Environmental quenching and hierarchical cluster assembly: evidence from spectroscopic ages of red-sequence galaxies in Coma

Russell J. Smith,[1]★ John R. Lucey,[1] James Price,[2] Michael J. Hudson[3,4] and Steven Phillipps[2]

[1]*Department of Physics, University of Durham, Durham DH1 3LE*
[2]*Astrophysics Group, H. H. Wills Physics Laboratory, University of Bristol, Tyndall Avenue, Bristol BS8 1TL*
[3]*Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada*
[4]*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario, N2L 2Y5, Canada*

ABSTRACT

We explore the variation in stellar population ages for Coma cluster galaxies as a function of projected cluster-centric distance, using a sample of 362 red-sequence galaxies with high signal-to-noise ratio spectroscopy. The sample spans a wide range in luminosity (0.02–$4\,L^*$) and extends from the cluster core to near the virial radius. We find a clear distinction in the observed trends of the giant and dwarf galaxies. The ages of red-sequence giants are primarily determined by galaxy mass, whether parametrized by velocity dispersion, luminosity, stellar mass or dynamical mass, with only weak modulation by environment, in the sense that galaxies at larger cluster-centric distance are slightly younger. For red-sequence dwarfs (with mass $\lesssim 10^{10}\,M_\odot$), the roles of mass and environment as predictors of age are reversed: there is little dependence on mass, but strong trends with projected cluster-centric radius are observed. The average age of dwarfs at the 2.5 Mpc limit of our sample is approximately half that of dwarfs near the cluster centre. The gradient in dwarf galaxy ages is a global cluster-centric trend, and is not driven by the ongoing merger of the NGC 4839 group to the south west of Coma. We interpret these results using environmental histories extracted from the *Millennium Simulation* for members of massive clusters. Hierarchical cluster assembly naturally leads to trends in the accretion times of galaxies as a function of projected cluster-centric radius. On average, simulated galaxies now located in cluster cores joined haloes above any given mass threshold earlier than those now in the outskirts of clusters. We test environmental quenching models, in which star formation is halted in galaxies when they enter haloes of a given mass, or when they become satellites. Our models broadly reproduce the gradients observed in Coma, but for dwarf galaxies the efficiency of environmental quenching must be very high to match the strong trends observed.

Key words: galaxies: clusters: general – galaxies: clusters: individual: Coma – galaxies: dwarf – galaxies: elliptical and lenticular, cD – galaxies: evolution.

1 INTRODUCTION

According to the current cosmological paradigm, galaxy formation and evolution takes place within a hierarchy of structure formation governed by the merging of dark-matter haloes. The richest galaxy clusters at $z = 0$, with total masses of $\gtrsim 10^{15}\,M_\odot$, occupy the apex of this hierarchy. By definition, such clusters are highly unrepresentative of the present-epoch Universe. Moreover, the galaxies now

collected within such clusters formed and evolved in regions of above-average density at *all* epochs of their history, leading to evolution that was accelerated compared to those in more representative regions.

The evolution of galaxies is linked to their surroundings through a rich variety of environmental processes (e.g. Boselli & Gavazzi 2006, and references therein). Most dramatically, galaxies may merge with others of similar mass, profoundly affecting their structure and triggering massive star-bursts which consume their gas supply on short time-scales. Mergers with less massive galaxies are individually less transformative, but cumulatively may dominate

★E-mail: russell.smith@durham.ac.uk

the assembly of stellar mass in some galaxies. Even where mergers do not occur, fly-by tidal interactions with other galaxies, or the larger-scale gravitational potential, can re-distribute stars leading to morphological transformation, or perhaps cause tidal instabilities leading to gas being shocked into star formation. Galaxies can also be influenced by interaction with their gaseous surroundings: rapid motion through a dense intra-cluster medium can lead to ram-pressure stripping of cold interstellar gas in discs as first discussed by Gunn & Gott (1972), or else the reservoir of hot halo gas may be stripped, eventually starving the disc of continued gas supply (Larson, Tinsley & Caldwell 1980; Balogh, Navarro & Morris 2000). Qualitatively, the end result of these processes is to exhaust or remove the gas from galaxies in denser environments, leading to a shutdown in star formation ('quenching'), as required to explain population correlations such as the morphology–density relation (e.g. Dressler 1980). At a more detailed level, however, the relative importance of the various processes, and how the contributions depend on local and global environment, and vary over cosmic time, remain crucial areas of uncertainty.

Despite being unrepresentative regions by definition, rich clusters are important because they provide observationally convenient galaxy samples for an intensive study of galaxy properties, and also because *some* of the possible quenching mechanisms, such as ram-pressure stripping of cold gas, are probably effective *only* in the densest, most massive systems. Galaxies with stripped tails directed away from the centres of clusters provide evidence for on-going stripping in massive clusters (e.g. Chung et al. 2007; Smith et al. 2010; Sun et al. 2010), showing that clusters are not merely passive recipients of dead galaxies. On the other hand, since many galaxies spend much of their histories in group-like environments, it is likely that these, rather than clusters, are dominant in driving evolution in the galaxy population globally, and that many cluster galaxies were 'pre-processed' in groups prior to accretion into their current environment (Zabludoff & Mulchaey 1998). Hence in clusters, we are likely observing the remnants of quenching processes that occurred in a range of environments, over a range of epochs, and subsequently mixed as the system assembles.

In this paper, we investigate the characteristic stellar ages of red-sequence (i.e. quenched) galaxies as a function of their location within a single very massive low-redshift cluster, Coma. Unlike 'instantaneous' measures of galaxy properties, such as current star formation rates (SFRs), or membership of the red-sequence or blue cloud, the stellar ages in principle probe the *history* of quenching in the galaxy population. Following previous spectroscopic work which suggested a strong spatial dependence of stellar populations in Coma galaxies (e.g. Guzmán et al. 1992; Caldwell et al. 1993; Carter et al. 2002), we conducted a high-signal-to-noise (S/N) ratio study of faint red-sequence galaxies in this cluster. Earlier results from our work (Smith et al. 2008, 2009a, hereafter S09) were based on observations in the core and in a region to the south west, where the NGC 4839 group is merging with the main cluster. We confirmed an age–radius relation for dwarfs which was much steeper than that recovered for more massive galaxies, in an ensemble of clusters, by Smith et al. (2006). However, the S09 analysis suffered from its limited spatial coverage, and in particular could not distinguish between a global age-versus-radius correlation and a localized excess of young galaxies associated with the sub-cluster merger. Studies of other clusters have provided additional hints for strong environmental trends for dwarfs (Chilingarian et al. 2008; Michielsen et al. 2008), which lend support to the former interpretation, but are based on samples of limited size and radial extent. Moreover, because S09 concentrated only on the dwarfs, the mass-

dependence of the environmental trends in Coma was not uniformly investigated.

In this paper, we extend the work of S09 using deep MMT observations for an enlarged sample of dwarfs, including additional outer fields which remove the bias towards the south-west region. Furthermore, where S09 used data for galaxies in a different location (the Shapley Supercluster) to provide the high-luminosity comparison sample, here we use spectra from the Sloan Digital Sky Survey (SDSS) which, from Data Release 7 (Abazajian et al. 2009), covers the whole Coma region. The treatment of the SDSS data has been described by Price et al. (2011). Together, the SDSS and MMT data sets cover a 6 mag range in luminosity, and cover most of the area inside the nominal virial radius of the cluster, enabling a comprehensive re-analysis of the cluster-centric dependence of passive galaxy ages.

We interpret the observational results explicitly in the context of hierarchical structure formation, through comparison with cluster assembly histories drawn from the *Millennium Simulation* (Springel et al. 2005). Specifically, we test models in which the cessation of star formation is linked to the incorporation of galaxies into haloes above a given threshold, or change of status from central galaxy to satellite within a halo. Such models lead generically to halo-centric gradients in stellar population age, since dynamical friction causes the earliest-accreted sub-haloes to sink towards the cluster centre, while the most recently added structures reside at larger radius (e.g. Gao et al. 2004). We determine how the strength of the predicted gradients depends on the details of the adopted quenching criteria and compare the predictions with our measurements in Coma.

