f_j (where j runs over the number of pixels in the spectrum, $N_{\rm pix}$). In this case, the log-likelihood function can again be written as

$$\ln(\mathcal{L}_{\rm spec}) \;=\; K - 0.5 \sum_{j}^{\rm N_{pix}} \left(\frac{y_j \;-\; m_j(\Theta)}{\sigma_j} \right)^2 \tag{4.3}$$

where K is a constant, $m_j(\Theta)$ is the model prediction for the pixel flux f_j, and σ_j is the corresponding uncertainty.

This approach has however been demonstrated to be less successful in describing spectroscopic data than photometry. Panter et al. (2003) note that the quality of the fits that they obtain are typically poor, and attribute this to both inadequacies in the models they fit to their data, and the difficulty of obtaining reliable uncertainties. Furthermore, in Panter et al. (2007), the authors report that improvements to the SDSS spectrophotometric calibration have significantly changed their inferred SFHs. These and other authors also note the challenges that exist in the exploration of the higher-dimensional parameter spaces of the more complex models required to fit spectroscopic data.

As the availability of high-quality panchromatic photometric data has increased, interest has grown in the joint analysis of spectroscopic and photometric data (e.g. Chevallard & Charlot 2016; Belli et al. 2019; Johnson et al. in prep). These kinds of analyses promise the ability to take advantage of both the broad wavelength coverage and excellent calibration of photometry, and the strong constraints on subtle physical parameters offered by spectroscopy.

4.4.2 A simplistic approach to fitting the joint datasets

When jointly analysing several datasets, the likelihood function is the product of the separate likelihoods. The log-likelihood function for the joint analysis of spectroscopy and photometry is therefore $\ln(\mathcal{L}) = \ln(\mathcal{L}_{\rm phot}) + \ln(\mathcal{L}_{\rm spec})$. In this section I attempt to jointly fit the VANDELS spectroscopic and photometric datasets using the physical model described in Section 4.3, and the ansatz for $\ln(\mathcal{L}_{\rm phot})$ and $\ln(\mathcal{L}_{\rm spec})$ given in Equations 4.2 and 4.3 respectively.

For the photometric data I employ the common methods discussed in Section 4.4.1 for dealing with systematic calibration uncertainties. I first apply the photometric zero-point offsets calculated for the VANDELS photometric catalogues by McLure et al. (2018a). I also assert that the uncertainty for each band must be ≥ 5 per cent of the observed flux, except for the two IRAC channels where a threshold of 10 per cent is used. The VANDELS DR2 spectra were binned by a factor of two to a sampling of 5Å. Fitting was carried out using the wavelength range from 5200Å $< \lambda <$ 9250Å where the detector sensitivity is high and sky-line contamination is minimal compared to longer wavelengths.

The results of this method applied to an example object from the VANDELS sample are shown in the middle panels of Fig. 4.1. It can be seen that the overall shape of the posterior distribution matches the spectroscopic data well. However, on closer inspection (and comparison to the above panel), it can be seen that the depths of individual absorption features are poorly reproduced. The most obvious failure however is in reproducing the observed photometry from 1−2 μm, where this model significantly overestimates the observed fluxes. In accordance with Panter et al. (2003), the quality of fit to both datasets is poor, leading to unrealistically tight constraints on model parameters. Issues of this nature were observed for the majority of galaxies in the sample.

The fact that the joint fit is incapable of matching both datasets means that the two are inconsistent under the assumptions made. The joint fit adopts a region of parameter space that best describes the spectroscopic data at the expense of the photometry because there are a larger number of spectral pixels than photometric bands, and hence more terms in the log-likelihood function that depend on the spectroscopy than the photometry. As an additional check, I fitted this model to both the photometric and spectroscopic data separately. I verified that this model is capable of producing high-quality fits to the observed photometry, including the data from 1−2 μm, however the fit to the spectroscopy does not improve when the photometric data are excluded.

Issues of this nature have been commonly observed in similar analyses, leading to suggestions that the spectroscopic data should be somehow down-weighted in the likelihood function, in order to give "equal consideration" to both datasets. However, apart from being statistically unjustified, this cannot solve the underlying issues of inconsistency between the datasets and poor quality of fit to spectroscopic data. Instead it is necessary to understand the causes of these issues, so that these effects can be included in the model fitted to the joint

datasets.

As described in Section 4.2.2, there is a known issue with the VANDELS calibration at the blue end of the spectra. Accurate spectrophotometric calibration is notoriously challenging, owing to the need to correct for a range of atmospheric and instrumental effects, such as differential atmospheric refraction, telluric contamination and characterisation of the sensitivity function of the detector. Even for well-calibrated spectra, wavelength-dependent uncertainties are known to exist, typically at levels of ~ 10 per cent (e.g. Moehler et al. 2014; Xiang et al. 2015; Yan et al. 2016).

The issue in VANDELS is likely to be related to one or more of these corrections, however the precise cause is still under investigation. When comparing the results of the separate spectroscopic and photometric fits, the stellar population ages recovered are significantly older in the spectroscopic fits. An underestimation of the blue flux is consistent with causing older stellar populations to be fitted, and hence the overestimation of the 1−2 μm photometry in the joint fits, which are dominated by the spectroscopic data.

Another possible reason for effects of this kind would be aperture-bias, as the photometry is measured within 2″−diameter apertures, whilst the spectroscopic observations use a 1″ slit. However, the magnitude of the effect is considerably greater than estimated by previous similar studies at lower observed redshifts, which find small aperture effects (e.g. Gallazzi et al. 2005, 2014). I would expect a smaller aperture-bias effect than these studies, due to the larger angular diameter distance to the VANDELS target redshift range, and the smaller physical sizes of quiescent galaxies at higher redshifts (e.g. McLure et al. 2013).

As an additional check, I obtained size measurements from van der Wel et al. (2014) for the 26 objects in the sample that have CANDELS imaging. The mean effective radius, $r_{\rm e}$, for these galaxies is 0.35″, and I confirm that no correlation exists between $r_{\rm e}$ and the degree of inconsistency between the two datasets. Aperture-bias also cannot explain the poor quality of fit observed when separately fitting the spectroscopic data. I therefore conclude that systematic calibration uncertainties in the VANDELS spectroscopic data are the cause of the issues identified, and move on to develop a model for these systematic uncertainties.

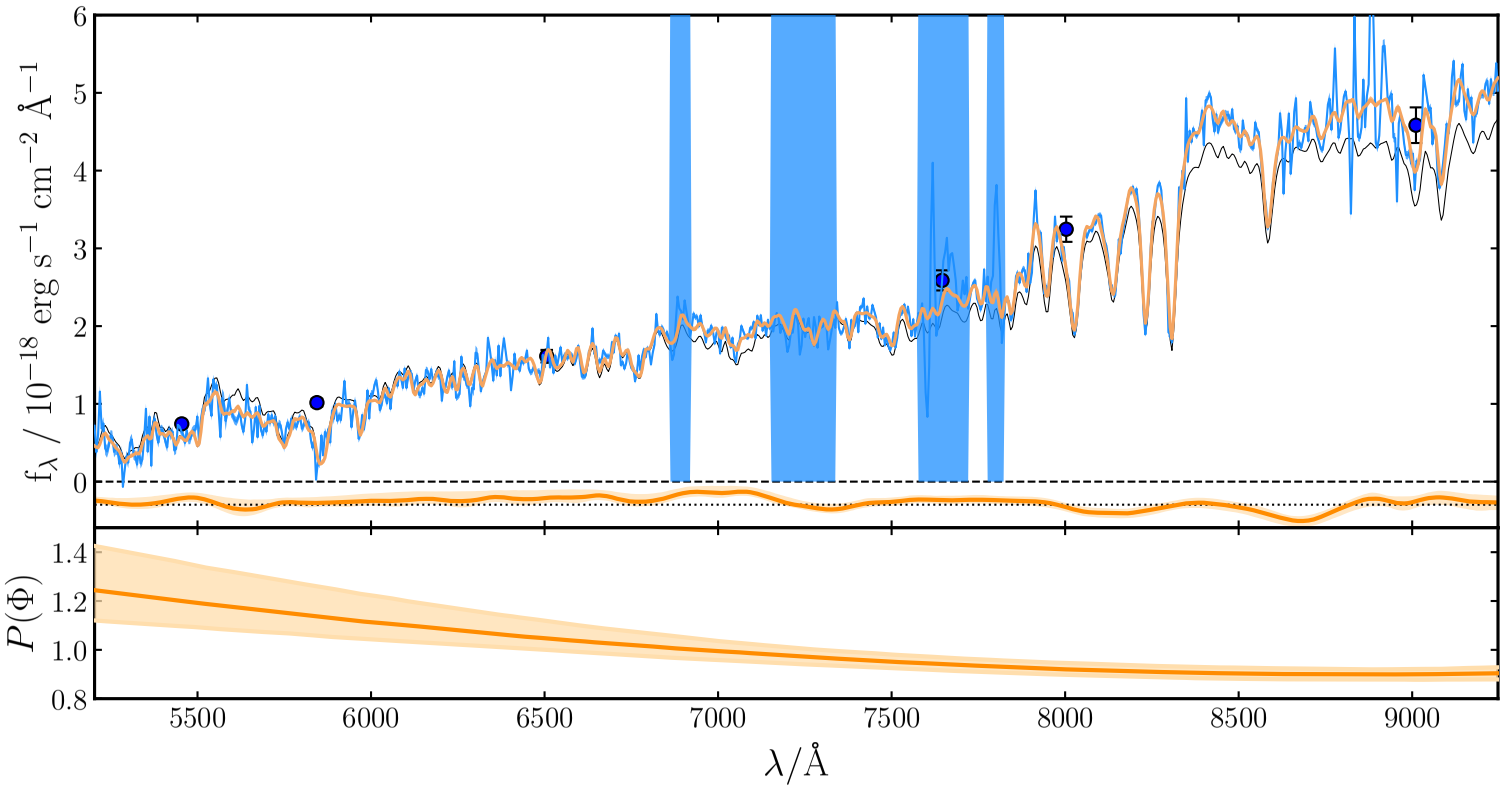

Figure 4.2 *Physical, noise and calibration models fitted to the same example object as shown in Fig. 4.1. The observed spectrum is shown in blue. The blue shaded regions have been masked, as described in Section 4.4.4. The orange line overlaid on the spectrum shows the posterior distribution for the full physical + calibration + noise model. The posterior distribution for the Gaussian process noise model is shown below in the same panel at an arbitrary vertical position. The zero-point about which it varies is shown with a dotted line. The posterior distribution for the polynomial calibration model is shown in the bottom panel. The posterior median for the physical model alone is shown in black. This is analogous to the spectrum corrected for systematic effects. To obtain the orange posterior in the top panel, the black line is divided by the polynomial at the bottom, then added to the Gaussian process model in the centre.*

4.4.3 Modelling spectroscopic systematic uncertainties

In general, as discussed in Section 4.4.2, it is extremely complex to construct a physical model for the atmospheric and instrumental effects to which both spectroscopic data and empirical stellar-population models are subjected. I therefore take a different approach, by constructing a flexible empirical model for systematic perturbations of the spectroscopic data about the physical model.

I will split these perturbations by the general form they take: either additive or multiplicative (e.g. Cappellari 2017). I will refer to these as noise and calibration offsets respectively. I then construct flexible models for these offsets by introducing nuisance parameters, $\mathbf{\Phi}$, into the spectroscopic log-likelihood function. I can then later marginalise these nuisance parameters out of the posterior distribution, in order to obtain a posterior for the physical parameters that includes uncertainties due to systematic effects.

I modify the log-likelihood function for spectroscopy presented in Equation 4.3 as follows. Firstly I generalise the model, $m(\mathbf{\Theta})$ to $m(\mathbf{\Theta}, \mathbf{\Phi})$ by dividing through by a multiplicative polynomial calibration model, $P_j(\mathbf{\Phi})$, such that

$$m_j(\mathbf{\Theta},\ \mathbf{\Phi}) = \frac{m_j(\mathbf{\Theta})}{P_j(\mathbf{\Phi})}. \tag{4.4}$$

This model will be discussed in Section 4.4.3.

Secondly, I drop the assumption that the uncertainties on the observed spectroscopic fluxes are independent, allowing additive correlated noise between spectral pixels. I hence replace the second term on the right of Equation 4.3 with a matrix equation, in which the inverse of the covariance matrix is multiplied on both sides by the vector of residuals between the observed and model fluxes. The covariance matrix, $\mathbf{C}(\mathbf{\Phi})$, will be drawn from a Gaussian process model, which is described in Section 4.4.3. The spectroscopic log-likelihood function is now

$$\ln(\mathcal{L}_{\rm spec}) \ = \ K - \ln\big(|\mathbf{C}(\mathbf{\Phi})|\big) \ - \Delta^{\rm T}\ \mathbf{C}(\mathbf{\Phi})^{-1}\ \Delta \tag{4.5}$$

where $\Delta \ = \ y_j \ - \ m_j(\mathbf{\Theta},\ \mathbf{\Phi})$ is the vector of residuals between the observed and model fluxes and K is a constant. Equation 4.5 is simply a generalisation of Equation 4.3, and reduces back to Equation 4.3 for the case in which $P_j(\mathbf{\Phi}) = 1$ and $\mathbf{C}(\mathbf{\Phi})$ is diagonal with elements σ_j^2.

Polynomial calibration model

A method recently adopted by several authors for addressing spectrophotometric calibration uncertainties is to perturb the spectroscopic data by a polynomial function of wavelength. One approach is to set the polynomial coefficients before fitting a physical model to the data, by comparing synthetic photometry derived from the spectrum to observed photometry in the same wavelength range (e.g. van der Wel et al. 2016). This is computationally simple, however it requires a significant number of photometric observations in the spectroscopic wavelength range. Another approach is to fit the polynomial coefficients at the same time as the parameters of the physical model, which has the advantage of incorporating calibration uncertainties into the uncertainties on physical parameters, but is more computationally expensive (e.g. Cappellari 2017; Belli et al. 2019).

I take the latter approach, parameterising $P_j(\Phi)$ in Equation 4.4 with a second-order Chebyshev polynomial. Because the physical model is divided by $P_j(\Phi)$, the polynomial posterior can be thought of as the multiplicative offset that would need to be applied to the spectroscopic data to correct its calibration. I caution however that any issues with the spectrophotometric calibration of the model spectra will also be incorporated into this polynomial.

The choice of a second order polynomial was made firstly due to computational constraints, and secondly because the calibration issues with the VANDELS spectra are known to be well approximated by a quadratic function of wavelength (see fig. 4 of Pentericci et al. 2018). A more-flexible choice for this model would be a multiplicative Gaussian process model (e.g. Johnson et al. in prep), and I intend to explore this option in future work.

I apply Gaussian priors to all three polynomial coefficients with standard deviations of $\sigma = 0.25$. The prior means are $\mu = 1$ for the zero order, and $\mu = 0$ for the first and second order. This means that the prior mean and median for $P_j(\Phi)$ are equal to 1 for all wavelengths, equivalent to no change in the calibration. The maximum deviation allowed for any polynomial order is 2σ from the mean.

Gaussian process noise model

As discussed in Section 4.4.1, it is common to assume that uncertainties on observational data are independently Gaussian distributed with well-known

variances. There are many good reasons to suspect that these assumptions do not hold in the case of spectroscopic data.

Firstly, it is standard practice to oversample the resolution element of the optical system by at least a factor of two, leading to local covariances between pixels, although this can be mitigated to some extent by binning adjacent pixels. Secondly, there is also good evidence that the error spectra determined from typical data reduction pipelines are underestimates of the true pixel variances (e.g. Panter et al. 2003, 2007; Belli et al. 2019). A common approach is to expand the variances for all spectral pixels by the median residual determined from an initial round of fitting (e.g. Belli et al. 2019). Finally, a range of effects from template mismatch to poor sky subtraction have the potential for introducing correlated additive offsets between the data and models being fitted (e.g. Cappellari 2017).

By modifying the log-likelihood function presented in Equation 4.3 to that of Equation 4.5, I have relaxed the assumption of independence in the spectroscopic uncertainties. I now parameterise the covariance matrix, $\mathbf{C}(\Phi)$, in terms of both independent (white) noise and covariant noise between pixels. I will fit these parameters alongside those of the physical and polynomial calibration models. The form I assume for the covariance matrix is

$$\mathbf{C}_{jk}(\Phi) \;=\; a^2\sigma_j\;\sigma_k\;\delta_{jk} \;+\; b^2\exp\left(-\frac{(\lambda_j-\lambda_k)^2}{2l^2}\right) \tag{4.6}$$

where $\sigma_{j,k}$ are the uncertainties on the pixel fluxes, $\lambda_{j,k}$ are the central wavelengths of the pixels, δ_{jk} is the Kronecker delta function, and a, b and l are free parameters to be fitted.

The first term in Equation 4.6 deals with the uncorrelated noise on the data. As I suspect that the uncertainties may be underestimated, I allow their magnitude to vary by a^2, where a is assigned a logarithmic prior between 0.1 and 10 (e.g. see section 6 of Hogg et al. 2010). This is similar to the iterative approaches of other authors, however the fact that this parameter is allowed to vary during fitting means that its uncertainty is propagated into the uncertainties on the physical parameters.

The second term in Equation 4.6 is drawn from a Gaussian process model, and allows for the modelling of covariant noise between spectral pixels. Gaussian process regression is implemented in Bagpipes using the George Python package (Ambikasaran et al. 2015). I adopt an exponential-squared kernel and

fit the normalisation b and correlation length l. I assign logarithmic priors to both of these quantities. I define b in units of the maximum flux in the observed spectrum, $f_{\rm max}$, and allow values from 10^{-4} to 1. The maximum flux is used as the unit of b such that the same range of prior values can be used for each spectrum. The mean or median flux value could also have been used.

Similarly, I define l in units of the wavelength range covered by the spectral data, $\Delta\lambda$ (in this case 4050Å), and allow values from 0.01 to 1. The minimum correlation length ($\sim$ 40Å) was chosen to prevent the Gaussian process model from reproducing individual absorption and emission features in the spectra. The Gaussian process model is intended to model poor sky subtraction and template mismatch between the models and data, as demonstrated with an additive polynomial by Cappellari (2017). As currently implemented, it cannot model covariances between adjacent spectral fluxes due to oversampling of the resolution element, or resampling from an initial non-uniform wavelength sampling. A term in Equation 4.6 to account for this is a possible extension to this model (e.g. Czekala et al. 2015).

4.4.4 Final fitting of the joint datasets

In this section I describe my final fitting methodology, from which all of the results presented in Section 4.5 are derived. I again fit the physical model described in Section 4.3, however I now also fit the models for systematic effects introduced in Section 4.4.3, by exchanging Equation 4.3 for Equation 4.5 in the log-likelihood function.

The photometric data are treated in the same way as described in Section 4.4.2, and I use the same wavelength range and binning for the spectroscopy. In addition, I also mask several spectral regions that experience strong telluric contamination, leading to residuals such as those visible in the top left panel of Fig. 4.1 at $\sim$ 7600Å. The regions masked are 6860−6920Å, 7150−7340Å, and 7575−7725Å. Finally, I mask the rest-frame region from 3702−3752Å, containing the [O II] emission line. This is because the excitation mechanism for low-level line emission in quiescent galaxies is still controversial, with AGN and ionization from old stars both thought to contribute (e.g. Yan et al. 2006; Lemaux et al. 2010; Singh et al. 2013; Herpich et al. 2018). By contrast, the only mechanism that can excite [O II] emission in the Bagpipes physical model is ionization from young stars, meaning that the inferred SFRs could be biased by [O II] emission

excited by other processes. I will compare the observed [O II] equivalent widths to the inferred specific star-formation rates (sSFRs) in Section 4.5.5.

The combined model has 15 free parameters, summarised in Table 4.1. Sampling from the posterior distribution with MultiNest therefore requires several million evaluations of the log-likelihood function, each of which is relatively computationally expensive, in particular the inversion of the covariance matrix. Fitting each galaxy therefore requires ~ 100 CPU hours, limiting the scalability of this method.

The posterior distribution for the combined model fitted to the object discussed in Section 4.4.2 is shown in the bottom panels of Fig. 4.1. Both the spectroscopic and photometric data can now be seen to be well matched by the posterior distribution. An expanded view of the bottom left panel of Fig. 4.1 is shown in Fig. 4.2. Additionally, the posterior distributions for the polynomial calibration and Gaussian process noise models are shown below the observed spectrum.

The posterior distribution for the polynomial calibration model can be seen to follow the expected form, as discussed in Section 4.4.2. The spectrophotometric calibration can be seen to be systematically high by ~ 10 per cent across most of the wavelength range. This is within the expected range, given that absolute flux calibration was performed by normalising the observed spectrum to the observed i-band flux. A drop in flux of ~ 30 per cent is observed at the blue end, within the expected range (see Section 4.2.2).

This is typical of the polynomial corrections recovered for objects in the sample, although the degree of correction at the blue end varies by ~ 30 per cent from object to object. The corrections introduced by the Gaussian process noise model can be seen to be small, however the extra flexibility in continuum shape allows the absorption features present in the observed spectrum to be well fitted by this model.

Finally, the black line in the main panel of Fig. 4.2 shows the posterior median for the physical model fitted to the observed spectrum. Assuming that systematics on the physical model fitted are negligible, this can be thought of as the best fit to the observational data corrected for systematic effects. For clarity, the black line in the top panel divided by the polynomial in the bottom panel then added to the Gaussian process model in the centre gives the orange posterior distribution in the top panel.

4.5 Results

In this section I present the results of the analysis described in Section 4.4.4, applied to the sample described in Section 4.2.3. I begin by splitting the sample into quiescent and green-valley sub-samples in Section 4.5.1. I then discuss the stellar ages of the quiescent sub-sample as a function of stellar mass in Section 4.5.2. I report the distribution of the VANDELS galaxies on the D4000 vs Hδ plane in Section 4.5.3 and on the UVJ diagram in Section 4.5.4. I discuss [O II] emission in Section 4.5.5, post-starburst and rejuvenated galaxies in Section 4.5.6, and AGN activity in Section 4.5.7. Finally, I present stacked spectra for both sub-samples in Section 4.5.8.