The remainder of the paper is organized as follows. Section 2 covers the observational results, including a summary of the sample and data employed (Section 2.1) and an investigation of the cluster-centric dependence of galaxy ages (Section 2.2). Section 3 presents the modelling results, starting by exploring how the typical accretion history of cluster galaxies relates to their eventual location within the cluster halo (Section 3.1), then using this as a basis for environmental quenching models (Section 3.2). The predictions are confronted with the measured trends in Section 3.3, and some caveats and limitations to the analysis are noted in Section 3.4. We conclude with a discussion of the results in the context of previous work in Section 4.

In converting to physical units, we adopt a distance of 100 Mpc for Coma (i.e. $h = 0.72$), so that $1°$ corresponds to 1.74 Mpc. For reference, the virial radius of Coma has been estimated to be 2.9 Mpc, and its mass within this radius is $\sim 1.4 \times 10^{15} \, M_\odot$ (Łokas & Mamon 2003).

2 COMA OBSERVATIONS AND AGE CORRELATIONS

2.1 Sample, data and parameter measurements

Broadly, the overall sample is based on spectra for 463 bright (~ 0.2–$4.0L^*$) Coma cluster members from the SDSS (analysed in Price et al. 2011) and 198 fainter members (~ 0.02–$0.20L^*$) observed using long integrations with the Hectospec multi-fibre spectrograph at the 6.5-m MMT (Fabricant et al. 2005). A detailed description of the sample construction, parameter measurements and methods, and data presentation, will be presented in a separate paper (Smith et al., in preparation, hereafter S11). S11 will also provide a full analysis of the scaling relations of metallicity and element abundance ratios.

Our MMT data set is an extended version of that reported in S09. In that paper, the data were drawn from two 1° diameter Hectospec fields, one centred on the cluster core and the other displaced nearly a degree south-west to sample the region around the NGC 4839 subcluster. As discussed above, this limited coverage impeded a clear interpretation of the results. To address this, we subsequently observed a further four outer fields in a hexagonal pattern around the core (west, north-west, north-east and east) to yield almost complete coverage of the region within the cluster virial radius (we were unable to observe the final planned field to the south-east).

Our Hectospec observations employed the 270 line mm^{-1} grating, which delivers spectra with wide wavelength coverage (3700–9000 Å) at a resolution of 4.5 Å full width at half-maximum. In each run, the high-S/N spectra required for the stellar population analysis were obtained in parallel with a redshift survey aimed at establishing cluster members to much deeper magnitude limits (Marzke et al., in preparation). Many fibre configurations were observed, each with 1 hour of integration, but fibres were repeatedly allocated to galaxies in the stellar populations sample, in order to build up total exposure times of 4–8 h per object. The galaxy sample for the stellar populations study was drawn from known cluster members with luminosity 2–4 mag fainter than M^*. We took into account optical colours to select passive galaxies, and the SDSS spectra where available, to exclude galaxies with emission at Hα (since this indicates nebular contamination of the higher-order Balmer lines used for age determination). The spectra were reduced as described in S09. Data from all runs were reprocessed together to ensure that the S09 data and the later observations were treated identically in the reduction stages. Error spectra were computed alongside the combined spectrum for each galaxy for use in estimating errors on the measured parameters.

As a complement to the Hectospec observations at low luminosity, we performed a broadly similar analysis of the available spectra from the SDSS DR7, covering approximately the same region in the Coma cluster. The treatment of the SDSS data is described in detail by Price et al. (2011). For building the sample used in this paper, we relaxed the magnitude and S/N limits imposed in the Price et al. analysis, so that we can analyse all available SDSS spectra within 2° radius of the Coma cluster. (We later impose a selection on age error to remove low-S/N measurements consistently across both data sources.)

The combined MMT and the SDSS sample has been homogeneously re-analysed to measure absorption line-strength indices and emission line equivalent widths. Stellar population ages ($T_{\rm spec}$), along with metallicities (Fe/H) and abundance ratios (Mg/Fe, Ca/Fe, C/Fe and N/Fe), were measured via the comparison of the index data against the Schiavon (2007) simple stellar population (SSP) models, using a new model inversion code described in S11. The measured quantity is strictly an SSP-equivalent age, i.e. the age of a single burst which best reproduces the observed indices. When comparing to model predictions in Section 3.3, we explicitly take into account the relationship between $T_{\rm spec}$ and the true (non-SSP) star formation history. To complement the line-strength data, velocity dispersions were compiled from the SDSS DR7 and a variety of literature sources (mainly Jørgensen 1999; Moore et al. 2002; Smith et al. 2004; Matković & Guzmán 2005; Cody et al. 2009), supplemented with new observations from Very Large Telescope (VLT)/FLAMES, and combined using observations from multiple data sources to determine relative systematic offsets. Seeing-corrected half-light radii were measured from the r-band SDSS images, taking into account improved background estimation methods. The radii are used for

applying aperture corrections, and for deriving approximate dynamical masses.

There are a total of 198 galaxies in the MMT sample and 463 in the SDSS sample, but some are common to both sets, some (especially in the SDSS set) have strong emission lines which contaminate the stellar Balmer absorption and some have insufficient S/N for a reliable stellar population analysis. We define the sample for analysis in this paper by requiring less than 0.3 dex error in the derived age, and an equivalent width of less than 1 Å in emission at Hα. Additionally, we remove four galaxies from the SDSS data set that have $g - r < 0.5$, substantially bluer than all other objects in the sample, and also one galaxy from MMT that is much fainter than the nominal limit of $r \approx 18$. For galaxies common to both data sources, the MMT measurement is used, since its S/N is always much greater than the SDSS spectrum (a factor of ∼3, on average). These cuts reduce the sample to 411 galaxies in total, of which 169 have ages measured from the MMT and 242 from the SDSS. Velocity dispersion data are available for a more restricted sample, comprising 355 galaxies, of which 120 have ages from the MMT and 235 from the SDSS.

Finally, for this paper, we further restrict our attention to the spatial regions of the Coma cluster that were targeted in the MMT observations, such that the sampling is fairly similar, as a function of magnitude, at all locations. Removing the SDSS galaxies that lie in regions without the MMT coverage (i.e. beyond a radius 1°.5, or beyond 0°.5 in the missing south-east wedge 110°–190°, measured east from north), the sample is reduced to 362 galaxies of which 310 have velocity dispersion measurements. For the final sample, the median S/N is 41 Å^{-1}, and the median error in age is 0.14 dex. Because the faint part of the sample is drawn from the higher-S/N MMT data, the data quality does not degrade rapidly for dwarf galaxies: for galaxies below the median luminosity of the sample ($r_{\rm petro} > 15.6$), the median S/N is 42 Å^{-1}, and the median error in age is 0.12 dex. Fig. 1 shows the colour–magnitude relation and sky distribution for the final sample.

2.2 Observed trends of age with environment

The characteristic stellar ages of red-sequence galaxies are known to depend strongly on their 'mass', as represented by various proxies such as velocity dispersion and luminosity (e.g. Caldwell, Rose & Concannon 2003; Nelan et al. 2005; Thomas et al. 2005; Smith, Lucey & Hudson 2009b), and there is a tendency for at least the most massive ellipticals to be located at the centres of clusters. Hence to recover secondary correlations with environment, for a sample spanning a wide mass range, it is desirable to control for the dependence on mass.

In this paper, we use four different mass-related observables to remove the mass dependence and hence isolate the environmental trends.

(i) Central velocity dispersion, σ, is generally found to be the quantity yielding the tightest scaling relations (e.g. Smith et al. 2009b), but is not measured for all of the galaxies in the sample.

(ii) Luminosity, L_r, is measured for all galaxies, and is cleanest in terms of the sample selection criteria, but favours younger, brighter galaxies at given stellar mass.