4.5.1 Quiescent and green-valley sub-samples

Two main methods have been used to define samples of quiescent galaxies: selection by sSFR and selection by rest-frame UVJ colours (typically evolving with observed redshift; e.g. Williams et al. 2009). Several recent studies define quiescence by a time-evolving criterion of sSFR $< 0.2/t_{\rm obs}$, where $t_{\rm obs}$ is the age of the Universe when the galaxy is observed (e.g. Gallazzi et al. 2014; Pacifici et al. 2016). This was demonstrated in Chapter 2 to produce good agreement with UVJ selection using a non-evolving colour criterion of $U-V > 0.88(V-J)$ + 0.69 at all redshifts.

As detailed in Section 4.2, the VANDELS UVJ selection uses the standard $1.0 < z < 2.0$ colour criterion of $U-V > 0.88(V-J) + 0.49$, which is more permissive. I therefore apply a further, sSFR-based selection to the sample in order to facilitate comparisons with other recent work. I apply the slightly modified method of Chapter 2, which uses the normalised SFR, or nSFR. This is defined as the SFR averaged over the most recent 100 Myr, $\rm SFR_{100}$, as a fraction of the average SFR over the whole history of the galaxy. This can be written as

$$\mathrm{nSFR} \;=\; \mathrm{SFR}_{100}\frac{t_{\rm obs}}{\int_0^{t_{\rm obs}} \mathrm{SFR}(t)\ \mathrm{d}t}. \tag{4.7}$$

In Chapter 2 I demonstrate that a selection criterion of nSFR < 0.1 produces good agreement with sSFR $< 0.2/t_{\rm obs}$ at all redshifts (see Fig. 2.6), whilst being insensitive to the age of the stellar population. This requires the ongoing SFR

of the galaxy at the redshift of observation to be less than 10 per cent of its historical average. By the application of this criterion I separate the VANDELS sample into 53 quiescent and 22 green-valley galaxies. This approach will be compared to UVJ-based selection criteria in Section 4.5.4.

4.5.2 Stellar mass vs formation redshift

From the star-formation histories inferred for the quiescent sub-sample I measure several quantities. Firstly, I calculate the mean time (measured forwards from the Big Bang) at which the stars in each galaxy formed, $t_{\rm form}$, given by

$$t_{\rm form} = \frac{\int_0^{t_{\rm obs}} t \ {\rm SFR}(t) \ {\rm d}t}{\int_0^{t_{\rm obs}} {\rm SFR}(t) \ {\rm d}t}. \tag{4.8}$$

This corresponds to the mean stellar age, or mass-weighted age, and is similar to the median formation time, t_{50}, used in some studies. I then calculate the redshift corresponding to $t_{\rm form}$, which I call the formation redshift.

I also calculate the history of the nSFR parameter over the inferred SFHs and extract the age of the Universe at which nSFR first falls below 0.1. This is the time at which the galaxy would first enter the quiescent sub-sample, and I therefore refer to this as the time of quenching, $t_{\rm quench}$.

Times of formation and quenching inferred using the double-power-law SFH model described in Section 4.3.1 were extensively validated using mock photometric observations of simulated quiescent galaxies in Chapter 2. However, evidence presented in Chapter 3 suggests that this SFH model under-predicts the ages of star-forming galaxies. I therefore do not report inferred ages for the green-valley sub-sample.

The times of formation and quenching inferred for the quiescent sub-sample are shown in Fig. 4.3 as a function of their inferred stellar masses. A clear trend towards earlier $t_{\rm form}$ and $t_{\rm quench}$ with increasing stellar mass is visible. Lower-mass galaxies are typically found to have formed their stellar populations at $z < 3$. By contrast, the formation of the stars in the most massive galaxies is found to have occurred at very early times, with the oldest objects found to have formed their stars at $z \sim 5$ and quenched by $z = 3$.

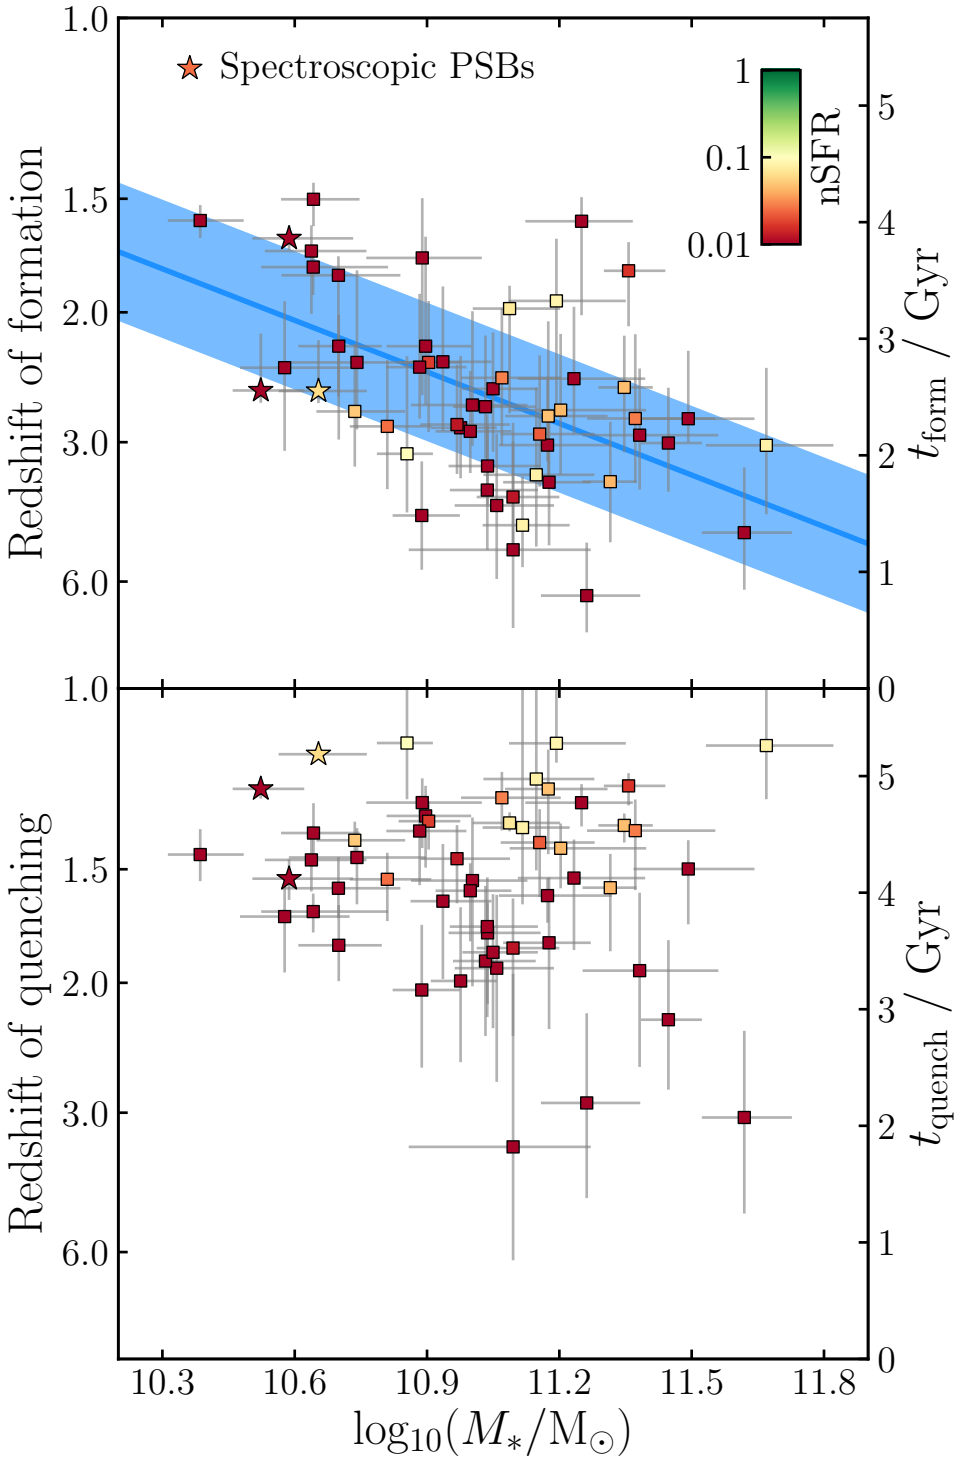

Figure 4.3 *Redshifts of formation and quenching for the quiescent sub-sample. The posterior median straight line fit to the data is shown in blue (Equation 4.9); the shaded region shows the posterior median intrinsic scatter. Individual SFHs are shown in Fig. 4.13.*

I fit a linear relationship to the recovered $t_{\rm form}$ values as a function of stellar mass, including an intrinsic scatter in $t_{\rm form}$ to account for effects unrelated to the stellar masses of these objects (e.g. galaxy environment). I find

$$\left(\frac{t_{\rm form}}{\rm Gyr}\right) = 2.56^{+0.12}_{-0.10} - 1.48^{+0.34}_{-0.39} \log_{10}\left(\frac{M_*}{10^{11}{\rm M}_\odot}\right) \tag{4.9}$$

with an intrinsic scatter of $0.58^{+0.09}_{-0.08}$ Gyr. The posterior median relationship is shown in blue in the top panel of Fig. 4.3, along with the posterior median intrinsic scatter. This result will be discussed in Section 4.6.1.

It is interesting to consider the distribution of SFRs for the objects shown in Fig. 4.3. For all objects with nSFR < 0.01 (those coloured dark red) the lower bound on the current level of star-formation is zero: I hence describe these objects as having no detectable star-formation. Galaxies with lighter colours have detectable star-formation at levels low enough that I still describe them as quiescent. It can be seen that no trend exists between $t_{\rm form}$ and nSFR in Fig. 4.3. This supports the findings of Belli et al. (2017), who attribute star-formation in quenched galaxies to stochastic processes such as minor mergers and rejuvenation events.

A trend is visible between $t_{\rm quench}$ and nSFR, with more-recently quenched galaxies having higher nSFR. However this result is not robust, as the double-power-law SFH model cannot reproduce rejuvenation events. Once star-formation drops to near zero it cannot rise again under this model, meaning quenching must be delayed until recent times to match any level of ongoing star formation. I hence conclude that $t_{\rm quench}$ cannot be reliably measured with the double-power-law model from UV spectroscopy for galaxies with detectable ongoing star-formation. A more advanced approach will be needed to model the details of these SFHs (e.g. Leja et al. 2018; Iyer et al. 2019; Lovell et al. 2019).

4.5.3 Distribution in D4000 vs Hδ

Historically, a common method for inferring galaxy ages and sSFRs from UV-optical spectroscopy has been to measure the strengths of the 4000Å break (D4000) and the Balmer delta (Hδ) absorption feature (e.g. Kauffmann et al. 2003; Brinchmann et al. 2004). I therefore report these spectral indices for the VANDELS galaxies, both to show the distribution of these parameters within the quiescent population at $1.0 < z < 1.3$, and to check that the results I infer from

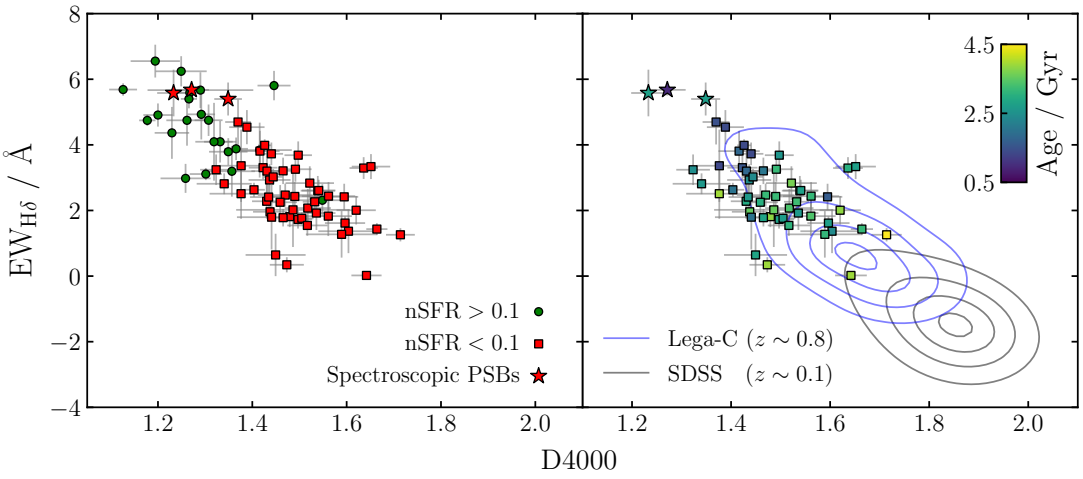

Figure 4.4 *Distribution of the VANDELS galaxies in $EW_{H\delta}$ vs D4000. To the left, the quiescent (red) and green-valley (green) sub-samples are shown, and can be seen to be cleanly separated in this parameter space at D4000 $\sim$ 1.3–1.4 and $EW_{H\delta}$ $\sim$ 4Å. To the right, the quiescent sub-sample is shown coloured by the inferred mass-weighted age. Contours are also shown marking the distributions of SDSS (gray) and Lega-C (blue) quiescent galaxies, also selected by nSFR < 0.1, from the samples of Wu et al. (2018b, priv. comm.).*

my full-spectral-fitting method are in agreement with the expected relationships between these parameters.

I measure D4000 from the VANDELS spectra as the ratio of average fluxes between 3850–3950Å and 4000–4100Å, whilst manually masking out pixels that experience significant sky-line contamination. I measure the rest-frame Hδ equivalent width, EW$_{H\delta}$, by fitting a first-order polynomial plus Gaussian model to the 100Å spectral region centred on Hδ. I then correct for nebular emission using the Hδ flux predicted by the fitted Bagpipes model (see Section 4.3.3). These corrections are small ($\lesssim$ 1Å) for the green valley sub-sample, and typically negligible for the quiescent sub-sample. For 8 of the 75 objects, Hδ is contaminated by sky lines, such that no measurement could be made. For these objects I apply the same fitting methodology to the posterior prediction for this spectral region from the fitted Bagpipes model.

The VANDELS sample is shown on the EW$_{H\delta}$ vs D4000 plane in the left panel of Fig. 4.4. The two sub-samples are significantly offset, with green valley objects having stronger Hδ absorption and a weaker 4000Å break. There is a clear transition at 1.3 < D4000 < 1.4 and EW$_{H\delta}$ $\sim$ 4Å, with almost all of the quiescent sub-sample at higher D4000 and lower EW$_{H\delta}$. Four significant outliers

are visible: three quiescent objects with low D4000 and strong Hδ absorption, and one green-valley object with a strong 4000Å break (D4000 $\sim$ 1.55). These objects will be discussed in Section 4.5.6.

In the right panel of Fig. 4.4 the quiescent sub-sample is shown coloured by their inferred mass-weighted ages (see Section 4.5.2). A trend in age with D4000 is visible as expected, with the oldest objects having D4000 $\simeq$ 1.7. Contours are plotted showing the distributions of quiescent galaxies at lower observed redshifts. Both samples are taken from Wu et al. (2018b, priv. comm.) with the additional imposition of the nSFR < 0.1 criterion. The SDSS sample, shown in gray, is at $0.04 < z < 0.14$; the Lega-C sample, shown in blue, is at $0.6 < z < 1.0$. A similar evolution of ~ 0.2 in the average D4000 value can be seen from $z \sim 0.1$ to $z \sim 0.8$ and from $z \sim 0.8$ to the VANDELS sample at an average redshift of $z \sim 1.15$. The cosmic time interval between SDSS and Lega-C is approximately four times that between Lega-C and VANDELS, demonstrating that the distribution of quiescent galaxies moves towards lower D4000 at an accelerating pace with increasing lookback time as expected.

4.5.4 Trends with rest-frame UVJ colours

In this section I consider the positions of the VANDELS galaxies on the UVJ diagram, and physical parameter trends with UVJ colours. The sample is shown on the UVJ diagram in Fig. 4.5.

Trends with star-formation rate

The top-left panel of Fig. 4.5 shows a direct comparison of the nSFR-based selection to both sets of UVJ selection criteria discussed in Section 4.5.1. Good agreement can be seen between the dashed criterion of $U-V > 0.88(V-J) + 0.69$ and the nSFR-based selection introduced in Chapter 2. None of the quiescent sub-sample fall below the dashed line, whereas seven green valley objects fall above. However, all of these objects are very close to the nSFR threshold: for four of the seven the 16th posterior percentile is at nSFR < 0.1.

The top-right panel of Fig. 4.5 shows the VANDELS sample coloured by nSFR. At an observed redshift of $z \sim 1$, a nSFR value of 0.1 is equivalent to $\log_{10}(\mathrm{sSFR}/\mathrm{yr}^{-1}) \simeq -10.5$. A a clear trend in nSFR (and hence sSFR) can

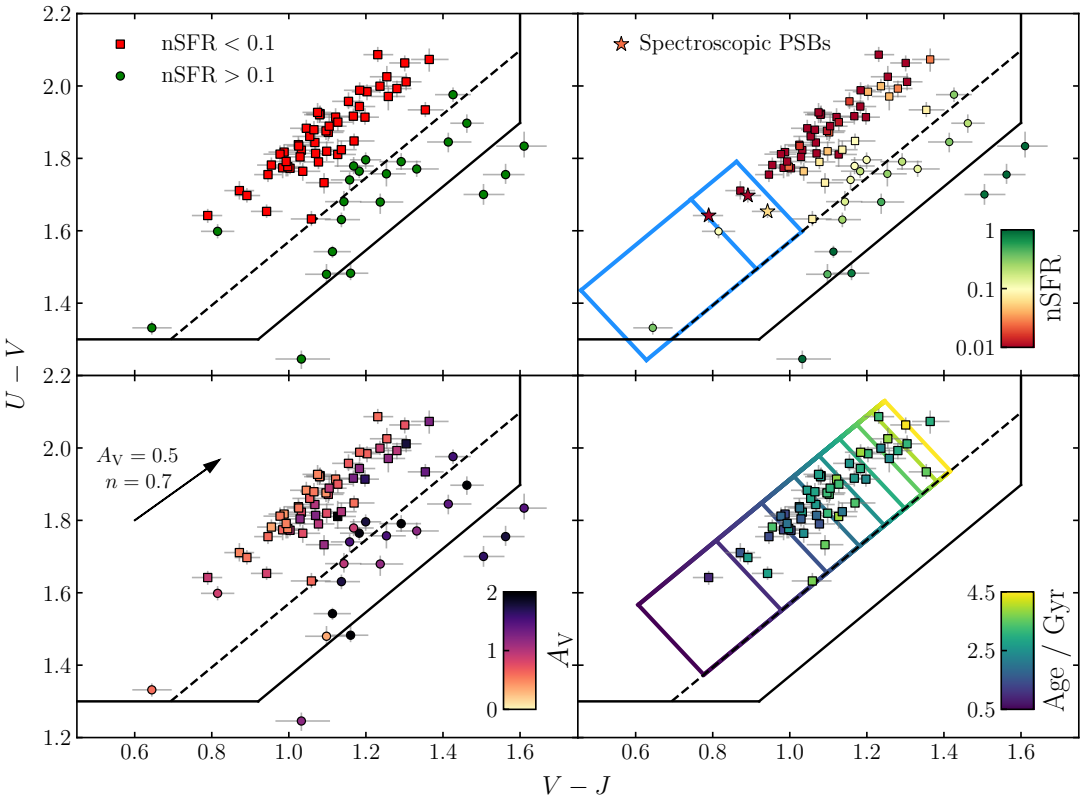

Figure 4.5 *The VANDELS sample on the UVJ diagram, coloured by nSFR, A_V and mass-weighted age. The VANDELS UVJ criteria (see Section 4.2.1) are shown by solid black lines; the dashed line is the stricter boundary of $U - V > 0.88(V - J) + 0.69$. In the top-right panel the larger blue box is the PSB selection of Belli et al. (2019), assuming that galaxies with median stellar ages of* 300–800 *Myr display PSB spectral properties. The upper box shows an extension of the PSB selection to a maximum age of 1.2 Gyr. In the bottom-right panel, the coloured grid shows predicted positions from Belli et al. (2019) for median stellar ages from* 0.5–4.5 *Gyr in 0.5 Gyr intervals.*

be observed perpendicular to the red sequence, in agreement with recent results (e.g. Fang et al. 2018). It is worth noting that the solid box representing the more-permissive UVJ selection criteria includes objects that are forming stars at up to $\sim$ 50−100 per cent of their historical average SFRs, meaning that these criteria alone should not be used to select high-redshift quiescent galaxies.

I therefore argue that the nSFR < 0.1 criterion is the most robust method for selecting quiescent galaxy samples. This method is truly redshift-independent, as it selects galaxies that are forming stars below a fixed fraction of their historical average SFRs. Both a fixed UVJ selection of $U - V > 0.88(V - J) + 0.69$ and sSFR selection proportional to t^{-1} produce results similar to nSFR < 0.1. By contrast, I argue that both the original redshift-dependent UVJ selection criteria, and selection using a fixed sSFR threshold are less appropriate, as they include galaxies that are proportionally more-highly star-forming at higher redshifts.

Trends with dust attenuation

The VANDELS sample is shown coloured by $A_{\rm V}$ in the bottom-left panel of Fig. 4.5. A strong trend in dust attenuation can be seem across the UVJ box, which follows the trend observed in sSFR, *perpendicular* to the dust reddening vector. These results are in accordance with those of Fang et al. (2018), who observe a drop in the dust contents of their star-forming galaxies as they move closer to the red sequence. The majority of objects within the solid UVJ selection box that have $A_{\rm V} > 1$ can be seen to be identified as in the green valley, rather than truly quiescent. This confirms that, within the UVJ box, dust attenuation is more-strongly related to sSFR than position along the dust-reddening vector. The quiescent sub-sample is typically found to be less dusty, however there is still a noticeable trend with distance from the edge of the UVJ selection box, as was found by Belli et al. (2019).