(iii) Estimated stellar mass, $M_{\rm stel}$, computed from L_r and the spectroscopic age and metallicity, is available for all galaxies, and is in principle a more physically meaningful quantity than luminosity. We use the stellar mass-to-light ratio from the models of Maraston (2005), assuming single-burst star formation histories and Kroupa
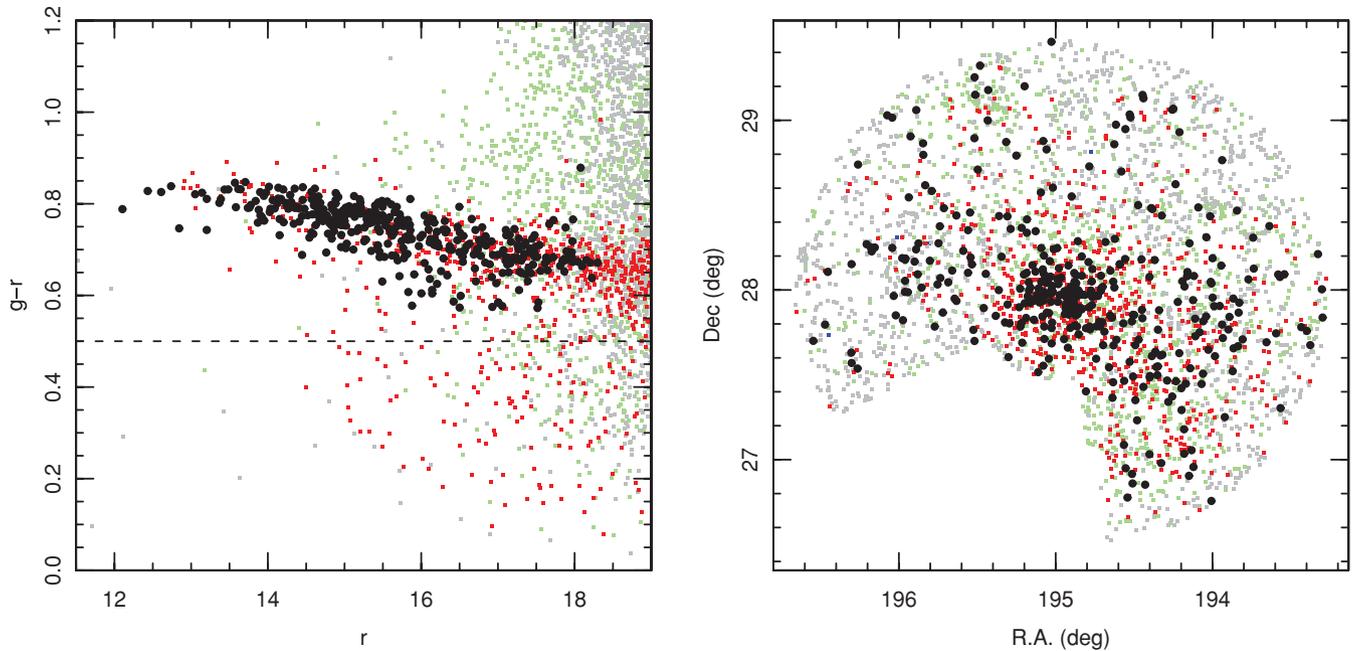
Figure 1. The combined SDSS and MMT stellar populations data set as analysed in this paper (black points), in colour–magnitude space (left) and projected on the sky (right). Photometry and positions are from the SDSS DR7 as used in the sample construction. (The galaxy at $r \approx 18$, $g - r \approx 0.9$ is affected by a poorly deblended star in DR7; its corrected colour in DR8 is 0.70.) For comparison, we show all other SDSS galaxies within the same spatial region, as small points coloured according to their cluster membership status: red points are spectroscopically confirmed cluster members (with $3000 < cz < 11000$ km s^{-1}), green points are confirmed non-members and grey points are galaxies with unknown redshift. Membership information includes data from our MMT/Hectospec redshift survey (Marzke et al., in preparation). The dashed line in the colour–magnitude diagram shows the colour cut of $g - r > 0.5$ used in defining the stellar populations sample.

(2001) initial mass function. Measurement errors in age translate into $M_{\rm stel}$, hence the errors are correlated.

(iv) Dynamical mass $M_{\rm dyn}$, computed from σ and the half-light radius $R_{\rm h}$, via $M_{\rm dyn} = 5R_{\rm h}\sigma^2/G$, is an alternative physical indicator for galaxy mass. This quantity does not suffer from error correlations, but is only available for those galaxies with velocity dispersion data.

The global age–mass correlations for the full sample are shown in the upper row of Fig. 2, where each panel shows results for a different mass proxy. As discussed in greater depth in S11, the sample shows an apparent break in the stellar population scaling relations, at a mass of $\sim 10^{10}$ M$_\odot$ or $\sigma \approx 90$ km s^{-1}. The trend towards younger stellar populations in less massive galaxies tends to flatten or even reverse at low mass, as previously noted by Allanson et al. (2009). To capture this behaviour, we adopt a broken-stick regression model, in which a different slope is allowed on each side of a threshold mass which is itself a free parameter in the fit. We weight galaxies in the fit by $(\sigma_i^2 + \sigma_0^2)^{-1}$ where σ_i are the age errors on the individual data points and σ_0 is a constant intrinsic scatter component determined iteratively by requiring $\chi_\nu^2 = 1$. The errors in the mass proxies are not accounted for in the fit.

In the subsequent rows of Fig. 2, we present the age residuals from the corresponding age–mass relation in the top row, as a function of the cluster-centric radius $R_{\rm proj}$, adopting a cluster centre mid-way between the central giant galaxies NGC 4874 and 4889. The slopes measured using each mass proxy are summarized in Table 1. In the second row of the figure, we show the residuals for all galaxies in the sample. A linear fit to the residuals yields a negative slope of 0.06–0.08 dex Mpc^{-1}. (We omit the sign in the text; all quoted gradients are negative.) More-

over, stronger trends are observed for the lowest-mass galaxies. To demonstrate this, we select and fit separately the 'dwarf' galaxies, defined as having masses in the lowest quartile of each mass proxy, specifically $\sigma < 56$ km s^{-1}; $L_r < 1.5 \times 10^9$ L$_\odot$; $M_{\rm stel} < 3.4 \times 10^9$ M$_\odot$; $M_{\rm dyn} < 5.8 \times 10^9$ M$_\odot$. Note that this cut is located at a lower mass than the break in the age–mass relations. For the dwarf subset, the recovered slopes are 0.13–0.15 dex Mpc^{-1}, while for 'giants' (defined to include all galaxies above the dwarf cut, and hence still extending to fairly low masses), the gradients are significantly weaker, at 0.04–0.07 dex Mpc^{-1}. The residual trends for giants and dwarfs are shown in the third and fourth rows of Fig. 2. These results are all robust with respect to the choice of variable used to remove the scaling with 'mass'. Indeed even fitting the gradient directly to the measured ages *without* controlling for mass, we recover qualitatively similar results, albeit with increased error due to larger scatter.

Cutting the sample further reveals some hints at the origin of the radial trend. First, limiting the fit to within ~ 1 Mpc steepens the environmental dependence to 0.08–0.12 dex Mpc^{-1}. This change can be traced to the dwarfs, which show a slope of 0.19–0.34 dex Mpc^{-1} inside 1 Mpc, though with larger error, given the restricted sample size and extent. The slope for the giants is unchanged by restricting the radial range. If the south-west sector (radius $0°.5$–$1°.5$ bounded at angles 190°–240° measured east from north) covering the NGC 4839 merging group is removed, the recovered trends are unchanged with respect to the original result, for giants and dwarfs alike. Fitting instead *only* the core ($<1°$) and the south-west sector, we recover slopes consistent with the original results for the giants and the full sample. For the dwarfs, however, this subset does show a slightly steeper gradient (0.15–0.20 dex Mpc^{-1}). Hence, while the trend is certainly not localized to the south-west, the age dependence for
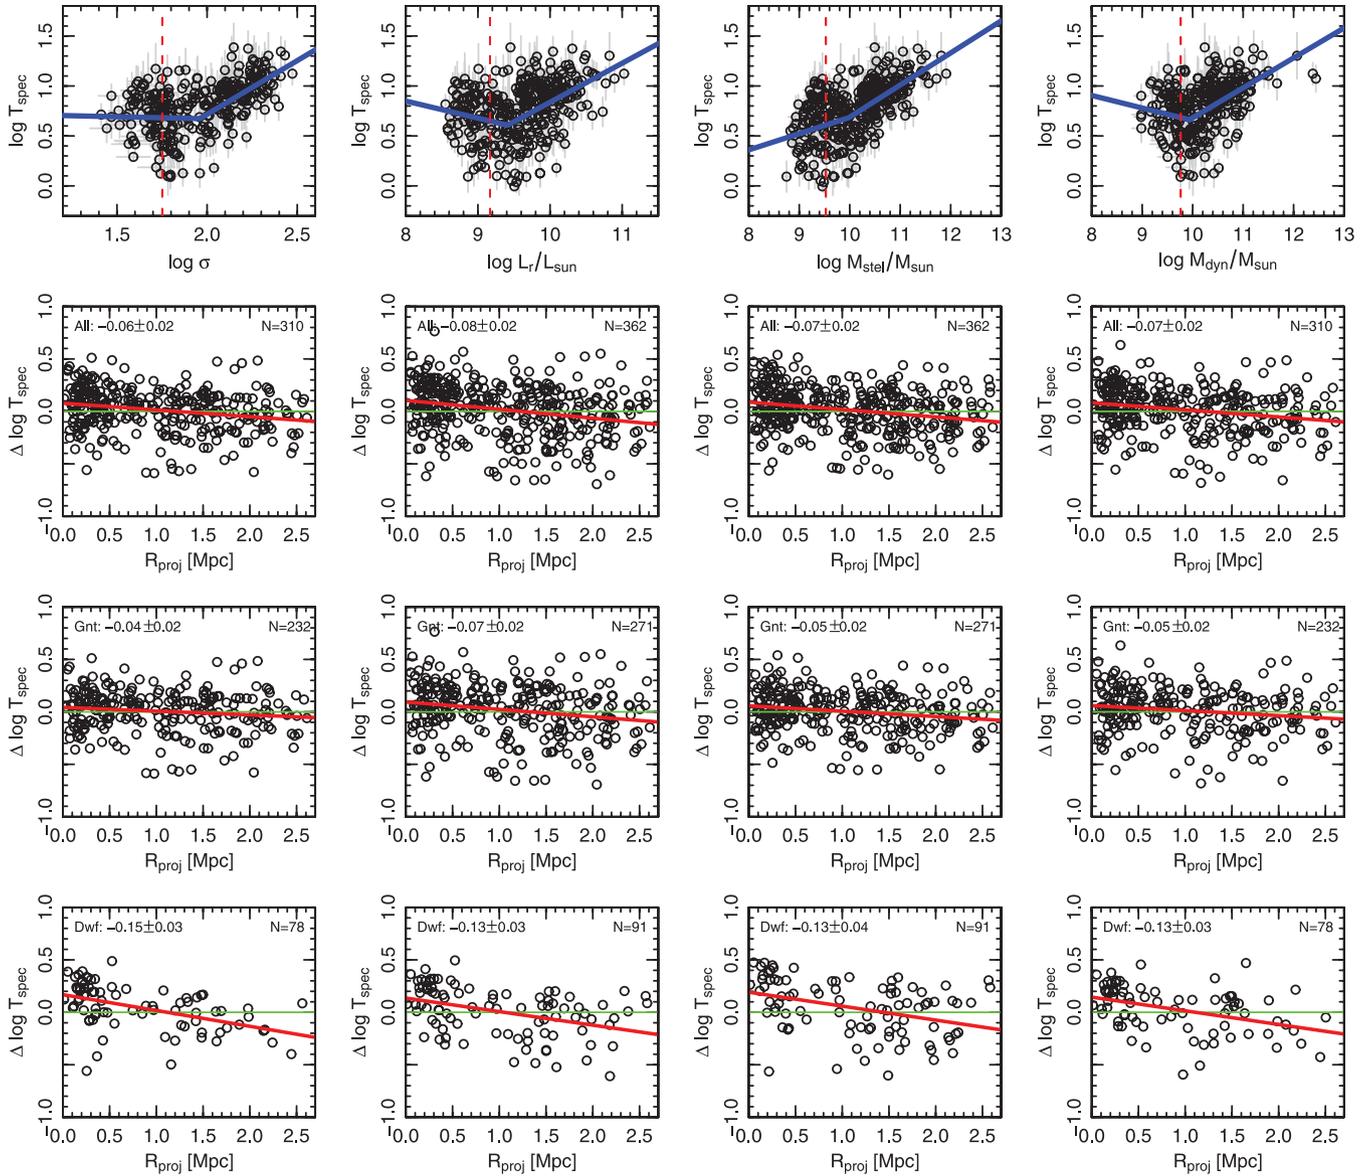
Figure 2. Top row: age versus mass relations, using a variety of 'mass-proxy' quantities, left to right: velocity dispersion, luminosity, stellar mass and dynamical mass. In each case, the heavy blue line segments show a broken-stick model fit to the data. The vertical dashed red line shows the position of the lower quartile of the mass distribution. Second row: residuals in age, relative to the corresponding age–mass relation in the top row, plotted against projected distance from the cluster centre, for all galaxies. The red line shows a linear fit to the trend, with slope noted above. A zero-slope reference line is shown in green. Third row: as for the second row, but showing only the galaxies in the upper three quartiles of the mass distribution ('giants'), and a fit to this sub-sample. Bottom row: equivalent for the galaxies in the lowest quartile of the mass distribution ('dwarfs'). A significant trend for younger ages in the outer galaxies, especially dwarfs, is recovered with similar amplitude regardless of which mass proxy is adopted.

dwarfs might be a little steeper in the direction of the infalling group than elsewhere.

We have seen that the relationship between age, mass and cluster-centric distance is complex, having noted at least three non-linear aspects to the correlations: (1) an apparent change in the dependence on mass, between giant and dwarf regimes; (2) steeper radial correlations in the inner part of the cluster and (3) stronger radial trends for dwarfs rather than giants. To illustrate this behaviour with minimal parametrization, Fig. 3 shows smoothed maps of the average spectroscopic age in the plane defined by galaxy 'mass' and R_{proj}. The age map provides a striking visual confirmation of the results described above: at high mass, the near-vertical contours describe a dominant age–mass relation, modulated by a weak radial trend, while at the low-mass end, the contours bend around to describe

a relation dominated by the cluster-centric radius dependence. As before, the key results are seen to be insensitive to which particular mass proxy is adopted.

Finally, in Fig. 4, we use the same method to map the residuals from the age–mass relationships as a function of position on the sky. These maps provide a way to visualize the variation of average ages with azimuthal angle, confirming that it is primarily the very core of the cluster, rather than the south-west region, that is 'different' from the rest of the cluster. In more detail, the maps suggest that the 'oldest' region of the cluster is displaced slightly west from NGC 4874, and extended towards the south-west. It is tempting to speculate that this feature in the age maps might be related to the extended plume of diffuse light in this area (Welch & Sastry 1971; Gregg & West 1998). For instance, perhaps this region harbours

Table 1. Observed gradients in $\Delta \log T_{spec}$ (the residual from the age–mass relation) as a function of projected cluster-centric distance. The units are decades in age per Mpc in cluster-centric distance. Results are given for the full sample and for various subsets, and using four different mass proxies to define the age residual, as described in the text. For each row, N_{gxy} gives the range in number of galaxies used in the four fits.