Trends with stellar age

The mass-weighted ages inferred for the quiescent sub-sample (see Section 4.5.2) are shown on the UVJ diagram in the bottom-right panel of Fig. 4.5. Also plotted are ages predicted by the relationship derived by Belli et al. (2019). The lines that run perpendicular to the dashed UVJ selection are lines of constant age, which are shown from 0.5−4.5 Gyr in intervals of 0.5 Gyr. The ages I derive

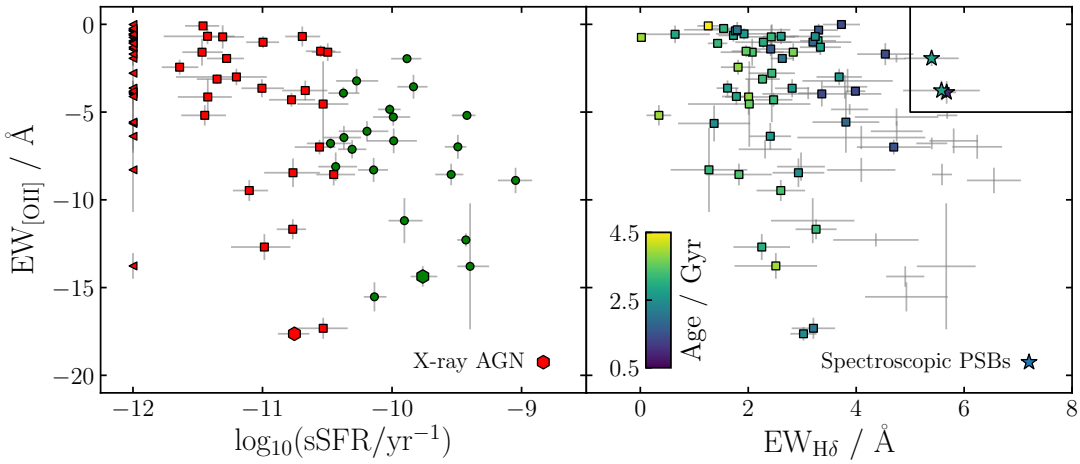

Figure 4.6 *The distribution of [O II] equivalent widths in the VANDELS sample. The left panel shows $EW_{[O\,II]}$ vs sSFR: symbols and colours are as defined in Fig. 4.4, except for galaxies with $log_{10}(sSFR/yr^{-1}) < -12$, which are shown as triangles, indicating an upper limit. Additionally, objects detected in the Chandra seven megasecond catalogue of Luo et al. (2017) are shown as hegaxons (see Section 4.5.7). The right panel shows the $EW_{H\delta}$ vs $EW_{[O\,II]}$ parameter space typically used to select post-starburst galaxies. Quiescent galaxies are coloured by mass-weighted age, green-valley galaxies are shown with gray errorbars. Commonly used spectroscopic PSB selection criteria are shown with solid lines (e.g. Tran et al. 2003; Maltby et al. 2016). Galaxies that meet these criteria are shown as stars (see Section 4.5.6).*

for VANDELS objects are in good agreement with the predictions of Belli et al. (2019), despite several methodological differences (median vs mean stellar age, different SFH models), demonstrating that the ages of quiescent galaxies are less model-dependent than those of star-forming objects.

It is remarkable that such a clear trend in stellar age can exist parallel to the dust reddening vector on the UVJ diagram. Galaxies along the top-left edge of the UVJ distribution can be seen to follow a pure age sequence, with no evolution in dust attenuation. A population of totally quenched objects with younger stellar populations that still retained significantly more dust would disrupt this, as well as the age trend found by Belli et al. (2019). This implies that quenching galaxies must lose most of their dust before their sSFRs drop to the extent where they can join the red sequence. These ideas will be explored further in Section 4.6.2.

4.5.5 [O II] emission properties

As described in Section 4.4.4, the [O II] 3727Å emission line was masked during the spectral fitting analysis due to uncertainties as to the excitation mechanism in quiescent galaxies. I now consider the distribution of [O II] emission in the VANDELS sample. I first measure the rest-frame equivalent width of the line, $EW_{[O\,II]}$, using the same method as was applied to the Hδ feature in Section 4.5.3.

A comparison of measured [O II] equivalent widths with inferred sSFRs is shown in the left panel of Fig. 4.6. The green-valley galaxies typically exhibit stronger [O II] emission, with 77 per cent having $EW_{[O\,II]} < -5$Å. By contrast, only 26 per cent of the quiescent sub-sample has [O II] emission stronger than this threshold. This is despite the dustier nature of the green valley sub-sample (see Section 4.5.4). Whilst a clear trend can be observed between the two sub-samples, it can be seen that at fixed $EW_{[O\,II]}$ the inferred sSFRs span a range as large as ~ 2 dex. This is in agreement with previous results that identify [O II] as a poor predictor of ongoing star-formation activity in quiescent galaxies (e.g. Lemaux et al. 2010).

Without rest-frame optical spectroscopy it is challenging to constrain possible AGN contributions to the observed [O II] fluxes. However, the majority of local quiescent galaxies with detectable [O II] emission have been shown to exhibit high [O II]/Hα ratios (e.g. Yan et al. 2006) that are inconsistent with excitation by ongoing star-formation (Kewley et al. 2004). I therefore identify galaxies in the sample with strong [O II] emission as likely hosts of low-level AGN activity, in particular the quiescent galaxies for which I find $EW_{[O\,II]} < -10$Å. I will further consider the possibility of AGN activity in the VANDELS galaxies in Section 4.5.7.

Ionization by hot low-mass stars has also been postulated as an explanation for line emission in quiescent galaxies (e.g. Singh et al. 2013). Recently, Herpich et al. (2018) reported a marginal difference in the stellar ages of local quiescent galaxies with and without visible emission lines. In principle, this measurement should be easier to make at $z \sim 1$, as the stellar populations of quiescent galaxies are considerably younger. The quiescent sub-sample is shown coloured by stellar age in the right panel of Fig. 4.6. No clear correlation between [O II] emission and stellar age is visible, however the VANDELS sample is considerably smaller than those available in the local Universe, which may preclude the detection of this subtle effect. Future large high-redshift spectroscopic surveys will be a valuable

tool for addressing this issue.

4.5.6 Post-starburst and rejuvenated galaxies

Post-starburst galaxies (PSBs) are widely identified as one of two major transitional states between the star-forming population and red sequence (see Section 4.1). A variety of methods have been used to identify samples of PSBs, ranging from spectroscopic selection based on strong Hδ or Hβ absorption and a lack of emission lines (e.g. Tran et al. 2003), to principal component analyses (e.g. Wild et al. 2007, 2014), to selection by rest-frame UVJ magnitudes (e.g. Belli et al. 2019). Comparisons between different methods have found significant overlap (e.g. Maltby et al. 2016), however a fully self-consistent set of criteria that returns objects with the desired properties is still to be agreed upon.

A detailed discussion of the physical properties of VANDELS PSBs will be presented by Wild et al. in prep. In this section I briefly discuss the distribution of PSBs within the VANDELS sample, and the degree of consistency between different selection methods. The right panel of Fig. 4.6 shows the $EW_{[O\,II]}$ vs $EW_{H\delta}$ parameter space often used to spectroscopically select PSBs (e.g. Tran et al. 2003; Maltby et al. 2016, submitted). I identify three spectroscopic PSBs, all of which are members of the quiescent sub-sample. These objects are marked with stars in Figs 4.3, 4.4, 4.5 and 4.6.

The spectroscopically identified PSBs are shown on the UVJ diagram in the top-right panel of Fig. 4.5. It can be seen that the PSBs occupy the region predicted by Wild et al. (2014), towards the bottom-left of the UVJ selection box. The larger of the two blue boxes is the PSB selection used by Belli et al. (2019), which assumes that PSB features are visible for quiescent galaxies with median stellar ages of 300−800 Myr. Both of the objects found within this box are part of the green-valley sub-sample and, whilst both are close to the PSB selection box in Fig. 4.6, neither fulfils the PSB selection criteria. The object close to the top of the box has strong Hδ absorption, however it has $EW_{[O\,II]} = -6.8 \pm 0.3$Å. By contrast, the object at the bottom of the box fulfils the [O II] criterion but has $EW_{H\delta} \sim 3$Å.

The objects spectroscopically identified as PSBs occupy a region slightly further up the red sequence, within the region for which Belli et al. (2019) predict ages of 800−1200 Myr. This region is highlighted with a smaller blue box above the

one used by Belli et al. (2019). These results suggest this is the region in which spectroscopic features usually associated with PSBs are strongest. The fourth object within the extended box, as well as the closest object above this box, have [O II] consistent with the spectroscopic selection criteria, but slightly weaker Hδ absorption. These results are in good agreement with Maltby et al. (2016), who find that ~ 50 per cent of galaxies identified by the photometric selection proposed by Wild et al. (2014), upon which the Belli et al. (2019) UVJ selection is based, exhibit [O II] and Hδ equivalent widths consistent with the spectroscopic criteria shown in Fig. 4.6.

These findings suggest that the contribution of the PSB quenching channel to the growth of the red sequence is towards the upper end of the range found by Belli et al. (2019). However, as will be discussed in Section 4.6.2, the timescale over which PSB features are visible is not necessarily determined by the time taken for a galaxy to traverse the blue box on Fig. 4.5 by passive evolution of its stellar population.

I finally note, based upon Fig. 4.4, the presence of one galaxy in the green valley sub-sample that displays D4000 > 1.5, but which has a blue continuum below this wavelength, and strong [O II] emission. For this object I find a stellar mass of $\log_{10}(\boldsymbol{M}_*/\mathrm{M}_\odot) = 11.25 \pm 0.15$ and a mass-weighted age of 3.6 ± 0.5 Gyr, amongst the oldest and most massive in the VANDELS sample. With currently available data it is challenging to discriminate between rejuvenated star-formation and AGN activity (though I find no evidence of AGN activity in this object using either X-ray or radio datasets, as discussed in Section 4.5.7). However, in either case, objects such as this are clearly of significant interest for assisting our understanding of continuing mass assembly in the oldest galaxies (e.g. Belli et al. 2017; Nelson et al. 2018). Assuming no AGN contribution I infer a SFR of $8.9^{+4.9}_{-3.2}$ $\mathrm{M}_\odot$ yr^{-1}, meaning the stellar mass of this object would increase by ~ 2–5 per cent over 100 Myr.

4.5.7 Evidence of AGN activity

As discussed in Section 4.5.5, it is challenging to determine whether or not galaxies in the sample host an AGN, given only UV-NIR photometry and rest-frame UV spectroscopy. Whilst line ratios from rest-frame optical spectroscopy would be the ideal tool for studying this, I here consider ancillary datasets in the X-ray, mid-infrared and radio to attempt to constrain any AGN contributions.

I first consider X-ray data from the *Chandra* seven megasecond source catalogue (Luo et al. 2017) and Subaru/*XMM-Newton* deep survey (Ueda et al. 2008). I find two matches within the VANDELS sample, both from the *Chandra* seven megasecond source catalogue. These objects are shown as hexagons in the left panel of Fig. 4.6, and can be seen to have two of the four strongest [O II] lines in the sample (the second highest is in the UDS). This suggests that strong [O II] emission is a good predictor of X-ray AGN activity in quiescent galaxies.

I then consider radio data at 1.4GHz from the Very Large Array, which is available for both fields (Simpson et al. 2006; Bonzini et al. 2013). I find that three of the VANDELS galaxies are detected, none of which are in common with the X-ray detected objects. All three are massive, with $\log_{10}(\boldsymbol{M}_*/\mathrm{M}_\odot) > 11$, and all three sit close to the nSFR threshold, with $\log_{10}(\mathrm{sSFR}/\mathrm{yr}^{-1}) \simeq -10.5$. There is no indication from their rest-frame UV continua or [O II] lines that these galaxies host AGN, as has typically been found for radio galaxies at high redshift (e.g. Dunlop et al. 1996).

I finally consider publicly available *Spitzer* MIPS 24μm imaging, which is available for 57 out of 75 galaxies (Dickinson et al. 2003a; Dunlop et al. 2007). For each of these objects I extract fluxes within 14″−diameter apertures and manually inspect the images, flagging objects as isolated (33 objects) or potentially confused (24 objects). Of the isolated sources, 21 are members of the quiescent sub-sample, and 12 are members of the green valley sub-sample.

Only one of the two objects identified above as X-ray sources has MIPS 24μm coverage: the galaxy in the green valley sub-sample. This object has the strongest isolated detection in the sample by approximately a factor of two. Of the 21 isolated quiescent galaxies, only one has a detection at $>3\sigma$, whereas 7 out of 12 isolated green valley objects are detected above this threshold. For the eight objects with robust, isolated detections I use the calibration of Kennicutt & Evans (2012) to convert their aperture-corrected fluxes into SFRs. I find that the X-ray detected source is the only object with significantly more 24μm-flux than expected, based on the SFRs I derive from my spectral fitting analysis.

The quiescent object that is detected at 24μm is one of the three identified as radio AGN above. This object is the most massive in the sample, with $\log_{10}(\boldsymbol{M}_*/\mathrm{M}_\odot) = 11.66 \pm 0.14$, and one of the oldest, with a mass-weighted age of 3.5±0.6 Gyr. However, I find no evidence of an AGN contribution to its 24μm flux: the inferred SFR is $13.2^{+5.6}_{-3.8}$ $\mathrm{M}_\odot$ yr^{-1}, whereas its 24μm-predicted SFR is 14.3 ± 2.5 $\mathrm{M}_\odot$ yr^{-1}.

This object is similar to the rejuvenated object discussed in Section 4.5.6, and seems consistent with the scenario proposed by Best et al. (2014), in which cooling of gas in the hot halo begins to provide fuel for jet-mode AGN activity and rejuvenated star-formation $\gtrsim 2$ Gyr after quenching has occurred.

The two 1.4GHz catalogues I consider jointly provide coverage down to 100μJy sensitivity for all of the VANDELS objects. This corresponds to a rest-frame 1.4GHz luminosity of $10^{23.6}$ W Hz^{-1} at $z = 1$, or $10^{23.9}$ W Hz^{-1} at $z = 1.3$. I find that 2 of the 4 objects with $\log_{10}(M_*/\mathrm{M}_\odot) > 11.5$ are detected above this level, whereas only 1 of the 39 objects with $11.0 < \log_{10}(M_*/\mathrm{M}_\odot) < 11.5$ is detected. This is broadly consistent with the local relationship between radio-loud AGN fraction and stellar mass (e.g. Best et al. 2005; Sabater et al. 2019), in agreement with other studies, which find little redshift evolution (e.g. Tasse et al. 2008; Simpson et al. 2013).

4.5.8 Stacking analyses

I finally perform a stacking analysis to demonstrate the average spectral properties of the galaxies in the two sub-samples. I de-redshift the VANDELS spectra and normalise over the rest-frame wavelength range from 3200−3600Å, which is in the centre of the observed spectral range and contains no strong features. I then resample the spectra to a common wavelength grid using Spectres (Carnall 2017). Median stacked green valley and quiescent spectra are shown in the top and bottom panels of Fig. 4.7 respectively. The stacks have total exposure times of 745 and 1971 hours respectively.

A clear transition from Balmer to 4000Å break can be seen between the two spectra, demonstrating the older stellar populations of the quiescent galaxies. As discussed in Section 4.5.5, the median [O II] flux is significantly higher in the green valley sub-sample. The continuum below 4000Å can also be seen to be bluer, with considerably more flux at < 3000Å, indicating higher levels of ongoing star-formation.

Median stacked 24μm images are also shown in Fig. 4.7, using all 57 objects for which MIPS coverage is available (see Section 4.5.7). The green valley stack shows a clear detection, with no corresponding detection in the quiescent stack.

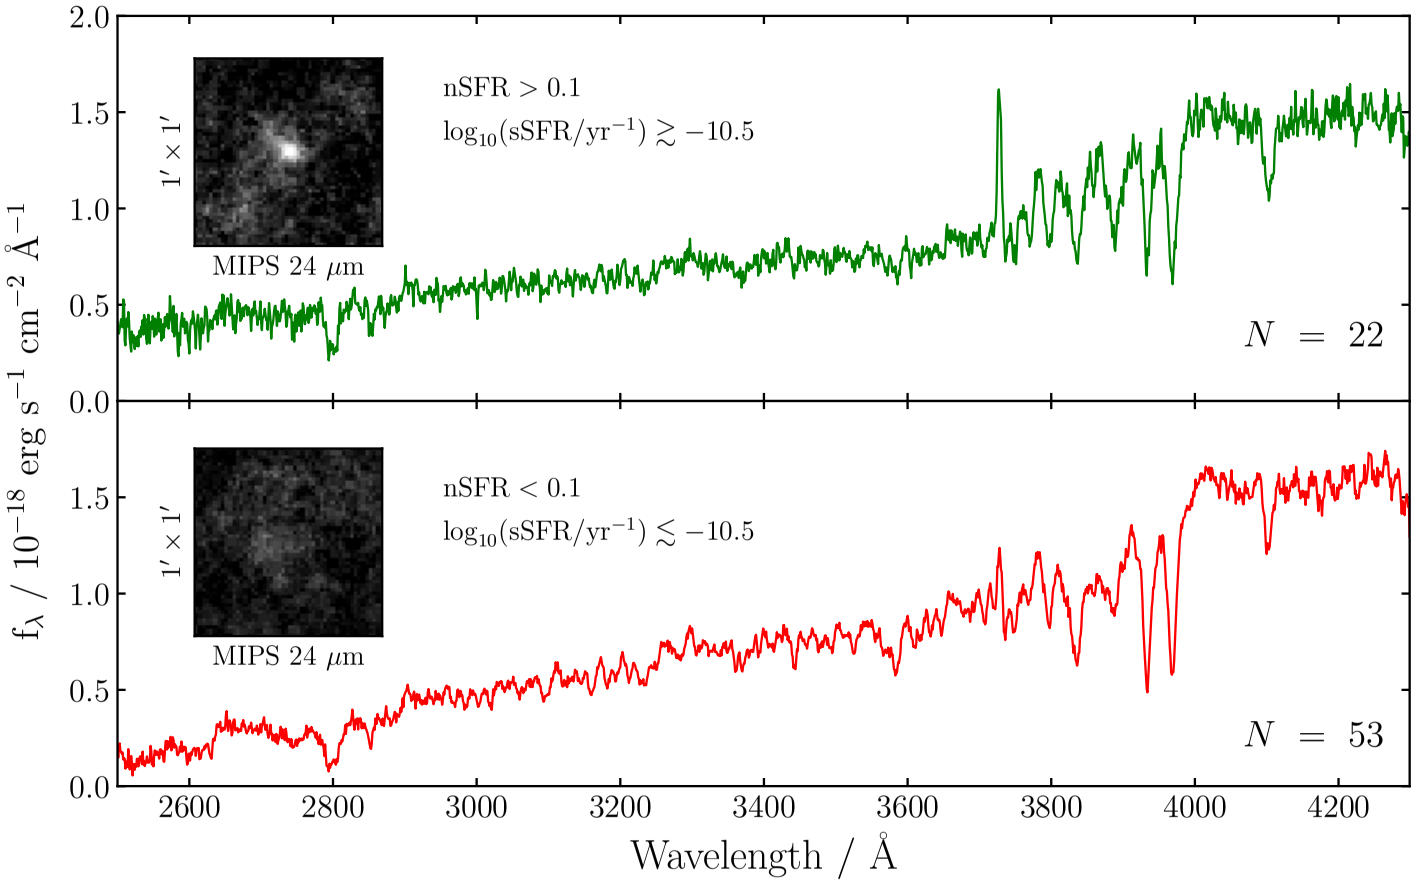

Figure 4.7 *Stacked spectra for the green-valley and quiescent sub-samples. Key differences include more flux below 3000Å and stronger [OII] emission for the green-valley galaxies, and a transition from Balmer to 4000Å break. Stacked 24 μm images are also shown, which include all 57 objects with MIPS coverage. The green valley stack shows a clear detection, consistent with greater ongoing star-formation.*

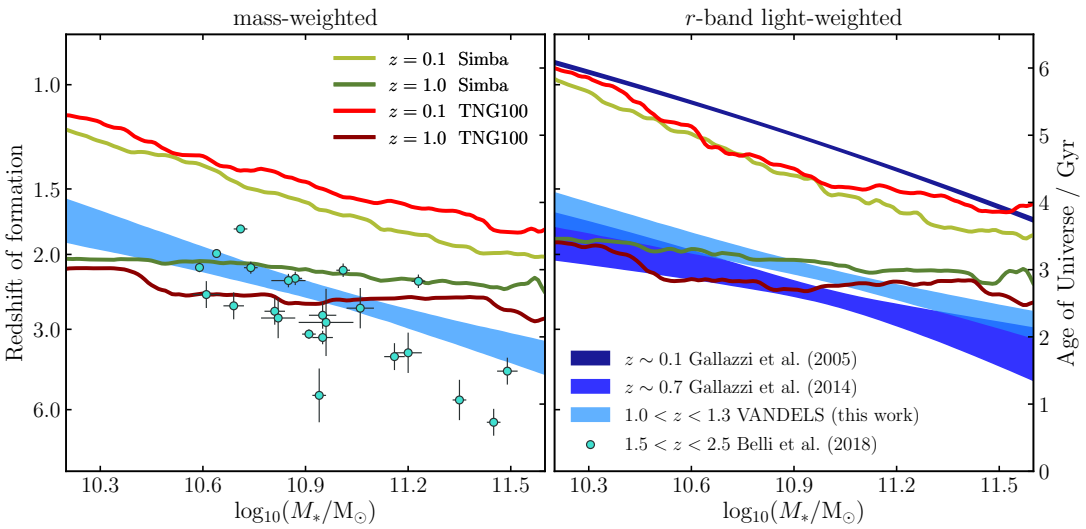

Figure 4.8 *A comparison of quiescent galaxy formation redshifts from spectroscopic studies and simulations at a range of observed redshifts. Formation redshifts in the left panel are mass-weighted, those in the right panel are r-band light-weighted. Results at $z = 0.1$ are measured within 3″−diameter circular apertures. Observational results at higher redshifts were obtained using 1″ slits. Results from simulations at $z = 1$ were extracted using 1″ square apertures. Systematic offsets exist in the vertical positions of the different observational relationships (e.g. due to the use of different SFH models), however the gradients can be seen to be in good agreement.*

4.6 Discussion

In this section I further discuss my results. In Section 4.6.1 I compare the stellar mass vs age relationship of Section 4.5.2 to results from the literature and recent simulations. In Section 4.6.2, I consider the connections between the quiescent, green-valley and post-starburst galaxies in the VANDELS sample. In Section 4.6.3 I discuss the shapes of the SFHs I infer for the quiescent sub-sample.

4.6.1 The stellar mass vs stellar age relationship

As discussed in Section 4.1, the epoch of formation as a function of galaxy stellar mass is a key observable property, which is strongly constraining on AGN-feedback models. In this section I compare my results to other observational studies, as well as making comparisons with simulations.

The stellar mass vs stellar age relationship I derive for the quiescent sub-sample is compared to results from the literature in Fig. 4.8. In the left panel the relationship between stellar mass and mass-weighted formation time derived in Section 4.5.2 is shown. The shaded region shows the 16^{th}–84^{th} percentiles of the posterior distribution given in Equation 4.9.

In order to facilitate comparisons with earlier work, I also calculate the same relationship using r-band light-weighted formation times, $t_{r\text{-band}}$ (e.g. Gallazzi et al. 2005). For this quantity I obtain a relationship of

$$\left(\frac{t_{r\text{-band}}}{\text{Gyr}}\right) = 2.91^{+0.08}_{-0.09} - 1.24^{+0.27}_{-0.30} \log_{10}\left(\frac{M_*}{10^{11}\text{M}_\odot}\right) \tag{4.10}$$

with an intrinsic scatter of $0.51^{+0.08}_{-0.07}$ Gyr. The slope of this relationship is slightly shallower than that which I obtain for t_{form} (Equation 4.9), and is offset towards later formation times by ~ 350 Myr at $\log_{10}(M_*/\text{M}_\odot) = 11$.

In the left panel of Fig. 4.8, I also show the 23 galaxies studied by Belli et al. (2019) at observed redshifts of $1.5 < z < 2.5$. The slope of the relationship can be seen to be similar in both studies, with the results of Belli et al. (2019) offset towards earlier formation by ~ 0.5 Gyr. I have shown in Section 4.5.4 that my inferred ages agree well with the predictions of Belli et al. (2019). This means that no significant systematic offsets should exist between the two samples. I therefore conclude that the relationship given in Equation 4.9 evolves towards earlier formation by ~ 0.5 Gyr between the VANDELS sample at $z \sim 1.15$ and the mean redshift of the Belli et al. (2019) sample at $z = 1.75$. This is inconsistent with pure-passive evolution across this redshift interval, in which case the formation redshifts should remain constant.

In the right panel, r-band light-weighted formation times for the quiescent sub-sample are compared to those found by Gallazzi et al. (2005, 2014) for quiescent galaxies at $z \sim 0.1$ and $z \sim 0.7$ respectively. Again, good agreement can be seen between the slope inferred for the quiescent sub-sample and those found by Gallazzi et al. (2005, 2014). I thus conclude that, at $\log_{10}(M_*/\text{M}_\odot) > 10.3$, an evolution of ~ 1.5 Gyr in formation time per decade in stellar mass is a robust result, which remains consistent from the local Universe to at least $z \sim 2$.

However, the relative vertical positions of the VANDELS relationship and the Gallazzi et al. (2014) relationship at $z \sim 0.7$ do not follow the expected trend of earlier formation with increasing observed redshift at fixed stellar mass. Both datasets are observed using 1″ slits, meaning that this difference is unlikely to be due to aperture bias. Instead, I identify methodological differences as the most likely cause. Gallazzi et al. (2014) fit indices derived from their rest-frame optical spectra, and use an exponentially declining SFH model. In this chapter I have applied a full-spectral-fitting approach to rest-frame UV spectra using a double-power-law SFH model. The use of different SFH models and priors is known to affect derived ages, which could plausibly give rise to this effect (see Chapter 3). It is additionally likely that the optical spectra of quiescent galaxies are dominated by an older population than dominates in the UV, meaning that fitting different spectral regions could result in different SFHs. A joint analysis of rest-frame UV and optical spectroscopy for the same sample would be extremely valuable for understanding the magnitude of this effect.

Stellar mass vs age from photometric studies

Several recent studies also attempt to probe the stellar mass vs age relationship using photometric data. This has the advantage of providing better statistics, as larger samples are available. However, as discussed in Section 4.1, the age-metallicity-dust degeneracy leads to larger uncertainties on individual age measurements. In this section I compare two recent photometric studies to the spectroscopic studies discussed in Section 4.6.1.

Pacifici et al. (2016) consider a sample of 845 objects with multi-band photometry from CANDELS, whereas, in Chapter 2, I consider a sample of 9289 galaxies from UltraVISTA (McCracken et al. 2012). The results of these two studies are summarised in Fig. 4.9. I derive $t_{\rm form}$ for the stacked SFHs of Pacifici et al. (2016) using the best-fitting double-power-law parameters reported in their table A1. Despite their very different methodologies, these studies can be seen to produce similar results. However, the slopes reported are considerably shallower than those shown in Fig. 4.8, at $\lesssim 0.5$ Gyr per decade in stellar mass.

This inconsistency can be attributed to the increased uncertainties on individual object SFHs when considering photometric data. In the presence of large uncertainties, a population of very massive, very old objects will be preferentially scattered towards later formation times. This is due to the constraint imposed by

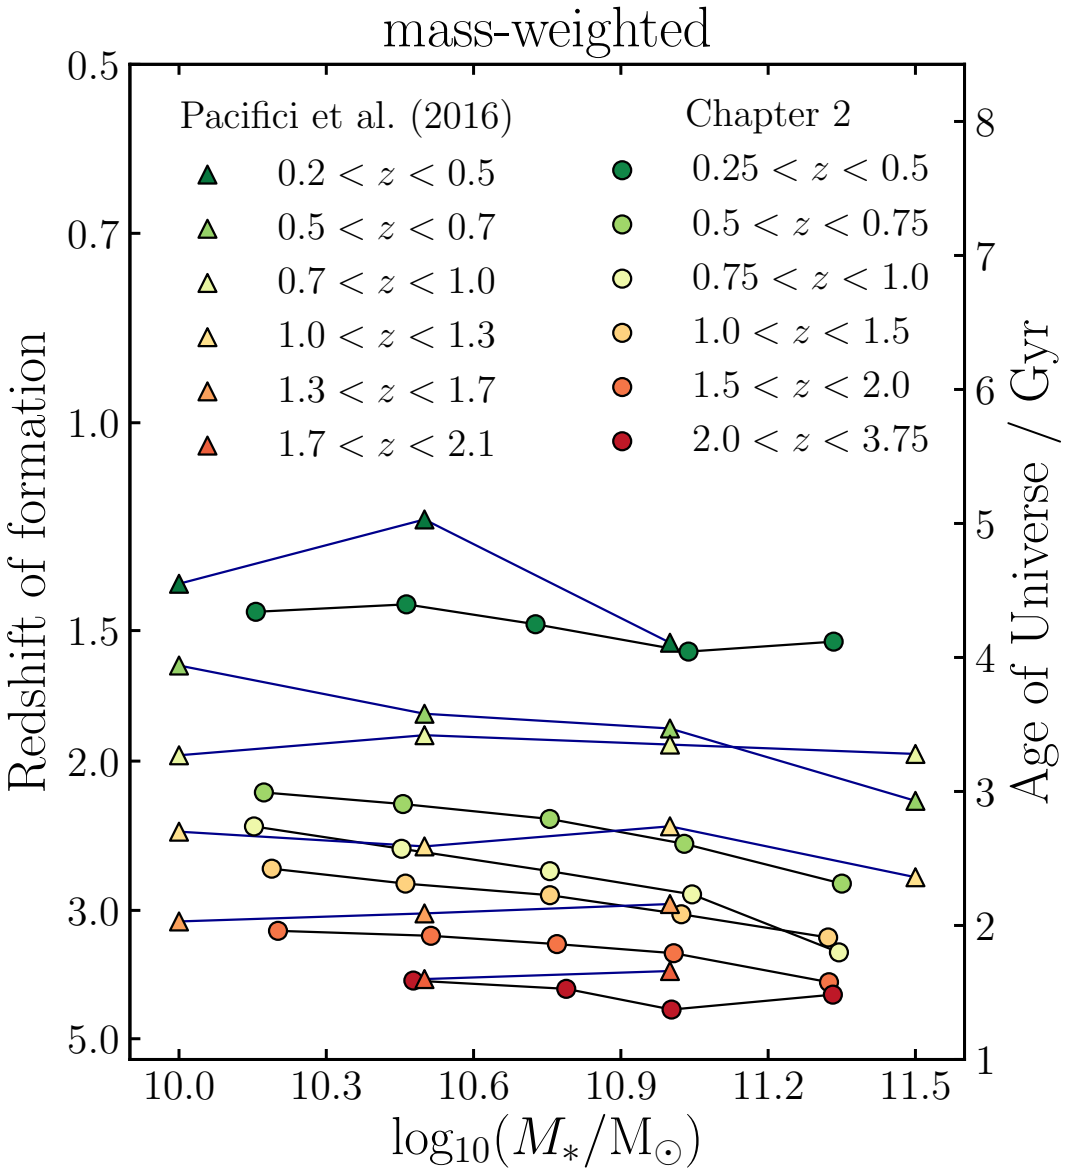

Figure 4.9 *A comparison of formation redshifts for massive quiescent galaxies from two recent photometric studies: Pacifici et al. (2016), and the study presented in Chapter 2. Whilst these two sets of results are in good agreement, considerably weaker stellar mass vs stellar age trends are recovered when compared to the spectroscopic studies shown in Fig. 4.8.*

the age of the Universe. Similarly, a population of younger, less massive quiescent galaxies will be preferentially scattered towards earlier formation times, as the time evolution of galaxy stellar populations is much more rapid at younger ages, meaning later formation times will be strongly inconsistent with the observed spectrum.

Both of these effects act to flatten the trends observed in Fig. 4.9 when compared to the spectroscopic analyses in Fig. 4.8. This finding highlights the importance of forthcoming large spectroscopic surveys, and analyses such as that presented in Section 4.4, in providing stronger constraints on the ages of individual objects than are available from photometric data alone (e.g. Pacifici et al. 2012; Thomas et al. 2017).

Stellar mass vs age in cosmological simulations

Given the conclusion in Section 4.6.1 that the strong stellar mass vs stellar age relationship observed for quiescent galaxies at low redshift is already in place by $z \sim 2$, it is interesting to consider whether this trend is reproduced by modern cosmological simulations. Historically, this relationship has been challenging to match, even in the local Universe, for both quiescent and star-forming galaxies (e.g. Somerville et al. 2008; Trager & Somerville 2009).

I consider predictions from the 100 h^{-1} Mpc box length runs of both Simba (Davé et al. 2019) and IllustrisTNG (e.g. Nelson et al. 2019), using snapshots at $z = 0.1$ and 1.0. In order to match these predictions as closely as possible with the observational studies discussed in Section 4.6.1, I apply apertures to the simulated galaxies as follows. For the $z = 0.1$ snapshots, I apply $3''-$diameter circular apertures to each galaxy, for consistency with the SDSS observations used by Gallazzi et al. (2005). For the $z = 1$ snapshots, I apply $1''$ square apertures, to match the $1''$ slits used by Gallazzi et al. (2014), Belli et al. (2019) and VANDELS (assuming a $\sim 1''$ region along the direction of the slit is extracted). I select all galaxies within these snapshots that meet the nSFR < 0.1 criterion (see Section 4.5.2), then calculate mass-weighted and r-band light-weighted formation times for the simulated galaxies. I use Bagpipes to predict the r-band flux from each star particle using their individual ages, stellar masses and metallicities.

In Fig. 4.8 I show the mass-weighted and r-band light-weighted formation times for each snapshot. At each point, the median formation time of galaxies within

a 0.25 dex mass range centred on that point is shown. The two simulations can be seen to be in reasonably good agreement, with discrepancies confined to levels of $\lesssim$ 250 Myr. The z = 0.1 relationships predicted by these simulations have slopes consistent with the $\sim$ 1.5 Gyr per decade in mass evolution found by the observational studies shown. Additionally, the normalisations of the z = 0.1 relationships in both simulations are in good agreement with the results of Gallazzi et al. (2005), as found for the general $z < 0.2$ galaxy population in IllustrisTNG by Nelson et al. (2018).

However, at $z = 1$ both simulations predict significantly weaker stellar mass vs age relationships. Using mass-weighted formation times, Simba predicts an evolution of 0.38 Gyr per decade in stellar mass across the interval shown in Fig. 4.8, whereas IllustrisTNG predicts an evolution of 0.51 Gyr per decade. This suggests that these simulations do not accurately reproduce the detailed physical properties of massive quiescent galaxies at $z > 1$. This is particularly interesting in the context of the results of Schreiber et al. (2018), who find that the precursors to these simulations (Mufasa and Illustris) significantly under-predict the number density of quiescent galaxies at $3 < z < 4$.

4.6.2 Connecting green-valley, post-starburst and quiescent galaxies

Much debate exists as to how galaxies evolve away from an initial position within the star-forming population, potentially through green-valley or post-starburst phases, to eventually join the red sequence. Several recent studies have considered this problem in the context of evolutionary tracks across the UVJ diagram (e.g. Belli et al. 2019; Morishita et al. 2018). One of the most challenging aspects of this approach is modelling the time evolution of dust attenuation, which has a significant impact on these evolutionary tracks.

Belli et al. (2019) consider the number of galaxies that pass through the larger blue PSB selection box drawn on the top-right panel of Fig. 4.5, arguing for separate fast and slow quenching mechanisms that do and do not pass through the box respectively. They find that fast quenching plays a more-important role at high redshift, in accordance with previous work (e.g. see Chapter 2; Wild et al. 2009, 2016; Schawinski et al. 2014; Pacifici et al. 2016). A toy model to describe both fast and slow quenching routes is presented in their fig. 12, assuming that

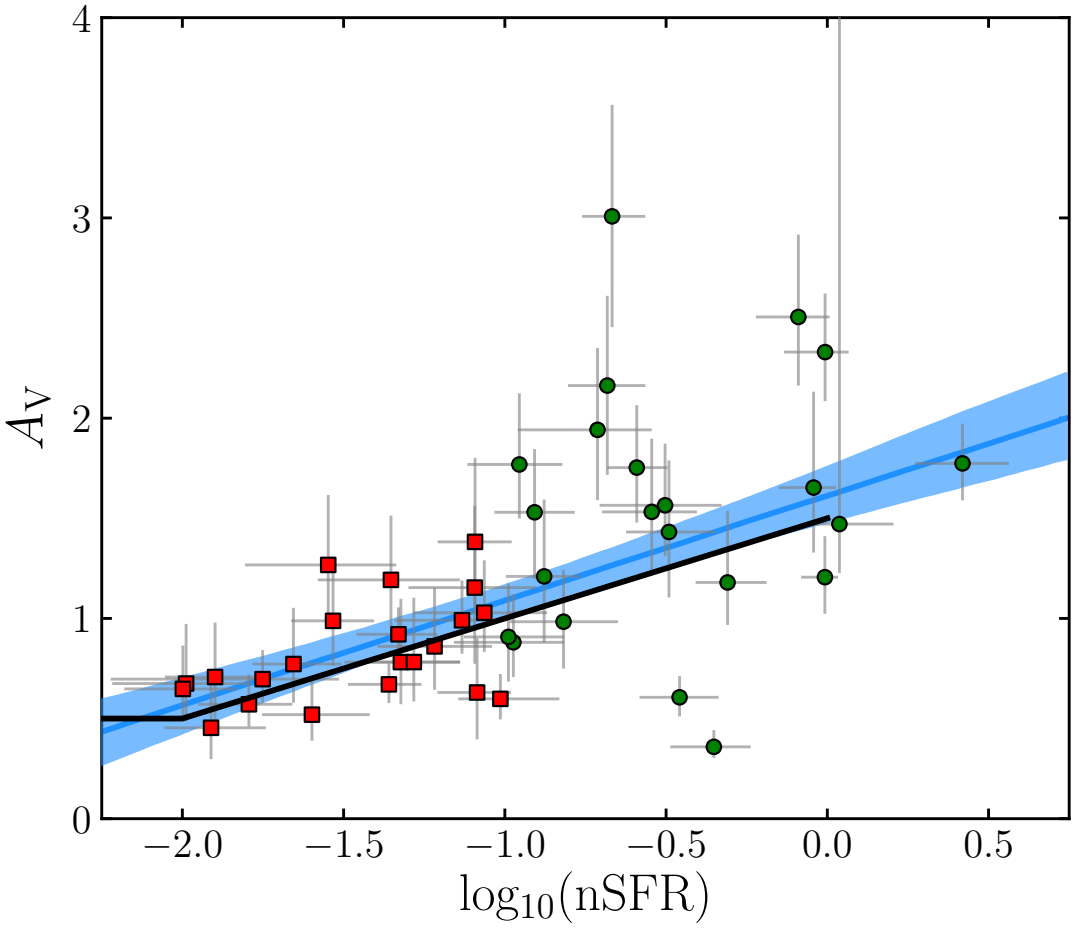

Figure 4.10 *Dust attenuation vs nSFR for galaxies in the VANDELS sample with detectable star-formation (nSFR > 0.01). The blue line shows the posterior median linear relationship reported in Equation 4.11, whereas the black line shows the ansatz used in Fig. 4.11.*

$A_V \propto$ SFR. For their exponentially declining SFH models this means galaxies lose most of their dust early on, and the tracks therefore approach the UVJ selection box from the bottom-left part of the diagram.

However, for the VANDELS green-valley sub-sample I find relatively high dust attenuation ($A_V \sim 1{-}2$) even very close to the dashed UVJ boundary in Fig. 4.5. This implies that these objects evolve differently to the scenarios proposed by Belli et al. (2019), approaching the UVJ selection box from further upwards and to the right on the UVJ diagram. This is expected, given their high masses and the fact that A_V is more-strongly correlated with stellar mass than SFR in star-forming galaxies (e.g. Garn & Best 2010; McLure et al. 2018b). In this section I discuss quenching scenarios for green-valley galaxies.

Modelling the evolution of UVJ colours

I begin by considering the time evolution of dust attenuation. As discussed in Section 4.5.4, strong trends in both nSFR and A_V perpendicular to the dashed UVJ selection can be observed in Fig. 4.5, meaning nSFR and A_V are correlated. This relationship is shown in Fig. 4.10 for galaxies with detectable levels of star-formation (nSFR > 0.01; see Section 4.5.2). I fit a linear relationship as described in Section 4.5.2, obtaining

$$A_V = 0.52^{+0.12}_{-0.11} \log_{10}(\mathrm{nSFR}) + 1.61^{+0.15}_{-0.13} \tag{4.11}$$

with an intrinsic scatter of $0.37^{+0.06}_{-0.05}$ magnitudes. At lower nSFR I find that A_V reaches a minimum value of ~ 0.5 as discussed in Section 4.5.4. At higher nSFR I do not have galaxies in the sample to test whether this relationship holds. In the discussion that follows I use a simple empirical ansatz for A_V, consistent with Equation 4.11. I assume that $A_V = 0.5 \log_{10}(\mathrm{nSFR}) + 1.5$ for $0.01 < \mathrm{nSFR} < 1$, then that A_V remains constant at 0.5 for lower nSFR (shown as a black line in Fig. 4.10). For simplicity I assume the Calzetti et al. (2000) dust attenuation law, with attenuation doubled towards stars formed in the last 10 Myr, as in Section 4.3.2.