| Sample | N_{gxy} | Gradient in $\Delta \log T_{spec}$ using mass proxy | | | | Median gradient |
		σ	L_r	M_{stel}	M_{dyn}	
All	310–362	-0.06 ± 0.02	-0.08 ± 0.02	-0.07 ± 0.02	-0.07 ± 0.02	-0.07
All inside 1 Mpc	166–178	-0.09 ± 0.06	-0.12 ± 0.07	-0.11 ± 0.06	-0.08 ± 0.06	-0.10
All excluding SW	269–317	-0.07 ± 0.02	-0.09 ± 0.02	-0.08 ± 0.02	-0.07 ± 0.02	-0.07
All in core and SW	197–209	-0.07 ± 0.03	-0.07 ± 0.03	-0.06 ± 0.02	-0.07 ± 0.03	-0.07
Dwarfs (<25 per centile mass)	78–91	-0.15 ± 0.03	-0.13 ± 0.03	-0.13 ± 0.04	-0.13 ± 0.03	-0.13
Dwarfs inside 1 Mpc	40–49	-0.19 ± 0.13	-0.20 ± 0.11	-0.26 ± 0.14	-0.34 ± 0.11	-0.23
Dwarfs excluding SW	65–78	-0.14 ± 0.03	-0.13 ± 0.03	-0.14 ± 0.04	-0.11 ± 0.04	-0.13
Dwarfs in core and SW	51–57	-0.20 ± 0.05	-0.17 ± 0.05	-0.15 ± 0.05	-0.20 ± 0.05	-0.19
Giants (>25 per centile mass)	271–232	-0.04 ± 0.02	-0.07 ± 0.03	-0.05 ± 0.02	-0.05 ± 0.02	-0.05
Giants inside 1 Mpc	118–138	-0.03 ± 0.07	-0.08 ± 0.08	-0.06 ± 0.06	$+0.02 \pm 0.08$	-0.04
Giants excluding SW	203–243	-0.04 ± 0.02	-0.08 ± 0.02	-0.05 ± 0.02	-0.06 ± 0.02	-0.06
Giants in core and SW	140–158	-0.02 ± 0.03	-0.04 ± 0.03	-0.03 ± 0.03	-0.01 ± 0.03	-0.03

the remains of an ancient group originally surrounding NGC 4874, which like the stellar envelope of that galaxy has been partly stripped and displaced by interactions with the sub-cluster centred on NGC 4889.

3 INTERPRETATION USING ENVIRONMENTAL QUENCHING MODELS

3.1 Environmental history of cluster members in models

As discussed in the Introduction, the increasing fraction of passive galaxies with increasing galaxy density or host group mass is interpreted as evidence for environment-driven cessation or 'quenching' of star formation. The results presented in Section 2 show that for galaxies that are now on the red sequence, the time *since* they were quenched also exhibits a modulation with environment, in this case measured by cluster-centric distance. Since the events responsible for shutting down star formation must have occurred several Gyr in the past, the *present* location of a galaxy may differ dramatically from the environment in which it was quenched. An important step in interpreting the observed trends is thus to determine how the average *environmental history* of cluster galaxies relates to their current location within the cluster.

In this section, we address this question using halo merger trees from the *Millennium Simulation* (Springel et al. 2005). We focus here explicitly on quantities related to dark-matter halo assembly (i.e. accretion into progressively larger haloes, change in status from central to satellite within halo, etc.), deferring until Section 3.2 any attempt to link these processes explicitly to the star formation histories of galaxies. While similar investigations have been made by Berrier et al. (2009) and McGee et al. (2009) for the 'global' assembly of clusters, the new element in our analysis is to address comprehensively how the environmental history is correlated with halo-centric radius at $z = 0$.[1]

As analogues to the Coma cluster, we selected the five most massive $z = 0$ haloes from the simulation, which have total halo masses $M_{halo} = 2$–5×10^{15} M$_\odot$ and (one-dimensional) velocity dispersions $\sigma_v = 900$–1300 km s^{-1}. From these haloes, we extracted the histories of all $\sim 10^4$ model galaxies having stellar mass $> 10^9$ M$_\odot$ at $z = 0$, as assigned by the semi-analytic model of Font et al. (2008). The semi-analytic machinery is only used in this step, to select galaxies of comparable final masses to those in our observed sample. As noted by McGee et al., using stellar masses from the Bower et al. (2006) model would not yield significantly different results. The stellar mass limit we adopt is an order of magnitude higher than the completeness limit of the simulation merger trees.

By tracing back the main progenitor branch of each galaxy, we determined the simulation time-step at which it first joined the 'main branch' of the assembling cluster (i.e. the most massive progenitor of the cluster at that epoch). We also recorded the time-step at which the galaxy first became a member of a halo of mass exceeding thresholds 10^{12} M$_\odot$, 10^{13} M$_\odot$ and 10^{14} M$_\odot$, roughly corresponding to poor groups, rich groups and Virgo-like clusters, respectively. No distinction is made here between cases when a galaxy crosses these halo mass thresholds by 'joining' a new halo, and cases in which the host halo grows 'around' galaxies that already belong to it. Finally, we extracted the latest time-step at which a galaxy was a 'central' within its own halo, and the first time-step at which it became a 'satellite' in a separate halo.[2]

Having extracted these key epochs from the environmental history of model cluster members, we examine their correlations with projected radius from their host halo centre at $z = 0$. Fig. 5 shows that *all* the events tracked are strongly dependent on radius, with slopes in the range 0.1–0.2 dex Mpc^{-1}, in the sense that galaxies further out in the cluster experienced all the events more recently than those in the core. The strongest radial trends are seen for the time of

---

[1] The radial dependence of accretion time in clusters has been touched on previously by Gao et al. (2004), who considered only accretion of dark-matter sub-haloes into the main branch of the assembling cluster, and by Weinmann, van den Bosch & Pasquali (2011), who examined only the time since becoming a satellite.