Armed with this relationship, I construct four representative double-power-law SFHs (see Equation 4.1) to describe a range of quenching scenarios. I consider "early quenching", which occurs at $z \sim 2$, and "late quenching" at $z \sim 1$. For each of these two scenarios I consider fast and slow quenching paths, which have

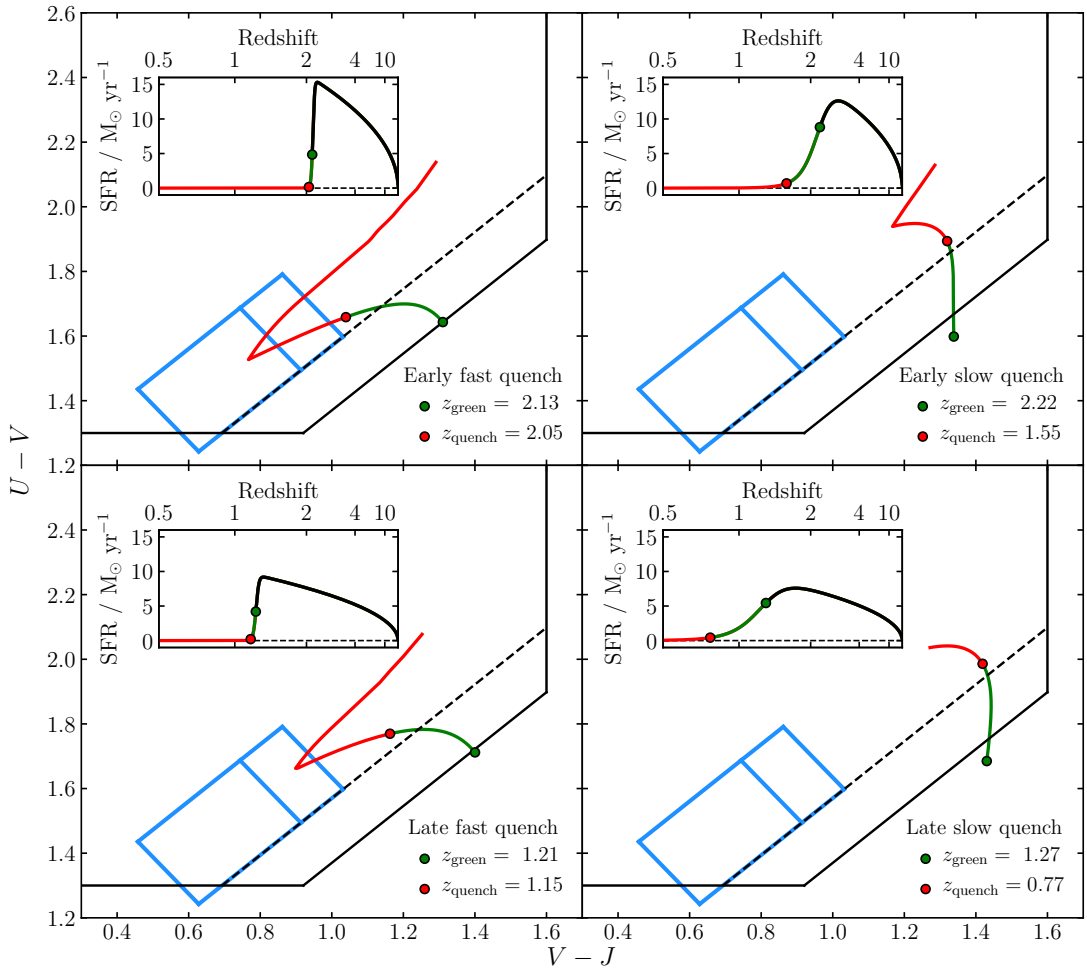

Figure 4.11 *UVJ tracks for four representative galaxy SFHs introduced in Section 4.6.2, assuming the relationship between dust attenuation and nSFR shown in black in Fig. 4.10. The SFH for each object is shown in the inset panel. The solid and dashed black lines, as well as the blue PSB selection box, are as described in the caption of Fig. 4.5. The green points show the redshifts at which these objects enter the green valley (nSFR = 1; when they begin to be tracked). The red points show the redshifts at which they quench (nSFR = 0.1).*

timescales of ~ 100 Myr and ~ 1 Gyr respectively, in accordance with Belli et al. (2019). The parameters of the four models are

- Early fast quench: $\tau = 3$ Gyr, $\beta = 0.5$, $\alpha = 100$
- Early slow quench: $\tau = 3$ Gyr, $\beta = 0.5$, $\alpha = 10$
- Late fast quench: $\tau = 5$ Gyr, $\beta = 0.5$, $\alpha = 100$
- Late slow quench: $\tau = 5$ Gyr, $\beta = 0.5$, $\alpha = 10$.

I model the evolution of the UVJ colours of these four galaxies forwards from the point at which their nSFR first falls below 1. this can be thought of as the time at which they enter the green valley, as this is roughly the highest nSFR found in the VANDELS green valley sub-sample (see Fig. 4.10).

Relating the green-valley, post-starburst and quiescent populations

The tracks the four mock galaxies follow across the UVJ diagram are shown in Fig. 4.11. Their SFHs are shown in the inset panels. I highlight with green and red points the redshifts at which nSFR falls below 1 and 0.1 respectively, corresponding to the times at which these galaxies enter the green valley and then quench. It can be seen that each of these objects at first follows a curving path, the shape of which is determined by both quenching speed and the duration of star-formation activity before quenching begins (as well as depending strongly on the assumptions made regarding dust attenuation). At the end of this curving track, galaxies begin to passively age along a straight path towards the upper right of the UVJ box. The "late slow quench" model has not reached this point by $z = 0.5$.

Galaxies that quench rapidly can be seen to briefly pass through the blue PSB selection box, both entering and leaving by the top-right edge. This picture is distinct from the PSB evolution channel of Belli et al. (2019), which enters the PSB box from the lower-left edge. This is a consequence of different assumptions regarding the evolution of dust: I assume that galaxies lose their dust rapidly during quenching, whereas in the Belli et al. (2019) scenario dust attenuation is already low on approach to the green valley. How far the rapidly quenched galaxies enter into the PSB box depends on how extended in time star-formation

activity is prior to quenching, as well as quenching speed. At high redshift, when star formation cannot be very extended, galaxies fall further into the box, and spend more time in the PSB phase. At lower redshifts, the older stellar populations present in galaxies prevent their UVJ colours becoming blue enough to fall as deep into the PSB box.

I interpret these findings in the context of the results of Wild et al. (2016), Almaini et al. (2017) and Maltby et al. (2018), who suggest a dual origin for the post-starburst population. High-redshift PSBs primarily follow a UVJ evolution similar to that suggested by Belli et al. (2019). They form in a brief, intense starburst, as suggested by Wild et al. (2016), but either never build up significant dust or shed their dust rapidly as star-formation begins to fall. This could be a consequence of strong AGN-driven outflows at high-redshift (e.g. Maiolino et al. 2012; Cimatti et al. 2013), which have been linked to $z \sim 1$ PSBs (Maltby et al. submitted). These galaxies enter the PSB box by the bottom-left (or possibly bottom-right) edge, and exit by the top-right. The morphological evidence presented by Almaini et al. (2017) and Maltby et al. (2018) supports this picture, with high-redshift PSBs found to be extremely compact, consistent with a major disruption event such as a merger.

By contrast, lower-redshift PSBs primarily form by the rapid quenching of normal star-forming galaxies (Wild et al. 2016), which shed their dust more slowly, following a UVJ evolution similar to that shown in the bottom-left panel of Fig. 4.11. These objects briefly enter and leave the PSB box by the top-right edge. This is again consistent with Maltby et al. (2018), who find that these low-redshift PSBs have less-concentrated structures, more similar to ordinary star-forming galaxies. At progressively lower redshifts, the progressively older stellar populations present in these objects prevent them from falling as far into the PSB box, explaining the decreasing number density of PSBs with redshift. This also explains the clustering of galaxies close to the top-right edge of the PSB box (e.g. see fig. 10 of Belli et al. 2019), which becomes more pronounced at lower redshifts.

A final piece of evidence for this scheme comes from the individual SFHs inferred for the three spectroscopic PSBs identified in Fig. 4.6. The posterior SFHs inferred for these objects are shown in Fig. 4.12. For two objects the SFHs are very extended before rapidly quenching. These objects are consistent with the cusp of the track shown in the bottom-left panel of Fig. 4.11. Whilst rapidly quenched, these objects are not, in a literal sense, "post-starburst". By contrast,

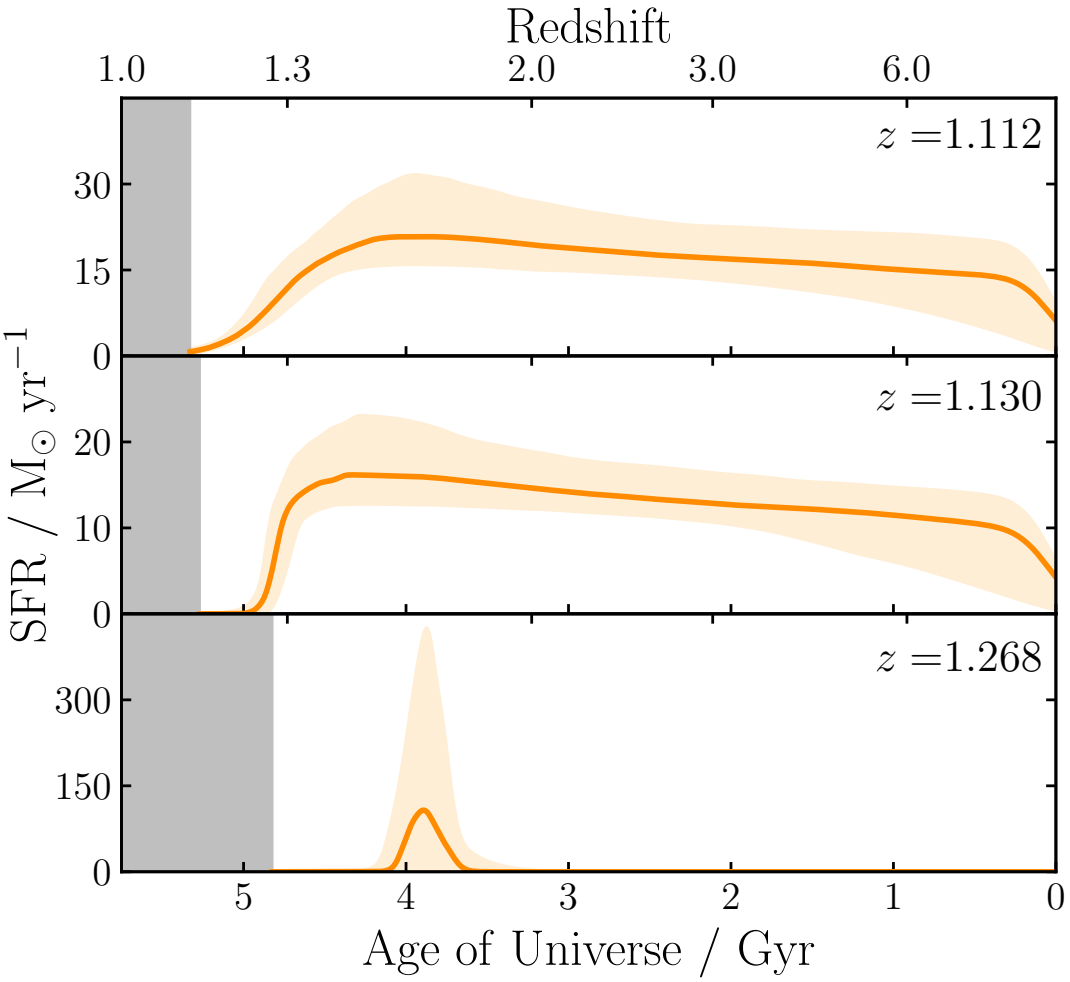

Figure 4.12 *SFHs inferred for the three spectroscopic PSBs identified in Fig. 4.6. Time after the observation of each galaxy is shaded gray. Two have SFHs that are significantly extended before rapidly quenching; the third is a literal post-starburst, formed at $z \sim 1.6$. The first two are consistent with the "late fast quench" scenario shown in Fig. 4.11. The third probably evolves similarly to the PSB track shown in fig. 12 of Belli et al. (2019).*

the third object *is* a literal post-starburst, with almost all of its stellar mass formed within a $\lesssim 500$ Myr period, approximately 1 Gyr before it is observed. This object has the youngest mass-weighted age of the galaxies in the quiescent sub-sample. It can be seen in Fig. 4.5 to the extreme lower left of the quiescent population, and is the spectroscopic PSB closest to the Belli et al. (2019) selection box. It also has the strongest Hδ absorption within the quiescent sub-sample (see Fig. 4.6). Further discussion of the SFHs and physical properties of VANDELS $z \sim 1$ PSBs will be presented by Wild et al. in prep.

The model galaxies that quench more slowly can be seen to follow a more conventional path in Fig. 4.11, crossing the green valley approximately perpendicular to the edge of the UVJ selection box, and joining the red sequence higher up. These tracks are plausibly linked to objects higher up the red sequence in Fig. 4.5 that still retain higher A_V. However, most of their dust has already been lost by the time they cross the dashed line, preserving the strong age trend along the red sequence found by Belli et al. (2019). These galaxies constitute a third path towards quiescence, which comes to dominate at the lowest redshifts. This picture is qualitatively similar to the findings of Chapter 2.

The morphological analysis of Gargiulo et al. (2017) also supports the picture we have presented in this section. They find that the build-up of massive quiescent galaxies at $z < 1$ is dominated by low-density objects, with no evolution observed in the numbers of high-density objects. This again suggests that the densest objects are oldest, and formed by mergers at high redshift, whereas the later build-up of the quiescent population is related to the secular quenching of normal star-forming galaxies. The final VANDELS sample should provide sufficient numbers of objects with *HST* imaging for a joint analysis of morphology and SFHs, meaning this scheme can be tested explicitly.

4.6.3 The star-formation histories of massive quiescent galaxies at $1.0 < z < 1.3$

In light of the discussion of Section 4.6.2, it is interesting to consider the details of the SFH shapes inferred for the quiescent sub-sample. The posterior median SFHs inferred are shown in Fig. 4.13, split into four mass bins. To demonstrate the magnitude of the typical uncertainties, in each bin a randomly selected object is highlighted in orange, and the 16^{th}–84^{th} percentiles of the posterior are shaded.

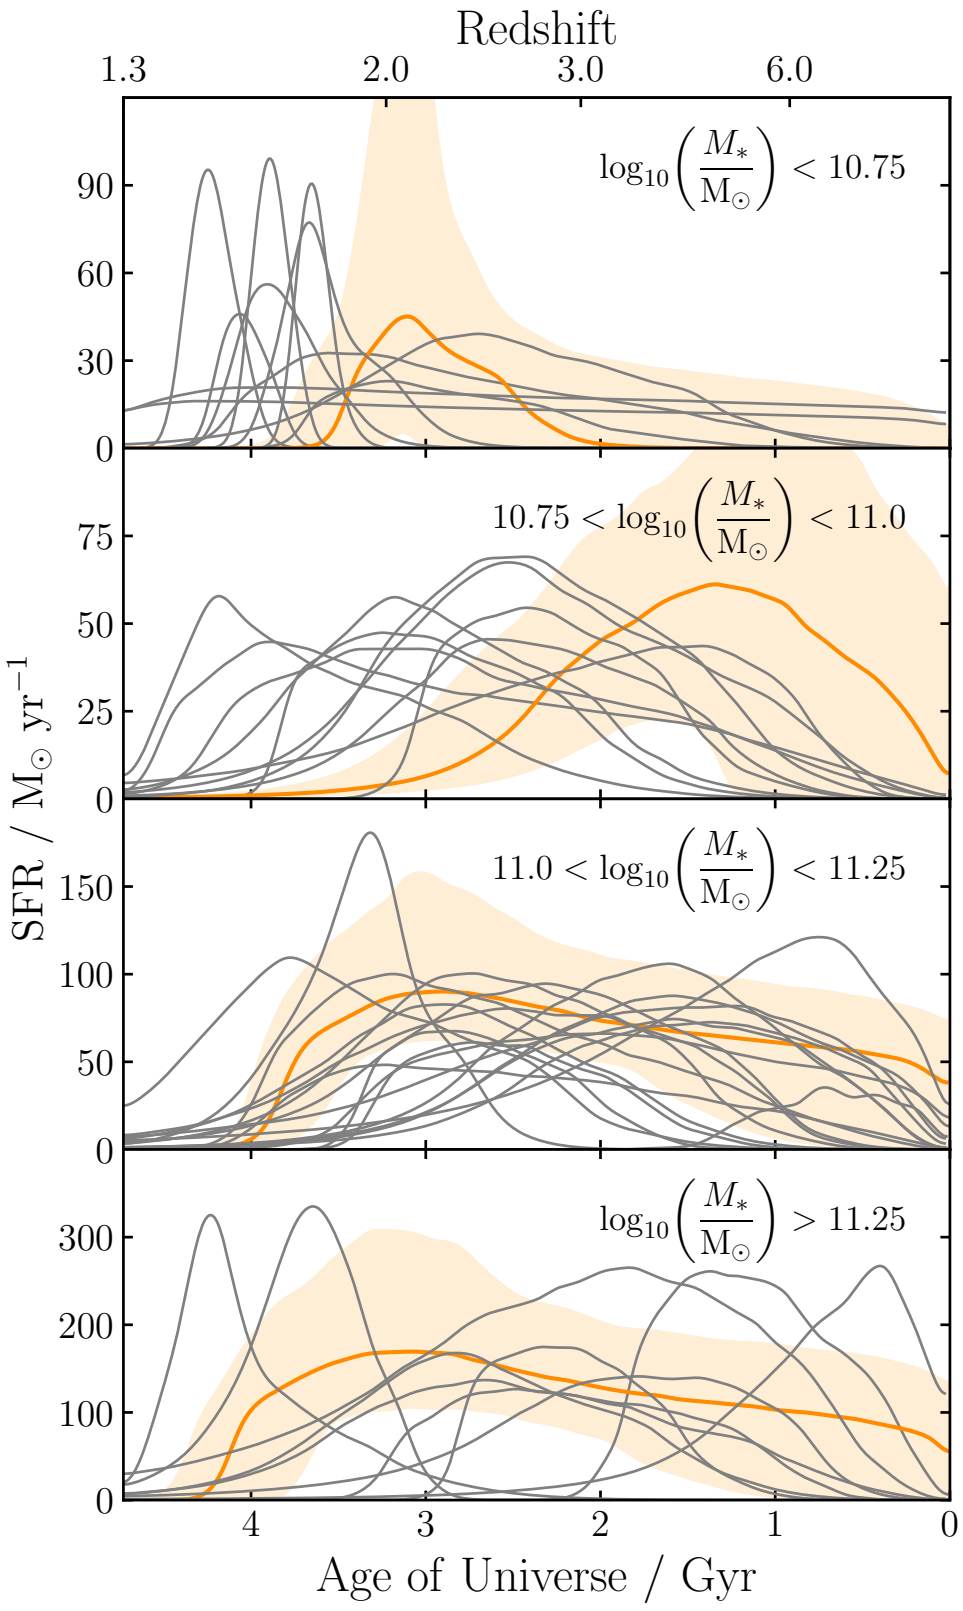

Figure 4.13 *Posterior median SFHs for the quiescent sub-sample, divided into four bins in stellar mass. In each bin, a randomly selected SFH is highlighted, and the* 16^{th}–84^{th} *percentiles of its posterior are shaded to demonstrate the typical uncertainties.*

The overall trend towards earlier formation with increasing stellar mass shown in Fig. 4.3 is clearly visible in Fig. 4.13, however there is significant variation within each mass bin. This demonstrates that, even at fixed stellar mass and observed redshift, the quiescent population contains galaxies with a wide range of formation histories. The fact that knowledge of the stellar mass is not sufficient to make strong predictions about the SFH suggests again that a range of physical processes contribute to the quenching of star-formation.

In the lowest-mass bin the majority of objects have very bursty SFHs, with short, $\lesssim 500$ Myr, formation timescales. A smaller number of similar objects are also present in the higher-mass bins, though they are a minority. The PSB shown in the bottom panel of Fig. 4.12 is the youngest member of this group, though other members form later in cosmic time, as they have lower observed redshifts. Many or all of these objects are likely to have experienced a PSB phase at an earlier stage in their evolution. They are also plausibly linked with the population of highly star-forming submillimetre galaxies, the redshift distribution of which peaks at $z \sim 2$ (e.g. Dunlop et al. 2017). The number of such objects present in the VANDELS DR2 sample does not allow us to place strong statistical constraints on their number density as a function of redshift. However the full VANDELS sample will contain approximately twice as many objects consistent with the selection criteria applied in Section 4.2.3.

The older galaxies in the quiescent sub-sample display a broad range of quenching timescales. Those that quench rapidly are likely to have experienced both green-valley and PSB phases, consistent with the "early fast quench" model in Fig. 4.11. A fraction however also have SFHs consistent with the "early slow quench" model, with quenching timescales of $\gtrsim 1$ Gyr. The two oldest galaxies in the sample, both of which have $\log_{10}(\boldsymbol{M}_*/\mathrm{M}_\odot) > 11$, quenched at $z > 3$. This corresponds to a comoving number density of $2 \pm 1 \times 10^{-5}$ Mpc^{-3} for quiescent galaxies at $z = 3$, consistent with Schreiber et al. 2018, who find a number density of $1.4 \pm 0.3 \times 10^{-5}$ Mpc^{-3} for quiescent galaxies with $\log_{10}(\boldsymbol{M}_*/\mathrm{M}_\odot) > 10.5$ at $3 < z < 4$.

The "early slow quench" objects may genuinely have experienced slow quenching at the highest redshifts, however two other scenarios are possible. Firstly, their SFHs may be composites of several systems that formed and quenched at different times, then merged to form the most massive quiescent galaxies. Secondly, it is possible that these systems were rapidly quenched, however, because their stellar populations are already old when observed, and hence slowly evolving, the VANDELS observations cannot rule out slower quenching scenarios. Future

instruments, such as NIRSpec on board the *James Webb Space Telescope* (JWST), which hold the promise of extending VANDELS-quality observations to quiescent galaxies at the highest redshifts, will provide invaluable constraints on the earlier evolution of these systems.

4.7 Conclusion

In this work I report SFHs for a sample of 75 UVJ-selected galaxies with stellar masses of $\log_{10}(M_*/\mathrm{M}_\odot) > 10.3$ at observed redshifts of $1.0 < z < 1.3$. As described in Section 4.2, the data consists of deep rest-frame UV spectroscopy from VANDELS, as well as multi-wavelength photometry. Using the Bagpipes code I fit the combined datasets with a sophisticated joint model for the physical properties of the galaxies and systematic uncertainties affecting the spectroscopic data, described in Sections 4.3 and 4.4. The combination of extremely deep VANDELS spectroscopy with this sophisticated fitting methodology allows me to significantly improve upon previous analyses, obtaining strong, yet realistic constraints on the SFHs of the target galaxies.

I firstly quantify the stellar mass vs stellar age relationship for massive quiescent galaxies at $1.0 < z < 1.3$ (Section 4.5.2). I find a strong trend towards earlier mass-weighted formation times with increasing stellar mass (downsizing) of $1.48^{+0.34}_{-0.39}$ Gyr per decade in mass (Equation 4.9). A slightly shallower trend of $1.24^{+0.27}_{-0.30}$ Gyr per decade in mass is observed for r-band light-weighted formation times (Equation 4.10). As shown in Fig. 4.8, the slope of this relationship is in agreement with spectroscopic results at $z \sim 0.1$ and $z \sim 0.7$ from Gallazzi et al. (2005, 2014), and at $z \sim 1.75$ from Belli et al. (2019). Recent photometric studies find weaker trends, of $\lesssim 0.5$ Gyr per decade in mass (Fig. 4.9), which can be attributed to larger uncertainties in individual age determinations due to the age-metallicity-dust degeneracy (see Section 4.6.1). I conclude that, at fixed observed redshift, an evolution in the stellar age vs stellar mass relationship of ~ 1.5 Gyr per decade in mass is a robust result, which holds across the observed redshift interval from $0 < z < 2$.

This result places strong constraints on the AGN-feedback models used in modern cosmological simulations. As such, in Fig. 4.8, I compare predictions from the Simba and IllustrisTNG simulations to my observational results. I conclude that the stellar mass vs stellar age relationships predicted by these simulations

at $z = 0.1$ are in good agreement with observations. However, at $z = 1$, the relationships predicted are considerably flatter than observational results, at $\lesssim 0.5$ Gyr per decade in mass. This conclusion supports recent findings that suggest that, whilst modern simulations now reproduce well the properties of local massive quiescent galaxies (e.g. Davé et al. 2017; Nelson et al. 2018), agreement is increasingly poor with increasing observed redshift (e.g. Schreiber et al. 2018).

By considering the distributions I find for galaxy physical parameters on the UVJ diagram (Fig. 4.5), and the SFH shapes I infer for the VANDELS sample (Fig. 4.13), I attempt to understand the connection between green-valley, post-starburst (PSB) and quiescent galaxies, and to constrain quenching mechanisms at $z > 1$. I demonstrate that typical green-valley galaxies, if rapidly quenched, pass through a PSB phase en route to quiescence, and show that SFHs consistent with this evolution exist within the VANDELS sample.

I additionally identify a class of predominantly $\log_{10}(\boldsymbol{M}_*/\mathrm{M}_\odot) \sim 10.5$ galaxies that formed and quenched at $z < 2$ in extreme starbursts with $\lesssim 500$ Myr timescales. These objects are consistent with mergers and associated strong, AGN-driven outflows, and are plausibly related to submillimetre galaxies. These objects also pass through a PSB phase, supporting a dual origin for the PSB population (e.g. Wild et al. 2016; Almaini et al. 2017; Maltby et al. 2018). I finally find that some of the most massive VANDELS galaxies appear to exhibit slow quenching at early times ($z > 2$), though their quenching timescales are harder to constrain due to their older stellar populations. These objects are plausibly the result of mergers between galaxies with stellar populations formed at different times. To understand the earlier evolution of these systems in detail, deep continuum spectroscopy must be extended to the highest redshifts.

These results demonstrate the power of large, high-redshift spectroscopic surveys for placing strong constraints on the evolution of subtle galaxy physical parameters across cosmic time. These results are important for furthering our understanding, as they are strongly constraining on models of galaxy formation. Upcoming instruments such as NIRSpec on JWST and MOONS at the VLT will greatly expand our high-redshift spectroscopic capabilities, and sophisticated fitting methodologies, such as presented in this work, will be key to realising their full potential.

Chapter 5

Conclusions and future work

In this section I summarise the most important conclusions reached in the preceding chapters. I also provide a brief outline of potential future work building on the contents of this thesis.

5.1 Thesis conclusions

In this section I summarise the main conclusions of this thesis. More detailed summaries are provided in the final sections of Chapters 2, 3 and 4.

5.1.1 The BAGPIPES code

Each of the projects presented in this thesis makes use of the Bagpipes spectral modelling and fitting software, introduced in Chapter 2. Bagpipes allows for the generation of sophisticated model galaxy spectra, including complex stellar populations, nebular emission, dust attenuation, dust emission, the effects of the IGM and velocity dispersion. The ability to easily and rapidly access and combine various options for modelling these physical effects within an intuitive Python wrapper provides a major quality-of-life improvement for those seeking to analyse galaxy spectra. Another significant advantage of Bagpipes is the ability to fit these complex models to combinations of photometric and spectroscopic data using a state-of-the-art Bayesian approach, implemented via the MultiNest nested sampling algorithm. Overall, Bagpipes provides an efficient and powerful

tool for extracting physical parameters from observational data, which is already gaining exposure in the literature (e.g. Cullen et al. 2018; Kemp et al. 2019; Wild et al. in prep.; Iyer et al. in prep.; Pacifici et al. in prep.).

5.1.2 Parametric star-formation-history models

The Bagpipes code is firstly used to carry out several analyses related to the use of parametric models to extract galaxy star-formation histories from broad-band photometric data. In Chapter 2, an initial exploration of parametric SFH models in the specific context of massive quiescent galaxies is presented. By constructing and fitting a mock photometric catalogue derived from the Mufasa simulation, I test the abilities of the exponentially declining and double power law parametric SFH models to recover stellar masses, times of quenching and times of formation. I find that the use of the exponentially declining model produces, on average, more-highly biased physical parameter estimates than the double power law model.

A more general exploration of the use of four parametric SFH models for recovering physical parameters from broad-band photometric data is presented in Chapter 3. I propose a general framework for such tests, applicable to any kind of SFH model, whereby the priors on physical parameters, such as the sSFR and mass-weighted age, are obtained by sampling a large number of SFHs from the chosen prior distributions on model parameters. This provides an understanding of the assumptions that are being made regarding physical parameters, which are often not obvious from the characteristics of the chosen model.

By the application of this method, I demonstrate that, at $z = 0$, the four parametric models considered all favour sSFRs within a narrow range around $\sim 10^{-10}$ yr^{-1}. This is significant as it could lead to the artificial imposition of a well-defined star-forming main sequence on populations of galaxies. I also demonstrate that the redshift evolution of this sSFR prior is in the same sense as the redshift evolution of the star-forming main sequence.

By fitting a volume-complete local sample from the Galaxy and Mass Assembly Survey, I then attempt to reconstruct the redshift evolution of the cosmic star-formation rate density using these parametric SFH models. The four models each suggest a peak at $z \sim 0.4$, approximately 6 Gyr later than found by Madau & Dickinson (2014) using direct measurements of high redshift galaxy SFRs. This

indicates that these models all impose a strong and unjustified prior preference for young galaxy stellar populations.

The results presented in Chapter 3 show that, whilst the use of parametric SFH models to fit broad-band photometric data can be effective in the recovery of photometric redshifts and stellar masses, all other galaxy physical parameters are subjected to significant biases, depending on the details of the model being fitted. I also demonstrate that broad-band photometric data cannot be used to discriminate between SFH models, meaning these data are of limited usefulness for understanding the evolutionary paths of individual galaxies.

Improvements can be made by the use of non-parametric SFH models, which allow the more direct imposition of specific, physically motivated priors on physical parameters. However, robust cross-checking, for example using the tests described in the previous paragraphs, are a necessary piece of due diligence for any such analysis. In order to obtain more-detailed information in a way that does not depend strongly on the applied priors, the use of more-sophisticated models must be combined with more-strongly constraining spectroscopic observational data.

5.1.3 The star-formation histories of massive quiescent galaxies

The main aim of this thesis is to place constraints on the processes that quench star-formation in massive quiescent galaxies via the inference of their star-formation histories. Two analyses are presented, firstly in Chapter 2, using a large sample of photometric data from the UltraVISTA survey, and secondly in Chapter 4 using a smaller, spectroscopic sample from VANDELS. Whilst the VANDELS sample is smaller and spans a reduced redshift range, the extremely deep spectroscopic data allows strong constraints to be placed on the SFHs of individual objects, complementing the large statistical sample from UltraVISTA.

The UltraVISTA analysis of Chapter 2 demonstrates the well-known stellar mass vs stellar age relationship, with more massive galaxies being, on average, older than their less massive counterparts at fixed observed redshift (downsizing). However, this relationship is found to be relatively weak, at ~ 0.5 Gyr per decade in stellar mass. At fixed stellar mass, the well-known trend towards lower average formation redshift with decreasing observed redshift is also recovered. This is a consequence of new galaxies quenching across the redshift interval

considered ($0.25 < z < 3.75$), adding younger stellar populations to the quiescent population. However, we also find significant evidence for continued evolution of the stellar populations of quiescent galaxies post-quenching, with significantly fewer extremely old objects at low redshift than would be predicted from high-redshift quiescent galaxy number densities, assuming pure-passive evolution.

I also consider the quenching timescales for UltraVISTA quiescent galaxies in Chapter 2, finding evidence for three separate quenching processes with different timescales, which change in relative importance as a function of stellar mass and observed redshift. At $z > 1$, a significant fraction of objects display extremely rapid quenching timescales of < 1 Gyr, with star-formation quickly rising before quenching takes place. These objects become less prevalent with decreasing observed redshift. A second class of objects experience much more extended star-formation epochs before also rapidly quenching, over timescales of $\sim 1 - 2$ Gyr. These objects form the majority of the sample, and are prevalent across the whole observed redshift and stellar mass ranges considered. The final class of objects, which are increasingly prevalent towards the lowest redshifts ($z < 1$), exhibit slow quenching, over timescales of $\gtrsim 3$ Gyr.

The VANDELS analysis of Chapter 4 also constrains the stellar mass vs stellar age relationship, however a much stronger trend is observed when analysing this spectroscopic dataset, of ~ 1.5 Gyr in formation time per decade in stellar mass. I show that this is consistent with other spectroscopic analyses from $0 < z < 2$ by Gallazzi et al. (2005, 2014) and Belli et al. (2019). The difference in the strength of this relationship between the two analyses presented is likely to stem from the increased uncertainties on individual age measurements when considering the UltraVISTA photometric data compared to the VANDELS spectroscopic data. These larger uncertainties act to flatten this relationship, highlighting the importance of large, high-quality spectroscopic surveys.

The quenching properties inferred from these two analyses are broadly consistent, with the $z > 1$ quiescent population split between objects that form in rapid $\lesssim 500$ Myr bursts, and objects that form more slowly, but quench almost equally as rapidly, over $\lesssim 1$ Gyr timescales. The former class of objects are most plausibly associated with large-scale disruption events, such as mergers and subsequent strong quasar-mode AGN feedback, which has the potential to lead to the extremely rapid bursting and quenching of star-formation. The latter class of objects are less clearly associated with a specific process, however their quenching timescales are in broad agreement with those for galaxies quenched by jet-mode

AGN feedback in modern cosmological hydrodynamic simulations (e.g. Davé et al. 2017; Nelson et al. 2018). These objects are also found down to $z = 0.25$ in the UltraVISTA analysis. The final class of slow-quenching objects found in the UltraVISTA sample at $z < 1$ are consistent with other recent observational results in the local Universe (e.g. Peng et al. 2015; Schreiber et al. 2016).

5.2 Future work

The work presented in this thesis forms part of a natural progression towards tackling more complex astrophysical questions using higher quality data and more advanced modelling approaches. In this section I present several ideas for potential future observations and follow-up projects that would build upon the work presented in Chapters 2, 3 and 4, further developing our understanding of the quenching of star-formation in galaxies.

5.2.1 Resolved studies of high-redshift quiescent galaxies with JWST

As discussed throughout this thesis, understanding the formation of quiescent galaxies is a key area of ongoing research in the field of galaxy evolution. The clear colour (e.g. Strateva et al. 2001) and morphological (e.g. Bell et al. 2012) bimodalities observed in the local galaxy population point towards a well-defined transition between star-forming and quiescent states. However, the physical processes that drive this quenching of star-formation are still relatively poorly understood. Progress has been made using cosmological hydrodynamic simulations, such as Mufasa (Davé et al. 2016), which are now able to accurately reproduce the local galaxy colour bimodality (e.g. Davé et al. 2017; Nelson et al. 2018), providing hints at the important physical ingredients.

In order to deepen our understanding, a range of more subtle predictions from simulations, such as the resolved stellar ages, metallicities and kinematics of quiescent galaxies must be subjected to rigorous observational testing. At low redshift, integral-field spectroscopic surveys such as SAURON (de Zeeuw et al. 2002), ATLAS^{3D} (Cappellari et al. 2011) and MaNGA (Bundy et al. 2015) have begun to provide observational constraints on the resolved properties of massive quiescent galaxies. A particular success has been the confirmation of the

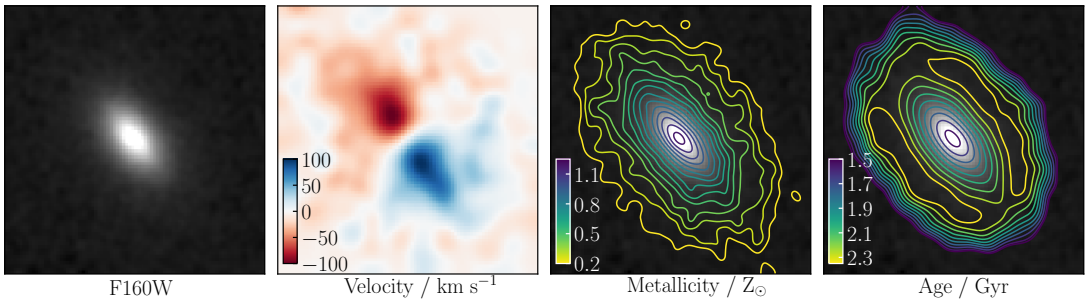

Figure 5.1 *The left panel shows a mock* $5'' \times 5''$ *CANDELS F160W cutout for the most massive quiescent galaxy at* $z = 1.5$ *in the medium resolution (25 Mpc/h box length) run of the* MUFASA *simulation. The other panels show the resolved stellar velocity field, stellar metallicity and stellar age maps for this object. Observations using the JWST NIRSpec IFU have the potential to recover these properties, facilitating more detailed comparisons with simulations.*

kinematic differentiation within the quenched population between fast and slow rotators (e.g. Emsellem et al. 2007), predicted by simulations as a consequence of formation through gas rich (wet) and gas poor (dry) mergers respectively (e.g. Khochfar & Burkert 2003; Khochfar et al. 2011).

The extension of these studies to higher redshifts is desirable for a number of reasons. Firstly, low redshift observations are not strongly constraining on the early-time evolution of galaxies. The weak time-evolution of the extremely old stellar populations of local quiescent galaxies makes detailed star-formation history recovery impossible, whereas their kinematics are dominated by the most recent interactions they have undergone. Secondly, simulations still struggle to reproduce correct number densities for $z \gtrsim 1$ quiescent galaxies (e.g. Davé et al. 2017; Schreiber et al. 2018). This, combined with recent observational evidence, including that presented in Chapters 2 and 4, suggests a change in quenching physics at $z \gtrsim 1$. This indicates that current simulations are incomplete, and raises questions as to the connection between the low and high redshift quiescent populations.

However, the extremely compact nature of high redshift quiescent galaxies (e.g. typical half-light radii of 3 kpc, equivalent to $\sim 0.3''$, for the most massive objects at $z = 1.4$; McLure et al. 2013), as well as the shifting of important rest-frame optical spectral features into the near-infrared, have made such extensions extremely challenging with existing instrumentation. The KMOS^{3D} Survey (Wisnioski et al. 2015) targets $0.7 < z < 2.7$ quiescent galaxies, but their mean spatial resolution of $\sim 0.6''$ is insufficient for extracting resolved properties. An

example of the resolved properties of a high-redshift massive quiescent galaxy from the MUFASA simulation is shown in Fig. 5.1. This object can be seen to be fast rotating, and to have strong radial age and metallicity gradients.

The transformative spatial resolution ($\lesssim 0.1''$) and sensitivity provided by the NIRSpec IFU onboard the *James Webb Space Telescope* (JWST) will allow resolved studies of the stellar-population and emission-line properties of massive quiescent galaxies to be extended to high redshifts for the first time. The scientific objectives of such a survey, as well as practical considerations for survey design, are discussed in the following sections.

What is the dominant quenching mechanism at high redshift?

The success of simulations in replicating the local galaxy colour bimodality is based on modelling low-accretion (jet) mode AGN feedback by injecting energy into the circumgalactic medium (e.g. Croton et al. 2006; Gabor et al. 2011). At low redshift the resulting quenching timescales of $1-2$ Gyr (Nelson et al. 2018) are in agreement with recent observational studies (e.g. Pacifici et al. 2016). However, at $z > 1$, observational evidence for more rapid quenching (e.g. Maiolino et al. 2012; Förster Schreiber et al. 2014; Schawinski et al. 2014), and underestimated number densities for quiescent galaxies in simulations (e.g. Schreiber et al. 2018), suggest a change in the dominant quenching mechanism to a faster mode that is not currently accounted for. One candidate is high-accretion (quasar) mode AGN feedback.

Resolved observations of massive quiescent galaxies at these redshifts with the JWST NIRSpec IFU would allow us to recover resolved star-formation histories with far greater precision than is possible in the local Universe, owing to the significantly younger stellar populations of these objects. This would allow us to discriminate between different physical scenarios for the quenching of star-formation. For example, the MUFASA object shown in Fig. 5.1 has a central age gradient of 1.04 Gyr/arcsec, indicating that star-formation became more concentrated at later times (outside-in quenching). Conversely, in the scenario of quasar-mode feedback, the very rapid expulsion of gas would result in a flat age gradient.

What drives the strong redshift evolution of the size-mass relationship?

As previously noted, the average physical sizes of quiescent galaxies have been found to be a strong function of redshift (e.g. McLure et al. 2013; van der Wel et al. 2014). Two scenarios have been proposed to explain this. Either existing quiescent galaxies expand (e.g due to dry minor mergers), or larger star-forming galaxies undergo quenching. It seems certain that both of these processes contribute to the evolution of the size-mass relationship, but their relative importances and environmental dependencies remain poorly understood. By obtaining precise resolved star-formation histories and metallicities for massive quiescent galaxies at high redshift across a range of environments, fresh insights would be provided into evolutionary differences between galaxy clusters and the field. Finely resolved morphologies would also allow us to understand better the role of minor merger events in the continuing evolution of quiescent galaxies.

Is the kinematic distinction between fast and slow rotators already in place by z = 1.5?

In the local Universe, the most massive quiescent galaxies are more than twice as massive as the most massive star-forming galaxies. Therefore, mergers between quiescent galaxies are necessary to produce the most massive systems. As noted above, simulations predict that galaxies formed in such dry mergers will not exhibit ordered rotation, in contrast to those formed in wet mergers. Extending resolved kinematic studies to high redshift would allow us to understand whether the observational result that the majority ($\sim$ 75%) of the most massive local quiescent galaxies are slow rotators (Emsellem et al. 2011) evolves with redshift.

What is the rate and dominant mechanism of mass growth in the quenched population at z > 1?

Recently, Belli et al. (2017) reported low-level ongoing star-formation activity in $\sim$ 17% of 120 KMOS^{3D} quiescent galaxies. Additional hints as to ongoing star-formation and AGN activity in $z > 1$ quiescent galaxies were also uncovered in Chapter 4. NIRSpec IFU observations have the potential to yield a full suite of rest-frame optical emission lines, allowing us to constrain the gas kinematics, metallicities and ionization states of massive quiescent galaxies at high redshift.

This will allow us to constrain any AGN activity, and to map the spatial distribution of ongoing star-formation. Mass growth rates through star-formation can be compared with mergers in order to determine which dominates, and the total growth rates can be compared with results from simulations.

Survey design

Both the VANDELS data and Bagpipes spectroscopic analysis techniques discussed in Chapter 4 provide natural starting points for the design of such a survey. VANDELS provides robust spectroscopic redshifts for massive quiescent galaxies from $1.0 < z < 2.5$ in the well-studied CANDELS UDS and CDFS fields, whereas the ability to fit sophisticated physical models provided by Bagpipes allows detailed physical properties to be extracted from spectroscopic data.

The objectives discussed above require resolved high-SNR spectroscopic observations of the stellar absorption features between rest-frame wavelengths of ~ 3700Å and ~ 5500Å to produce the strongest possible constraints on resolved star-formation histories, stellar metallicities and stellar kinematics. Observations of the strong optical emission lines Hβ at 4861Å, [O III] at 5007Å and Hα at 6563Å would also be valuable in order to constrain the kinematic and physical properties of the ionized gas in potential targets, and hence the roles of ongoing star-formation and AGN activity.

A natural choice of target redshift range is therefore $1.60 < z < 1.88$, as all of the relevant features fall between $0.97 - 1.89$ μm, and so can be observed simultaneously through the F100LP filter (see Fig. 1.6). Given this optimum redshift range, the CANDELS UDS field is a natural choice, due to the presence of a known rich cluster at $z = 1.62$ (Papovich et al. 2010).