[2] For some galaxies, these are not the same event. For example, a galaxy G1 that is central within a group may be accreted into a slightly larger group, and may spend some time as a satellite before merging with the central galaxy, G2, of the larger group. If the stellar mass of G1 is larger than that of G2 (which is quite possible given the substantial scatter in the M_{stel}/M_{halo} ratio for centrals), then after the galaxies merge G1 will become the central galaxy of the halo. At some later time, the merged halo may itself be accreted into a larger cluster, and G1 will again become a satellite.
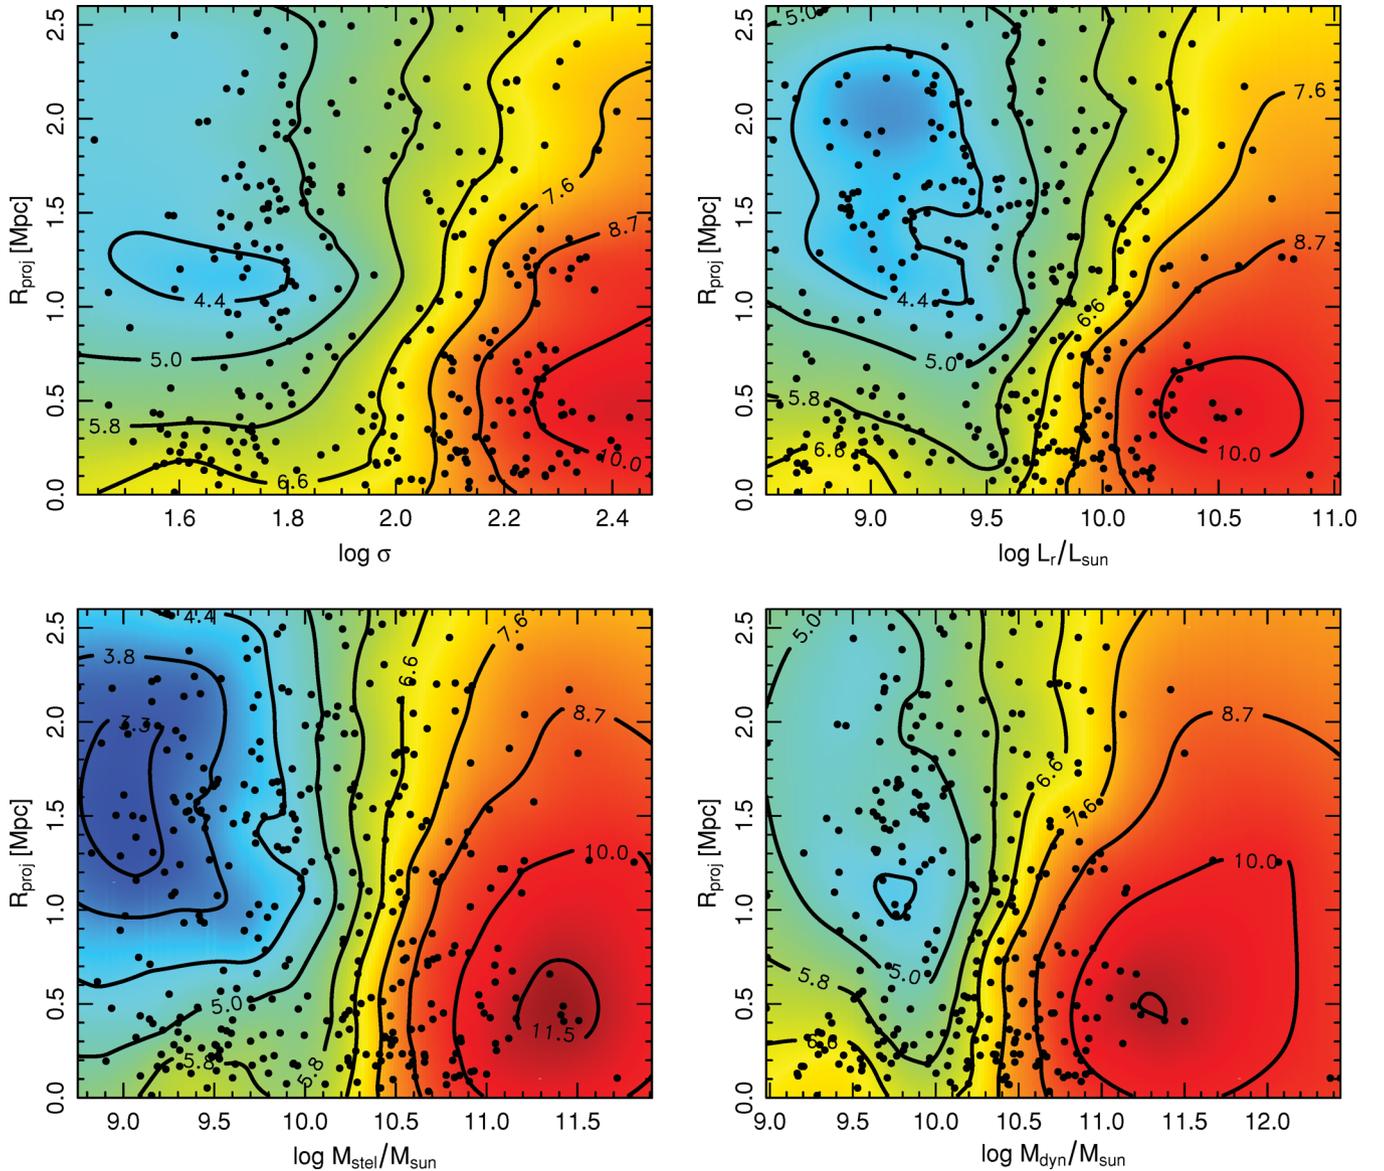
Figure 3. Contour maps of average age as a function of 'mass' and projected distance from the cluster centre, for four alternative choices of the mass-related quantity. The maps were created using a Kriging model (e.g. Cressie 1993). Black points show the location of the sample galaxies within this plane. Contours are drawn at intervals of 0.06 dex (15 per cent) in age. The key results of Section 2.2 are evident in all four panels: the ages of 'giants' are related primarily to 'mass' (with some weak environmental modulation), while the ages of 'dwarfs' depend mainly on their location in the cluster.

incorporation into the 'main branch' of the assembling halo, i.e. the final accretion event for a galaxy, and for the time of incorporation into *any* halo above $10^{14}\,M_\odot$ (upper panels in Fig. 5). (For such a high threshold halo mass, these often refer to the same event.) As an example, an average galaxy observed projected near the cluster core first joined a Virgo-sized halo 8.9 Gyr ago (i.e. at $z = 1.3$), but at a projected distance 2 Mpc, the typical galaxy encountered this cluster-like environment only 4.8 Gyr ago ($z = 0.5$). For these events, there is a large scatter among the five haloes studied, since they are sensitive to a small number of late mergers between high-mass haloes.

Perhaps more surprisingly, the central panels in Fig. 5 show that the time of accretion into *smaller* haloes, corresponding to galaxy groups, is also correlated with the final location in the cluster. For example, an average galaxy observed projected near the cluster core first joined an $M_{halo} > 10^{13}\,M_\odot$ halo 10.1 Gyr ago (i.e. at

$z = 1.9$), but a typical galaxy at 2 Mpc radius encountered this group-like environment only 6.2 Gyr ago ($z = 0.7$). This result can be elaborated by splitting the galaxies according to their halo mass immediately prior to accretion into the main branch. Some 56 per cent of the galaxies enter the main branch without passing through any previous halo of mass greater than $10^{13}\,M_\odot$ (in agreement with McGee et al. 2009), so that their accretion into the main branch and first accretion into a group-scale halo refer to the same event. The radial trend in accretion time (into groups) for these galaxies is similar to that obtained for all objects (though by construction they are displaced to more recent accretion times). For the 44 per cent of galaxies which *were* accreted through $M_{halo} > 10^{13}\,M_\odot$ haloes, the accretion time (into the group) is still correlated with location in the final cluster, although the slope of the correlation is reduced by a factor of 2. The trend of group-accretion epoch with ultimate cluster-centric radius can thus be attributed to a combination
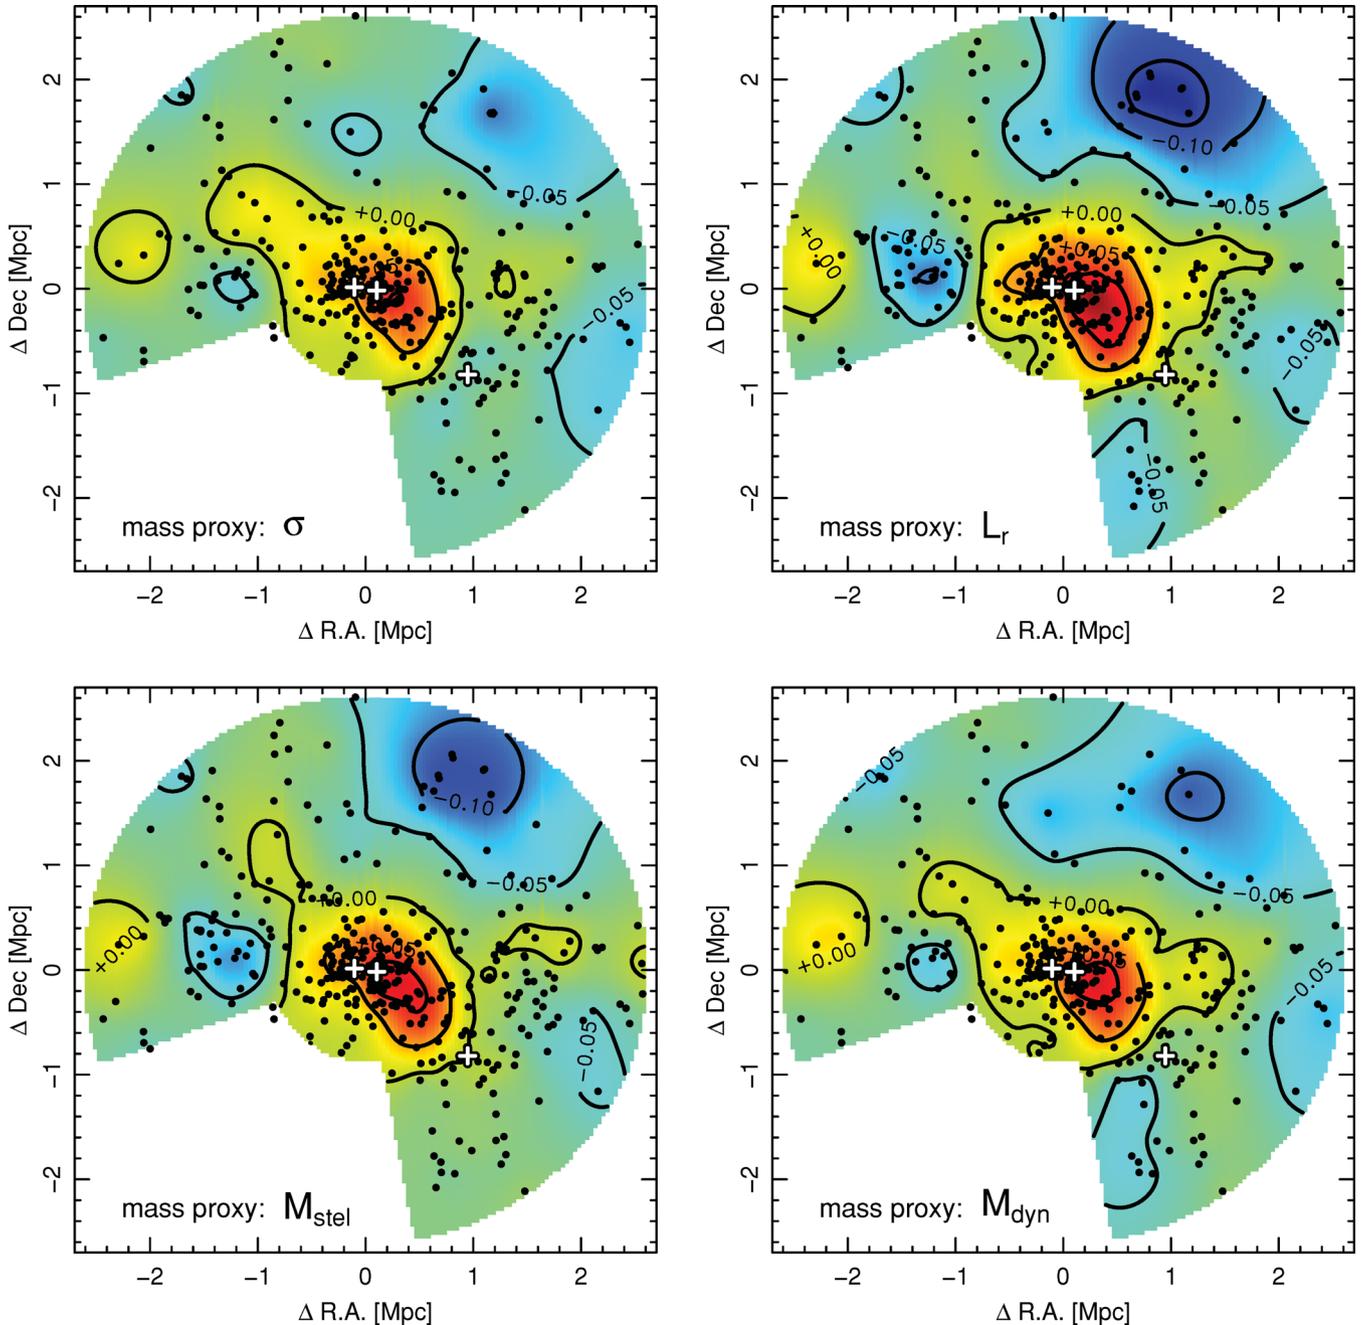
Figure 4. Residuals from the age-versus-mass relations, mapped as a function of location within the cluster. Black points show the positions of the sample galaxies, while the colour-map and contours (spaced by 0.05 dex) give the average residual from the broken-stick age–mass relations. Each panel shows the residuals computed using a different mass-proxy variable as noted in the lower left. White crosses mark the positions of the three giant galaxies NGC 4839, 4874 and 4889 (from right to left). North is up and east is to the left.