5.2.2 The stellar mass-metallicity relationship for VANDELS galaxies

VANDELS provides extremely deep ($\sim$ 45 hours on average) rest-frame UV spectra for ~ 250 massive quiescent galaxies at $1.0 < z < 1.5$ over a combined area of ~ 0.2 deg^2 in UDS and CDFS. These spectra contain a wealth of age and metallicity sensitive features, which I have used in Chapter 4 to extend studies of massive quiescent galaxy physical properties to a statistical sample at $z > 1$ for

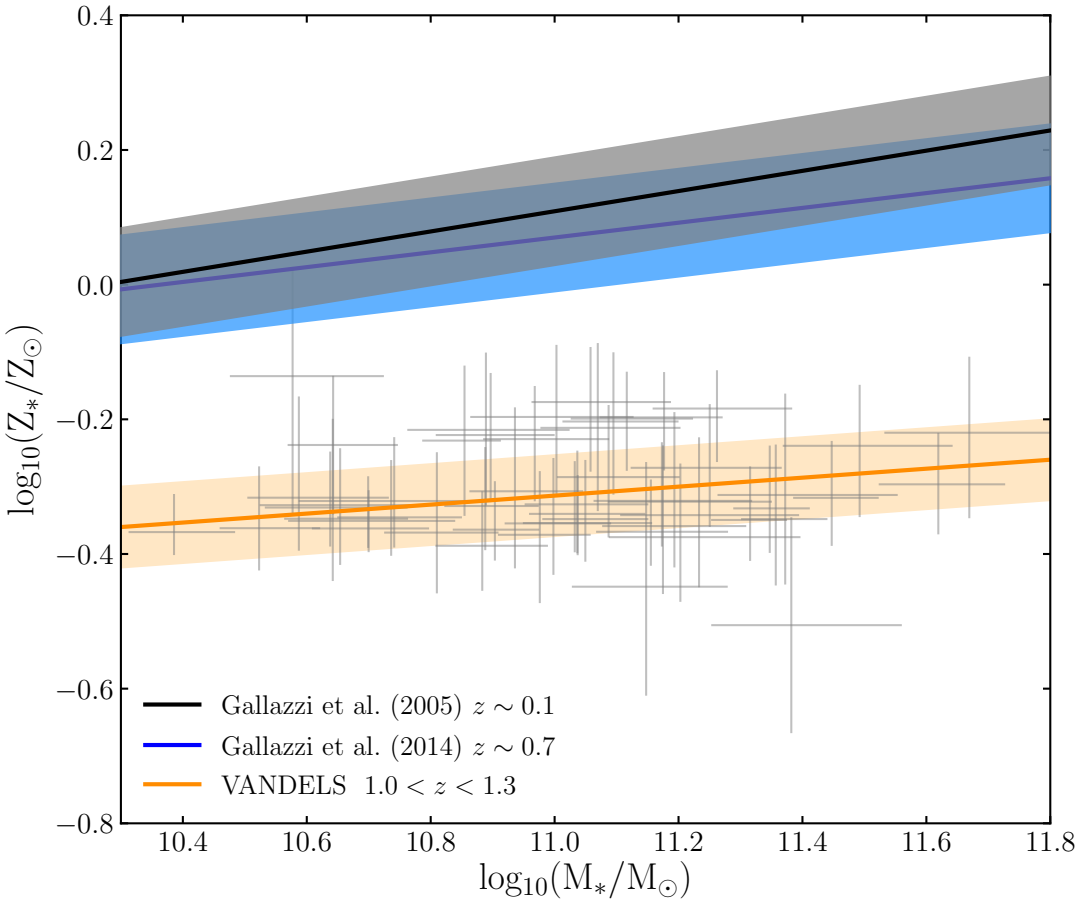

Figure 5.2 *The rest-ultraviolet stellar mass-metallicity relationship for VANDELS massive quiescent galaxies, from the analysis presented in Chapter 4. The rest-optical stellar mass-metallicity relationships from Gallazzi et al. (2005, 2014) at $z \sim 0.1$ and $z \sim 0.7$ are also shown. My results indicate an evolution of approximately -0.4 dex.*

the first time.

Whilst the star-formation histories of the VANDELS galaxies were the main focus of Chapter 4, my analysis also constrains their stellar metallicities, based on their rest-frame ultraviolet spectra. Fig. 5.2 shows the inferred rest-ultraviolet stellar metallicities, compared to rest-optical studies at lower redshift. It can be seen that I recover typical stellar metallicities of $Z_* \sim 0.5$ $Z_\odot$, indicating an evolution of approximately -0.4 dex from Gallazzi et al. (2014) at $z \sim 0.7$. However, since these estimates are based on different sets of metallicity sensitive features over different wavelength ranges, significant systematic uncertainties potentially exist between these studies.

Future observations could build upon the success of VANDELS by following up a

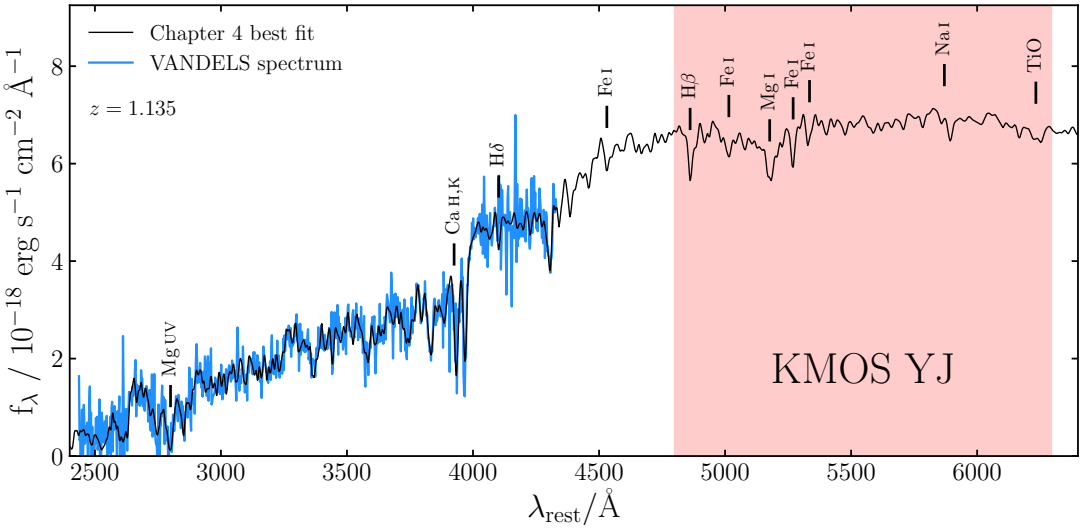

Figure 5.3 *Example KMOS YJ-band observation for an object from the VANDELS sample. The best fit (black) from Chapter 4 to the VANDELS data (blue) has a stellar metallicity, Z_*, of $0.51^{+0.09}_{-0.08}$ $Z_\odot$, based largely on Mg_{UV} at $\sim 2800Å$. KMOS YJ-band observations (red) at $1.0 < z < 1.5$ would probe the core set of Mg and Fe indices at $4500Å < \lambda < 5500Å$ typically used to obtain optical stellar metallicities in the local Universe (e.g. Gallazzi et al. 2005). The most important age and metallicity sensitive stellar absorption features in this wavelength range are labelled.*

subset of these objects in the KMOS YJ-band, giving us access to the rest-frame optical Fe and Mg absorption features used to study stellar metallicities in the local Universe. The scientific objectives of such a survey, as well as practical considerations for survey design, are discussed in the following sections.

Extending the stellar mass-metallicity relation to high redshift

The relationship between stellar mass and stellar metallicity as a function of cosmic time provides a fundamental constraint on models of galaxy formation (e.g. Maiolino & Mannucci 2019). This relationship has been extensively studied in the local Universe, primarily using rest-frame optical Fe and Mg absorption features (e.g. Gallazzi et al. 2005; Panter et al. 2008; Zahid et al. 2017). For massive quiescent galaxies, a relatively flat, slightly super-Solar mass-metallicity relation has been observed, as shown in Fig. 5.2. Studies at higher redshifts find little evolution with cosmic time, at least as far as $z \sim 0.7$ (Gallazzi et al. 2014).

In Chapter 4, I present a full-spectral fitting analysis designed to constrain the

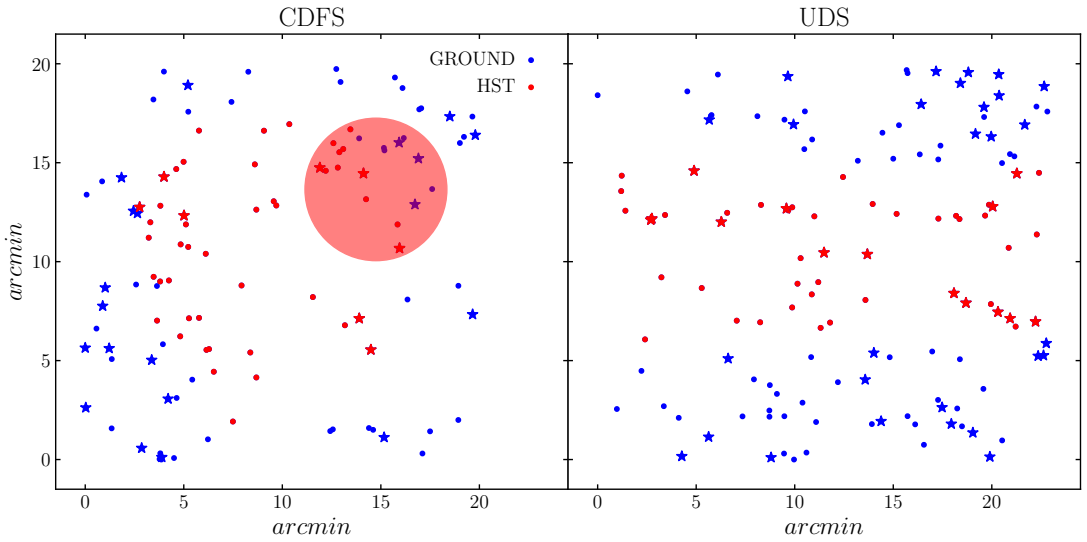

Figure 5.4 *The distribution of VANDELS quiescent galaxies from* $1.0 < z < 1.5$ *in CDFS and UDS. Red targets have HST imaging, blue targets have ground-based imaging. The 75 objects considered in Chapter 4 are marked with stars. An example KMOS pointing is shown as a red circle, containing 21 objects. This demonstrates that the target density is high enough to fill all of the available KMOS arms.*

star-formation histories for a subset of the VANDELS passive sample at $1.0 < z < 1.3$. Fig. 5.3 shows an example VANDELS spectrum (blue) and the best-fitting model (black), with several key features labelled. Near-infrared follow-up of these galaxies would allow us to self-consistently extend our knowledge of the stellar mass-metallicity relation to $z > 1$.

If these results supported significant evolution in the rest-optical stellar mass-metallicity relationship at $z > 0.7$, this would be strong evidence that the transition towards slower quenching timescales observed at $z \sim 1$ in Chapter 2 (see also Schawinski et al. 2014; Schreiber et al. 2016; Wild et al. 2016) is associated with a change in the dominant physical process that quenches star-formation. Alternatively, if these rest-optical features are consistent with super-Solar metallicities as found at lower redshift, this would indicate that the ultraviolet and optical spectra of massive quiescent galaxies at $1.0 < z < 1.5$ are dominated by stellar populations with significantly different metallicities (see below). The combination of robust spectroscopic redshifts and rich multi-wavelength ancillary data makes VANDELS the ideal sample for extending studies of this key relationship out to $1.0 < z < 1.5$.

Rejuvenated star-formation in massive quiescent galaxies

Recently, Belli et al. (2017) presented evidence for low-level ongoing star-formation in massive quiescent galaxies from $0.7 < z < 2.7$. Crucially, they find lower gas-phase metallicities in these galaxies than are observed for star-forming galaxies with similar stellar masses. The light-weighted ages of stellar populations are a strong function of wavelength, with younger populations dominating the ultraviolet and older populations dominating the optical and infrared. Therefore, as well as potentially indicating redshift evolution, the results shown in Fig. 5.2 could indicate lower-metallicity star formation at later stages in the evolution of quiescent galaxies.

Such low-metallicity rejuvenation events must be fuelled by inflowing gas, potentially associated with minor-merger events, meaning that the scale and frequency of such events is strongly connected to the processes by which star-formation remains quenched in quiescent galaxies. These processes are challenging to reproduce in cosmological simulations, with a constant input of energy required to prevent galaxies quenched in violent events at high redshift re-igniting star-formation on $1 - 2$ Gyr timescales (Gabor et al. 2010, 2011).

Rest-frame optical data would allow us to directly compare the optical and ultraviolet stellar metallicities of VANDELS galaxies, constraining the role of rejuvenated low-metallicity star-formation activity. Additionally, as discussed in Chapter 4, a significant fraction of the VANDELS sample exhibits [O II] emission, and rest-optical follow-up observations would cover the [O III] and Hβ emission lines (as well as Hα for a small number of objects with the lowest redshifts). Where present, these lines would provide us with additional insights as to the level of ongoing star formation and role of AGN activity in $z > 1$ quiescent galaxies.

Survey design

In order to obtain strong constraints on the stellar mass-metallicity relationship, a statistical sample of objects is necessary ($\sim$ 100; e.g. Gallazzi et al. 2014). As shown in Fig. 5.4, the target density of VANDELS quiescent galaxies from $1.0 < z < 1.5$ in UDS and CDFS is high enough that all of the available KMOS arms could be filled in each pointing. Approximately 5 pointings would therefore be required to obtain a statistical sample of $\sim$ 100 objects. To obtain strong constraints on stellar metallicities, a SNR of $\sim$ 20 per resolution element at R =

1000 is required (e.g. Gallazzi et al. 2005; Pacifici et al. 2012). The VANDELS targets are extremely compact, with a median effective radius of $\sim 0.35''$, and can therefore be approximated as point sources for the purposes of ground-based observations. Assuming $0.8''$ seeing and an airmass of 1.2, an integration time of 8 hours is required for an object with a J-band magnitude of 21.9, which is the median value for the VANDELS sample shown in Fig. 5.4.

5.2.3 Re-analysis of archival data from SDSS, Lega-C and VANDELS

The rapid increase in the volume of data available in astronomy in recent years, combined with the rapid development of new and more sophisticated analysis techniques, have led to issues when attempting to compare results obtained at different times from different datasets. For example, a large number of studies of the star-forming main sequence in the late 2000s and early 2010s found apparently contradictory results, which were shown by Speagle et al. (2014) to be mostly due to differences in sample selection and analysis techniques.

In the case of the stellar age vs stellar mass relationship, as discussed in Chapter 4, analyses conducted at low redshift using SDSS data and spectral indices methods differ significantly from the more sophisticated methods that are now being applied to new, higher redshift data (e.g. Chauke et al. 2018; Belli et al. 2019). Similar issues exist for the stellar mass-metallicity relationship, as discussed in Section 5.2.2. For this reason, a self-consistent re-analysis of spectroscopic data across the whole available redshift range from SDSS, Lega-C and VANDELS using several independent methods would be extremely valuable. Firstly, this would allow for the development of a clear understanding of the time-evolution of the relationships between stellar mass, age and metallicity. Secondly, this would allow us to develop an understanding of systematic differences that arise due to the application of different methods.

Bibliography

Abramson L. E., Gladders M. D., Dressler A., Oemler Jr. A., Poggianti B., Vulcani B., 2016, ApJ, 832, 7

Acquaviva V., Gawiser E., Guaita L., 2011, ApJ, 737, 47

Allen M. G., Groves B. A., Dopita M. A., Sutherland R. S., Kewley L. J., 2008, The Astrophysical Journal Supplement Series, 178, 20

Almaini O., et al., 2017, MNRAS, 472, 1401

Alpher R. A., Herman R. C., 1948, Physical Review, 74, 1737

Alpher R. A., Bethe H., Gamow G., 1948, Physical Review, 73, 803

Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2015, IEEE Transactions on Pattern Analysis and Machine Intelligence, 38, 252

Anders P., Fritze-v. Alvensleben U., 2003, A&A, 401, 1063

Anders E., Grevesse N., 1989, Geochimica Cosmochimica Acta, 53, 197

Ashby M. L. N., et al., 2013, ApJ, 769, 80

Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, ARA&A, 47, 481

Astropy Collaboration et al., 2013, A&A, 558, A33

Bañados E., et al., 2018, Nature, 553, 473

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

Baes M., Trčka A., Camps P., Nersesian A., Trayford J., Theuns T., Dobbels W., 2019, MNRAS, 484, 4069

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

Baldry I. K., et al., 2018, MNRAS, 474, 3875

Baldwin J. A., Phillips M. M., Terlevich R., 1981, PASP, 93, 5

Barro G., et al., 2013, ApJ, 765, 104

Bayes T., Price R., 1763, Philosophical Transactions (1683-1775), 53, 370

Becker R. H., et al., 2001, AJ, 122, 2850

Behroozi P. S., Wechsler R. H., Conroy C., 2013a, ApJ, 762, L31

Behroozi P. S., Wechsler R. H., Conroy C., 2013b, ApJ, 770, 57

Behroozi P., Wechsler R., Hearin A., Conroy C., 2018, preprint, (arXiv:1806.07893)

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

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

Bell E. F., et al., 2004, ApJ, 608, 752

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

Belli S., et al., 2017, ApJ, 841, L6

Belli S., Newman A. B., Ellis R. S., 2019, ApJ, 874, 17

Best P. N., Kauffmann G., Heckman T. M., Brinchmann J., Charlot S., Ivezić Ž., White S. D. M., 2005, MNRAS, 362, 25

Best P. N., Ker L. M., Simpson C., Rigby E. E., Sabater J., 2014, MNRAS, 445, 955

Blanton M. R., et al., 2003, ApJ, 594, 186

Bonzini M., Padovani P., Mainieri V., Kellermann K. I., Miller N., Rosati P., Tozzi P., Vattakunnel S., 2013, MNRAS, 436, 3759

Bowler R. A. A., et al., 2014, MNRAS, 440, 2810

Bowler R. A. A., et al., 2015, MNRAS, 452, 1817

Brammer G. B., van Dokkum P. G., Coppi P., 2008, ApJ, 686, 1503

Brammer G. B., et al., 2009, ApJ, 706, L173

Brammer G. B., et al., 2011, ApJ, 739, 24

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

Brinchmann J., Charlot S., Kauffmann G., Heckman T., White S. D. M., Tremonti C., 2013, MNRAS, 432, 2112

Broadhurst T. J., Ellis R. S., Shanks T., 1988, MNRAS, 235, 827

Bromm V., Yoshida N., 2011, Annual Review of Astronomy and Astrophysics, 49, 373

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

Buat V., et al., 2014, A&A, 561, A39

Buchner J., et al., 2014, A&A, 564, A125

Bundy K., et al., 2015, ApJ, 798, 7

Burbidge E. M., Burbidge G. R., Fowler W. A., Hoyle F., 1957, Reviews of Modern Physics, 29, 547

Byler N., Dalcanton J. J., Conroy C., Johnson B. D., 2017, ApJ, 840, 44

Byrd R., Lu P., Nocedal J., Zhu C., 1995, SIAM Journal of Scientific Computing, 16, 1190

Calzetti D., Kinney A. L., Storchi-Bergmann T., 1994, ApJ, 429, 582

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

Cappellari M., 2017, MNRAS, 466, 798

Cappellari M., et al., 2011, MNRAS, 413, 813

Cardelli J. A., Clayton G. C., Mathis J. S., 1989, ApJ, 345, 245

Carilli C. L., Walter F., 2013, Annual Review of Astronomy and Astrophysics, 51, 105

Carnall A. C., 2017, preprint, (arXiv:1705.05165)

Carnall A. C., et al., 2015, MNRAS, 451, L16

Carnall A. C., McLure R. J., Dunlop J. S., Davé R., 2018, MNRAS,

Carnall A. C., Leja J., Johnson B. D., McLure R. J., Dunlop J. S., Conroy C., 2019a, ApJ, 873, 44

Carnall A. C., et al., 2019b, preprint, 1903.11082 (arXiv:1903.11082)

Carson D. P., Nichol R. C., 2010, MNRAS, 408, 213

Casey C. M., 2012, MNRAS, 425, 3094

Chabrier G., 2003, ApJ, 586, L133

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

Charlot S., Longhetti M., 2001, MNRAS, 323, 887

Chauke P., et al., 2018, ApJ, 861, 13

Chevallard J., Charlot S., 2016, MNRAS, 462, 1415

Chevallard J., et al., 2019, MNRAS, 483, 2621

Choi J., Conroy C., Moustakas J., Graves G. J., Holden B. P., Brodwin M., Brown M. J. I., van Dokkum P. G., 2014, ApJ, 792, 95

Choi J., Conroy C., Johnson B. D., 2019, preprint, (arXiv:1901.06391)

Chuter R. W., et al., 2011, MNRAS, 413, 1678

Cid Fernandes R., Mateus A., Sodré L., Stasińska G., Gomes J. M., 2005, MNRAS, 358, 363

Ciesla L., et al., 2015, A&A, 576, A10

Ciesla L., Elbaz D., Fensch J., 2017, A&A, 608, A41

Cimatti A., et al., 2004, Nature, 430, 184

Cimatti A., et al., 2008, A&A, 482, 21

Cimatti A., et al., 2013, ApJ, 779, L13

Citro A., Pozzetti L., Moresco M., Cimatti A., 2016, A&A, 592, A19

Cochrane R. K., Best P. N., 2018, MNRAS, 480, 864

Cohn J. D., 2018, MNRAS, 478, 2291

Colless M., 1999, Philosophical Transactions of the Royal Society of London Series A, 357, 105

Colless M., Ellis R. S., Taylor K., Hook R. N., 1990, MNRAS, 244, 408

Conroy C., 2013, Annual Review of Astronomy and Astrophysics, 51, 393

Conroy C., Wechsler R. H., 2009, ApJ, 696, 620

Conroy C., Villaume A., van Dokkum P. G., Lind K., 2018, ApJ, 854, 139

Croton D. J., et al., 2006, MNRAS, 365, 11

Cullen F., Cirasuolo M., McLure R. J., Dunlop J. S., Bowler R. A. A., 2014, MNRAS, 440, 2300

Cullen F., McLure R. J., Khochfar S., Dunlop J. S., Dalla Vecchia C., 2017, MNRAS, 470, 3006

Cullen F., et al., 2018, MNRAS, 476, 3218

Curtis H. D., 1917, PASP, 29, 206

Czekala I., Andrews S. M., Mandel K. S., Hogg D. W., Green G. M., 2015, ApJ, 812, 128

Daddi E., et al., 2005, ApJ, 626, 680

Davé R., Thompson R., Hopkins P. F., 2016, MNRAS, 462, 3265

Davé R., Rafieferantsoa M. H., Thompson R. J., 2017, MNRAS, 471, 1671

Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, preprint, (arXiv:1901.10203)

Dey A., Spinrad H., Stern D., Graham J. R., Chaffee F. H., 1998, ApJ, 498, L93

Dickinson M., Giavalisco M., GOODS Team 2003a, in The Mass of Galaxies at Low and High Redshift. p. 324 (arXiv:astro-ph/0204213), doi:10.1007/10899892_78

Dickinson M., Papovich C., Ferguson H. C., Budavári T., 2003b, ApJ, 587, 25

Diemer B., Sparre M., Abramson L. E., Torrey P., 2017, ApJ, 839, 26

Domínguez A., et al., 2013, ApJ, 763, 145

Donnari M., et al., 2019, MNRAS, 485, 4817

Dopita M. A., Kewley L. J., Heisler C. A., Sutherland R. S., 2000, ApJ, 542, 224

Draine B. T., Li A., 2007, ApJ, 657, 810

Driver S. P., et al., 2009, Astronomy and Geophysics, 50, 5.12

Driver S. P., et al., 2016, MNRAS, 455, 3911

Driver S. P., et al., 2018, MNRAS, 475, 2891

Dunlop J. S., 2013, in Wiklind T., Mobasher B., Bromm V., eds, Astrophysics and Space Science Library Vol. 396, The First Galaxies. p. 223 (arXiv:1205.1543), doi:10.1007/978-3-642-32362-1_5