of two causes: (1) around half of the eventual members join the cluster without passing through a massive group stage, and (2) even for members which were accreted via a massive group, the time of accretion into that group was earlier among galaxies which end up in the cluster core. The latter occurs because galaxies in the cluster core were formed in an overdense region that experienced accelerated hierarchical growth, through all mass thresholds, compared to the initially more representative parts of the Universe which fall into the cluster at later times. The trend found for the lower mass threshold 10^{12} $M_{\odot}$, corresponding to small groups, is probably similarly explained.

Finally, the time at which galaxies change status, from central within their own halo to satellites within other haloes, is shown by the lower panels of Fig. 5. Note that in this case there is no threshold mass; we are simply tracking the time at which a galaxy was accreted into a halo larger than its own. Again we find a correlation in accretion epoch with cluster-centric distance: an average galaxy observed projected near the cluster core last ceased to be central within its own halo 10.6 Gyr ago ($z = 2.2$), but at a projected distance 2 Mpc, the typical galaxy last ceased to be a central only 6.6 Gyr ago ($z = 0.75$). The slope of this correlation is comparable to that shown by Weinmann et al. (2011).
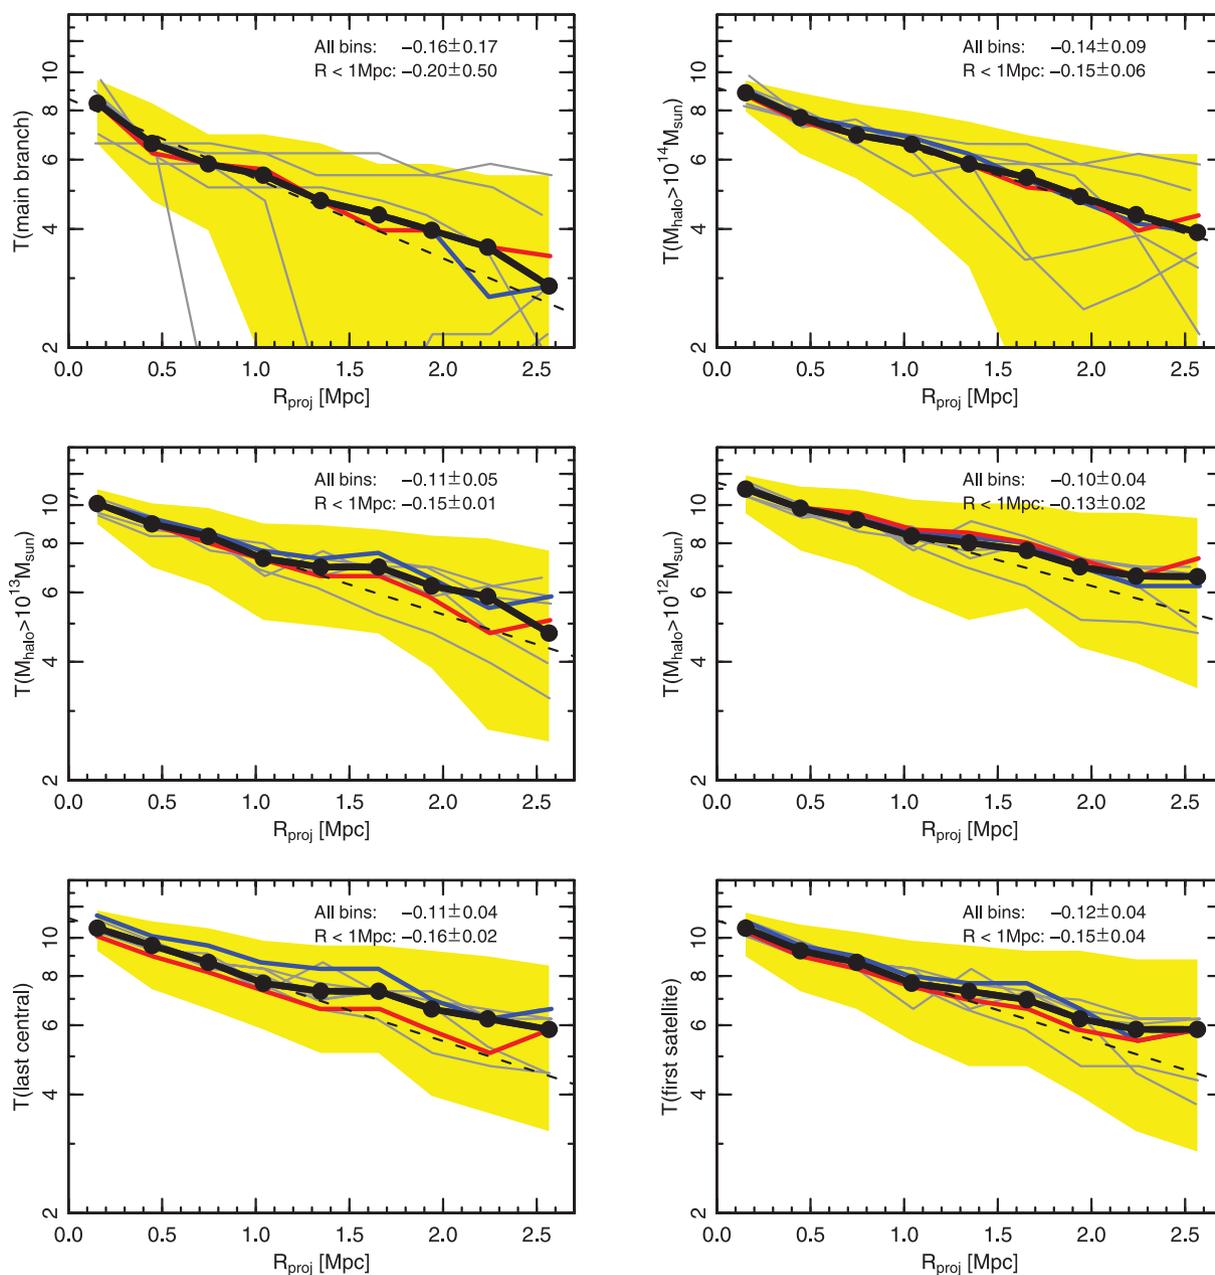
Figure 5. Environmental 'ages' of *Millennium Simulation* galaxies in $10^{15} M_\odot$ haloes, as a function of projected radius from the centre of their host halo centre at $z = 0$. Each panel shows results for a different 'event' describing the accretion of galaxies into haloes. The black points and line show the average time since the accretion event, in a set of radial bins, with the yellow region indicating the interquartile range computed over all galaxies in the same bins. The thin grey lines show the results for the individual haloes; the red and blue lines show the trends for galaxies in the upper and lower tercile of the galaxy stellar mass distribution, respectively. The legend reports the halo-centric gradient computed in two different radial ranges, in units of decades in age per Mpc in cluster-centric distance. Errors here are derived from the range of results among individual clusters. The fit to the inner bins is also shown by a dashed line.