Dunlop J., Peacock J., Spinrad H., Dey A., Jimenez R., Stern D., Windhorst R., 1996, Nature, 381, 581

Dunlop J., et al., 2007, A Spitzer Public Legacy survey of the UKIDSS Ultra Deep Survey, Spitzer Proposal

Dunlop J. S., et al., 2017, MNRAS, 466, 861

Efstathiou G., Ellis R. S., Peterson B. A., 1988, MNRAS, 232, 431

Einstein A., 1917, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin, pp 142–152

Eisenstein D. J., et al., 2005, ApJ, 633, 560

Eldridge J. J., Stanway E. R., 2009, MNRAS, 400, 1019

Ellis R. S., 2008, Observations of the High Redshift Universe. pp 259–364, doi:10.1007/978-3-540-74163-3_3

Ellis R. S., Colless M., Broadhurst T., Heyl J., Glazebrook K., 1996, MNRAS, 280, 235

Ellis R., Santos M. R., Kneib J.-P., Kuijken K., 2001, ApJ, 560, L119

Ellis R. S., et al., 2013, ApJ, 763, L7

Emsellem E., et al., 2007, MNRAS, 379, 401

Emsellem E., et al., 2011, MNRAS, 414, 888

Estrada-Carpenter V., et al., 2019, ApJ, 870, 133

Faber S. M., 1972, A&A, 20, 361

Faber S. M., Friel E. D., Burstein D., Gaskell C. M., 1985, ApJS, 57, 711

Faber S. M., et al., 2007, ApJ, 665, 265

Falcón-Barroso J., Sánchez-Blázquez P., Vazdekis A., Ricciardelli E., Cardiel N., Cenarro A. J., Gorgas J., Peletier R. F., 2011, A&A, 532, A95

Fan X., et al., 2006, AJ, 132, 117

Fang J. J., et al., 2018, ApJ, 858, 100

Ferland G. J., et al., 2017, Rev. Mex. Astron. Astrofis., 53, 385

Feroz F., Hobson M. P., 2008, MNRAS, 384, 449

Feroz F., Hobson M. P., Bridges M., 2009, MNRAS, 398, 1601

Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2013, preprint, (arXiv:1306.2144)

Fitzpatrick E. L., 1986, AJ, 92, 1068

Foreman-Mackey D., 2016, The Journal of Open Source Software, 24

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306

Förster Schreiber N. M., et al., 2014, ApJ, 787, 38

Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, ApJ, 327, 507

Friedmann A., 1922, Zeitschrift fur Physik, 10, 377

Fumagalli M., et al., 2016, ApJ, 822, 1

Furusawa H., et al., 2016, ApJ, 822, 46

Gabor J. M., Davé R., 2015, MNRAS, 447, 374

Gabor J. M., Davé R., Finlator K., Oppenheimer B. D., 2010, MNRAS, 407, 749

Gabor J. M., Davé R., Oppenheimer B. D., Finlator K., 2011, MNRAS, 417, 2676

Gallazzi A., Charlot S., Brinchmann J., White S. D. M., Tremonti C. A., 2005, MNRAS, 362, 41

Gallazzi A., Brinchmann J., Charlot S., White S. D. M., 2008, MNRAS, 383, 1439

Gallazzi A., Bell E. F., Zibetti S., Brinchmann J., Kelson D. D., 2014, ApJ, 788, 72

Gargiulo A., et al., 2017, A&A, 606, A113

Garn T., Best P. N., 2010, MNRAS, 409, 421

Gladders M. D., Oemler A., Dressler A., Poggianti B., Vulcani B., Abramson L., 2013, ApJ, 770, 64

Glazebrook K., et al., 2017, Nature, 544, 71

Goodman J., Weare J., 2010, Communications in Applied Mathematics and Computational Science, Vol.˜5, No.˜1, p.˜65-80, 2010, 5, 65

Gorgas J., Faber S. M., Burstein D., Gonzalez J. J., Courteau S., Prosser C., 1993, ApJS, 86, 153

Grogin N. A., et al., 2011, The Astrophysical Journal Supplement Series, 197, 35

Gunn J. E., Peterson B. A., 1965, ApJ, 142, 1633

Han Y., Han Z., 2019, The Astrophysical Journal Supplement Series, 240, 3

Hartley W. G., et al., 2010, MNRAS, 407, 1212

Heavens A. F., Jimenez R., Lahav O., 2000, MNRAS, 317, 965

Heavens A., Panter B., Jimenez R., Dunlop J., 2004, Nature, 428, 625

Heavens A., Fantaye Y., Mootoovaloo A., Eggers H., Hosenie Z., Kroon S., Sellentin E., 2017, preprint, (arXiv:1704.03472)

Heckman T. M., Robert C., Leitherer C., Garnett D. R., van der Rydt F., 1998, ApJ, 503, 646

Herpich F., Stasińska G., Mateus A., Vale Asari N., Cid Fernandes R., 2018, MNRAS, 481, 1774

Hertzsprung E., 1911, Publikationen des Astrophysikalischen Observatoriums zu Potsdam, 63

Hildebrand R. H., 1983, QJRAS, 24, 267

Hogg D. W., 1999, arXiv e-prints, pp astro–ph/9905116

Hogg D. W., Bovy J., Lang D., 2010, preprint, (arXiv:1008.4686)

Hopkins A. M., Connolly A. J., Haarsma D. B., Cram L. E., 2001, AJ, 122, 288

Hoyle F., 1946, MNRAS, 106, 343

Hoyle F., 1954, ApJS, 1, 121

Hu E. M., McMahon R. G., Cowie L. L., 1999, ApJ, 522, L9

Hubble E. P., 1926, ApJ, 64, 321

Hubble E., 1929, Proceedings of the National Academy of Science, 15, 168

Hudelot P., et al., 2012, VizieR Online Data Catalog, 2317

Inoue A. K., Shimizu I., Iwata I., Tanaka M., 2014, MNRAS, 442, 1805

Iyer K., Gawiser E., 2017, ApJ, 838, 127

Iyer K. G., Gawiser E., Faber S. M., Ferguson H. C., Koekemoer A. M., Pacifici C., Somerville R., 2019, preprint, (arXiv:1901.02877)

Jørgensen I., Chiboucas K., 2013, AJ, 145, 77

Kant I., 1755, Allgemeine Naturgeschichte und Theorie des Himmels

Kauffmann G., et al., 2003, MNRAS, 341, 33

Kemp T. W., Dunlop J. S., McLure R. J., Schreiber C., Carnall A. C., Cullen F., 2019, preprint, (arXiv:1903.08169)

Kennicutt Robert C. J., 1998, ApJ, 498, 541

Kennicutt R. C., Evans N. J., 2012, ARA&A, 50, 531

Kewley L. J., Geller M. J., Jansen R. A., 2004, AJ, 127, 2002

Kewley L. J., Dopita M. A., Leitherer C., Davé R., Yuan T., Allen M., Groves B., Sutherland R., 2013, ApJ, 774, 100

Khochfar S., Burkert A., 2003, ApJ, 597, L117

Khochfar S., Silk J., 2009, MNRAS, 397, 506

Khochfar S., et al., 2011, MNRAS, 417, 845

Koekemoer A. M., et al., 2011, The Astrophysical Journal Supplement Series, 197, 36

Kowal C. T., 1968, AJ, 73, 1021

Kriek M., Conroy C., 2013, ApJ, 775, L16

Kroupa P., 2001, MNRAS, 322, 231

Lange R., et al., 2015, MNRAS, 447, 2603

Law-Smith J., Eisenstein D. J., 2017, ApJ, 836, 87

Leavitt H. S., 1908, Annals of Harvard College Observatory, 60, 87

Lee H.-c., Worthey G., Trager S. C., Faber S. M., 2007, ApJ, 664, 215

Lee S.-K., Idzi R., Ferguson H. C., Somerville R. S., Wiklind T., Giavalisco M., 2009, ApJS, 184, 100

Leitner S. N., 2012, ApJ, 745, 149

Leja J., van Dokkum P. G., Franx M., Whitaker K. E., 2015, ApJ, 798, 115

Leja J., Johnson B. D., Conroy C., van Dokkum P. G., Byler N., 2017, ApJ, 837, 170

Leja J., Carnall A. C., Johnson B. D., Conroy C., Speagle J. S., 2018, preprint, (arXiv:1811.03637)

Lemaître G., 1927, Annales de la Société Scientifique de Bruxelles, 47, 49

Lemaux B. C., Lubin L. M., Shapley A., Kocevski D., Gal R. R., Squires G. K., 2010, ApJ, 716, 970

Liddle A., 2003, An Introduction to Modern Cosmology, Second Edition

Lilly S. J., Tresse L., Hammer F., Crampton D., Le Fevre O., 1995, ApJ, 455, 108

Lilly S. J., Le Fevre O., Hammer F., Crampton D., 1996, ApJ, 460, L1

Lilly S. J., Carollo C. M., Pipino A., Renzini A., Peng Y., 2013, ApJ, 772, 119

Loeb A., 2010, How Did the First Stars and Galaxies Form?

Lonoce I., Longhetti M., Saracco P., Gargiulo A., Tamburri S., 2014, MNRAS, 444, 2048

Lotz J. M., Jonsson P., Cox T. J., Croton D., Primack J. R., Somerville R. S., Stewart K., 2011, ApJ, 742, 103

Lovell C. C., Acquaviva V., Thomas P. A., Iyer K. G., Gawiser E., Wilkins S. M., 2019, preprint, (arXiv:1903.10457)

Luo B., et al., 2017, The Astrophysical Journal Supplement Series, 228, 2

Madau P., 1995, ApJ, 441, 18

Madau P., Dickinson M., 2014, ARA&A, 52, 415

Madau P., Ferguson H. C., Dickinson M. E., Giavalisco M., Steidel C. C., Fruchter A., 1996, MNRAS, 283, 1388

Madau P., Pozzetti L., Dickinson M., 1998, ApJ, 498, 106

Maiolino R., Mannucci F., 2019, Astronomy and Astrophysics Review, 27, 3

Maiolino R., et al., 2008, A&A, 488, 463

Maiolino R., et al., 2012, MNRAS, 425, L66

Maltby D. T., et al., 2016, MNRAS, 459, L114

Maltby D. T., Almaini O., Wild V., Hatch N. A., Hartley W. G., Simpson C., Rowlands K., Socolovsky M., 2018, MNRAS, 480, 381

Mannucci F., Cresci G., Maiolino R., Marconi A., Gnerucci A., 2010, MNRAS, 408, 2115

Maraston C., Pforr J., Renzini A., Daddi E., Dickinson M., Cimatti A., Tonini C., 2010, MNRAS, 407, 830

McCracken H. J., et al., 2012, A&A, 544, A156

McLeod D. J., McLure R. J., Dunlop J. S., Robertson B. E., Ellis R. S., Targett T. A., 2015, MNRAS, 450, 3032

McLeod D. J., McLure R. J., Dunlop J. S., 2016, MNRAS, 459, 3812

McLure R. J., Dunlop J. S., Cirasuolo M., Koekemoer A. M., Sabbi E., Stark D. P., Targett T. A., Ellis R. S., 2010, MNRAS, 403, 960

McLure R. J., et al., 2011, MNRAS, 418, 2074

McLure R. J., et al., 2013, MNRAS, 428, 1088

McLure R. J., et al., 2018a, MNRAS,

McLure R. J., et al., 2018b, MNRAS, 476, 3991

Merlin E., et al., 2018, MNRAS, 473, 2098

Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, The Journal of Chemical Physics, 21, 1087

Mobasher B., et al., 2015, ApJ, 808, 101

Moehler S., et al., 2014, A&A, 568, A9

Moresco M., et al., 2010, A&A, 524, A67

Morgan W. W., Keenan P. C., Kellman E., 1943, An atlas of stellar spectra, with an outline of spectral classification

Morishita T., et al., 2018, preprint, (arXiv:1812.06980)

Mortlock D. J., et al., 2011, Nature, 474, 616

Mortlock A., McLure R. J., Bowler R. A. A., McLeod D. J., Mármol-Queraltó E., Parsa S., Dunlop J. S., Bruce V. A., 2017, MNRAS, 465, 672

Moustakas J., Kennicutt Jr. R. C., Tremonti C. A., 2006, ApJ, 642, 775

Muzzin A., et al., 2013, ApJS, 206, 8

Narayanan D., Conroy C., Davé R., Johnson B. D., Popping G., 2018, ApJ, 869, 70

Nelson D., et al., 2018, MNRAS, 475, 624

Nelson D., et al., 2019, preprint, (arXiv:1902.05554)

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

Nogueira-Cavalcante J. P., Gonçalves T. S., Menéndez-Delmestre K., Sheth K., 2018, MNRAS, 473, 1346

Ocvirk P., Pichon C., Lançon A., Thiébaut E., 2006, MNRAS, 365, 46

Oke J. B., Gunn J. E., 1983, ApJ, 266, 713

Onodera M., et al., 2012, ApJ, 755, 26

Onodera M., et al., 2015, ApJ, 808, 161

Oort J. H., 1940, ApJ, 91, 273

Osterbrock D. E., 1989, Astrophysics of gaseous nebulae and active galactic nuclei

Pacifici C., Charlot S., Blaizot J., Brinchmann J., 2012, MNRAS, 421, 2002

Pacifici C., et al., 2015, MNRAS, 447, 786

Pacifici C., et al., 2016, ApJ, 832, 79

Panter B., Heavens A. F., Jimenez R., 2003, MNRAS, 343, 1145

Panter B., Jimenez R., Heavens A. F., Charlot S., 2007, MNRAS, 378, 1550

Panter B., Jimenez R., Heavens A. F., Charlot S., 2008, MNRAS, 391, 1117

Papovich C., Dickinson M., Ferguson H. C., 2001, ApJ, 559, 620

Papovich C., et al., 2010, ApJ, 716, 1503

Papovich C., et al., 2012, ApJ, 750, 93

Peacock J. A., Jimenez R., Dunlop J. S., Waddington I., Spinrad H., Stern D., Dey A., Windhorst R. A., 1998, MNRAS, 296, 1089

Peacock J. A., et al., 2001, Nature, 410, 169

Pearson K., 1900, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 50, 157

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

Peng Y., Maiolino R., Cochrane R., 2015, Nature, 521, 192

Pentericci L., et al., 2018, A&A, 616, A174

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

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

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

Pillepich A., et al., 2018, MNRAS, 473, 4077

Planck Collaboration et al., 2018, preprint, (arXiv:1807.06209)

Press W. H., Schechter P., 1974, ApJ, 187, 425

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

Rayleigh J. W., 1879, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 8, 261

Reddy N. A., Pettini M., Steidel C. C., Shapley A. E., Erb D. K., Law D. R., 2012, ApJ, 754, 25

Reddy N. A., et al., 2015, ApJ, 806, 259

Reddy N. A., et al., 2018, ApJ, 853, 56

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

Robertson B. E., et al., 2013, ApJ, 768, 71

Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, ApJ, 802, L19

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

Russell H. N., 1914, Popular Astronomy, 22, 275

Sabater J., et al., 2019, A&A, 622, A17

Salim S., et al., 2016, ApJS, 227, 2

Salmon B., et al., 2015, ApJ, 799, 183

Salmon B., et al., 2016, ApJ, 827, 20

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

Sawicki M. J., Lin H., Yee H. K. C., 1997, AJ, 113, 1

Schawinski K., et al., 2014, MNRAS, 440, 889

Schechter P., 1976, ApJ, 203, 297

Schreiber C., Elbaz D., Pannella M., Ciesla L., Wang T., Koekemoer A., Rafelski M., Daddi E., 2016, A&A, 589, A35

Schreiber C., et al., 2018, A&A, 618, A85

Shanks T., Stevenson P. R. F., Fong R., MacGillivray H. T., 1984, MNRAS, 206, 767

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

Simpson C., et al., 2006, MNRAS, 372, 741

Simpson C., Westoby P., Arumugam V., Ivison R., Hartley W., Almaini O., 2013, MNRAS, 433, 2647

Simpson F., Jimenez R., Pena-Garay C., Verde L., 2017, J. Cosmology Astropart. Phys., 6, 029

Singh R., et al., 2013, A&A, 558, A43

Siudek M., et al., 2017, A&A, 597, A107

Skilling J., 2006, Bayesian Anal., 1, 833

Slipher V. M., 1913, Lowell Observatory Bulletin, 1, 56

Slipher V. M., 1917, Popular Astronomy, 25, 36

Smethurst R. J., et al., 2018, MNRAS, 473, 2679

Somerville R. S., Davé R., 2015, Annual Review of Astronomy and Astrophysics, 53, 51

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

Speagle J. S., 2019, preprint, p. arXiv:1904.02180 (arXiv:1904.02180)

Speagle J. S., Steinhardt C. L., Capak P. L., Silverman J. D., 2014, ApJS, 214, 15

Spinrad H., Stern D., Bunker A., Dey A., Lanzetta K., Yahil A., Pascarelle S., Fernández-Soto A., 1998, AJ, 116, 2617

Springel V., Di Matteo T., Hernquist L., 2005, MNRAS, 361, 776

Steidel C. C., Strom A. L., Pettini M., Rudie G. C., Reddy N. A., Trainor R. F., 2016, ApJ, 826, 159

Steinhardt C. L., et al., 2014, ApJ, 791, L25

Straatman C. M. S., et al., 2014, ApJ, 783, L14

Straatman C. M. S., et al., 2016, ApJ, 830, 51

Straatman C. M. S., et al., 2018, The Astrophysical Journal Supplement Series, 239, 27

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

Strazzullo V., et al., 2013, ApJ, 772, 118

Tasse C., Best P. N., Röttgering H., Le Borgne D., 2008, A&A, 490, 893

Taylor E. N., et al., 2011, MNRAS, 418, 1587

Thomas D., Maraston C., Schawinski K., Sarzi M., Silk J., 2010, MNRAS, 404, 1775

Thomas R., et al., 2017, A&A, 602, A35

Tojeiro R., Heavens A. F., Jimenez R., Panter B., 2007, MNRAS, 381, 1252

Tomczak A. R., et al., 2014, ApJ, 783, 85

Trager S. C., Somerville R. S., 2009, MNRAS, 395, 608

Tran K.-V. H., Franx M., Illingworth G., Kelson D. D., van Dokkum P., 2003, ApJ, 599, 865

Trayford J. W., et al., 2017, MNRAS, 470, 771

Tremonti C. A., et al., 2004, ApJ, 613, 898

Trotta R., 2008, Contemporary Physics, 49, 71

Ueda Y., et al., 2008, ApJS, 179, 124

Vogelsberger M., et al., 2014, MNRAS, 444, 1518

Wechsler R. H., Tinker J. L., 2018, Annual Review of Astronomy and Astrophysics, 56, 435

Wechsler R. H., Bullock J. S., Primack J. R., Kravtsov A. V., Dekel A., 2002, ApJ, 568, 52

Weinberger R., et al., 2017, MNRAS, 465, 3291

Whitaker K. E., et al., 2011, ApJ, 735, 86

Whitaker K. E., van Dokkum P. G., Brammer G., Franx M., 2012, ApJ, 754, L29

Whitaker K. E., et al., 2013, ApJ, 770, L39

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

Wild V., Kauffmann G., Heckman T., Charlot S., Lemson G., Brinchmann J., Reichard T., Pasquali A., 2007, MNRAS, 381, 543

Wild V., Walcher C. J., Johansson P. H., Tresse L., Charlot S., Pollo A., Le Fèvre O., de Ravel L., 2009, MNRAS, 395, 144

Wild V., et al., 2014, MNRAS, 440, 1880

Wild V., Almaini O., Dunlop J., Simpson C., Rowlands K., Bowler R., Maltby D., McLure R., 2016, MNRAS, 463, 832

Wilkinson D. M., Maraston C., Goddard D., Thomas D., Parikh T., 2017, MNRAS, 472, 4297

Williams R. J., Quadri R. F., Franx M., van Dokkum P., Labbé I., 2009, ApJ, 691, 1879

Wisnioski E., et al., 2015, ApJ, 799, 209

Worthey G., Faber S. M., Gonzalez J. J., Burstein D., 1994, ApJS, 94, 687

Wright T., 1750, An Original Theory or New Hypothesis of the Universe

Wright A. H., et al., 2016, MNRAS, 460, 765

Wright A. H., et al., 2017, MNRAS, 470, 283

Wu P.-F., et al., 2018a, ApJ, 855, 85

Wu P.-F., et al., 2018b, ApJ, 855, 85

Wu P.-F., et al., 2018c, ApJ, 868, 37

Wuyts S., Franx M., Cox T. J., Hernquist L., Hopkins P. F., Robertson B. E., van Dokkum P. G., 2009, ApJ, 696, 348

Wuyts S., et al., 2011, ApJ, 738, 106

Xiang M. S., et al., 2015, MNRAS, 448, 90

Yan R., Newman J. A., Faber S. M., Konidaris N., Koo D., Davis M., 2006, ApJ, 648, 281

Yan R., et al., 2016, AJ, 151, 8

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

Younger J. D., et al., 2009, MNRAS, 394, 1685

Zahid H. J., Kudritzki R.-P., Conroy C., Andrews B., Ho I. T., 2017, ApJ, 847, 18

Zwicky F., 1937, ApJ, 86, 217

al-Sufi A., ca. 964, Book of Fixed Stars

da Cunha E., Charlot S., Elbaz D., 2008, MNRAS, 388, 1595

de Zeeuw P. T., et al., 2002, MNRAS, 329, 513

van der Wel A., et al., 2014, ApJ, 788, 28

van der Wel A., et al., 2016, ApJS, 223, 29