The variation in the results from cluster to cluster is substantial, particularly at large radii, and especially for the main-branch accretion and the $10^{14} M_\odot$ halo mass threshold. However, the overall trend of more recent 'accretion' (however defined) at a larger distance is recovered in all five clusters individually. The halo-centric trends are slightly stronger within the central region of the eventual cluster (within $\sim$1 Mpc), but remain evident at larger radius. A linear fit of $\log T$ versus R_{proj} provides a better description than a fit against $\log R_{proj}$. Dividing the analysis according to the galaxy stellar mass as assigned by the Font et al. model (red and blue lines in Fig. 5), we find little mass dependence in the environmental

histories. The only exception is in the time since a galaxy was last a central galaxy, which shows an offset at all radii, such that high-mass galaxies ceased to be centrals later than low-mass galaxies at the same projected radius. This presumably arises because galaxies continue to form stars while they are centrals, so that longer duration as a central galaxy leads to higher final stellar mass.

3.2 Quenching models

In Section 3.1, we examined the accretion histories of simulated galaxies and found that the timing of the key accretion events

remains correlated with projected cluster-centric distance at $z = 0$. To connect this result to the trends measured in Section 2.2, we now construct simple models to describe the quenching of star formation in relation to the environment, and hence predict the resulting gradients of stellar population age within clusters.

Although semi-analytic models such as Bower et al. (2006) or Font et al. (2008) provide prescriptions for the star formation histories of galaxies in the *Millennium Simulation*, they do not adequately describe all of the environmental quenching processes. A particular limitation, in the present context, is that even the Font et al. model, which incorporates halo gas stripping more realistically, does not track the orbits of satellite galaxies within haloes consistently with the dark-matter simulation. Instead, accreted galaxies are assigned orbital parameters randomly, from a physically motivated parent distribution (see section 2.2 of Font et al.). Since the assigned orbit dictates the stripping of halo gas, the semi-analytic star formation histories of individual galaxies are not correctly related to their locations recorded from the simulation. It would therefore be misleading to compare the measured age gradients directly to the existing semi-analytic predictions.

To link the environmental history to the predicted spectroscopic ages, we instead impose a simplistic 'quenching' model for the star formation history. The model is defined by (1) a 'quenching *criterion*', i.e. the environmental conditions which trigger quenching and (2) a 'quenching *scenario*' i.e. the form of the star formation history, and how it relates to the time of quenching. The quenching criterion describes the 'event' in the environmental history of a galaxy at which the quenching is imposed. Following the analysis in the previous section, we consider two types of quenching criteria.

(i) *Halo mass threshold:* the galaxy is quenched when it becomes part of a halo with mass above $10^{14}\,\mathrm{M_\odot}$, $10^{13}\,\mathrm{M_\odot}$ or $10^{12}\,\mathrm{M_\odot}$ (by either accretion or halo growth). Hence there is an *absolute* halo mass scale responsible for the processes causing quenching.

(ii) *Central-to-satellite status change:* the galaxy is quenched when it ceases to be a central galaxy in its own halo, and becomes instead a satellite within another halo. In this case, quenching occurs when the galaxy falls into any larger halo, regardless of its absolute mass.

The above criteria account for five of the six environmental events shown in Fig. 5. We do not consider the sixth case, of quenching only on accretion into the main branch of the assembling halo, which would imply that the 'main' progenitor of the final cluster was in some sense different from other haloes of the same mass. This picture is clearly incorrect, for instance, in the case of late mergers between haloes of comparable mass.

For a chosen quenching criterion, we can extract the quenching time, T_Q, for each galaxy from the simulation data. The quenching *scenario* describes the form of the star formation history, and hence how T_Q maps on to the SSP-equivalent age which would be measured spectroscopically, T_{spec}. We consider three types of scenarios.

(i) *Burst scenario:* here we assume that quenching is caused by a star-burst which rapidly consumes all remaining cold gas in the galaxy. No further star formation takes place after the burst, and that the burst is sufficiently strong to dominate the 'observed' spectrum. Physically, this could be caused by tidal interaction with the cluster potential driving cold gas to the centre where a star-burst results (e.g. Byrd & Valtonen 1990). This scenario has the advantage of simplicity, since by construction the measured SSP-equivalent age is approximately equal to the quenching age, i.e. $T_{spec} \approx T_Q$.

(ii) *Stripping scenario:* in this case, we assume a constant SFR beginning at very high redshift, followed by an abrupt quenching event at which all star formation ceases immediately. This scenario could represent a rapid ram-pressure stripping of cold gas from the disc of the infalling galaxy (e.g. Gunn & Gott 1972; Quilis, Moore & Bower 2000). The spectroscopic (SSP-equivalent) age in this case is *not* equal to the quenching age, since all stars formed *before* the quenching event. However, the spectroscopic age is weighted towards the youngest stellar populations, both because they are more luminous, and through the non-linear age dependence of the Balmer lines (Serra & Trager 2007). In S09, building on the work by Allanson et al. (2009), we determined a conversion between quenching time and T_{spec} for the abruptly quenched constant SFR model. [A similar result was obtained by Trager, Faber & Dressler (2008).] Inverting the form of the relation quoted in S09 gives $\log T_{spec} = \sqrt{1.087 \log T_Q + 0.190} - 0.07$.

(iii) *Strangulation scenario:* the third possibility is similar to the stripping case in having a constant SFR prior to quenching, but now the quenching event is followed by an exponentially declining SFR, with an *e*-folding time of 1 Gyr. Physically, this is intended to mimic the strangulation picture (e.g. Balogh et al. 2000) in which the hot halo gas is removed from a satellite galaxy through interaction with the ambient intra-halo medium, while the remaining cold gas in the galaxy disc is consumed over a longer time-scale by star formation. The spectroscopic age is younger than for an abrupt quenching at the same epoch, since the post-quenching tail of young stars compensates for the pre-quenching population. In fact, S09 showed that for the particular case of a 1 Gyr decline time, the SSP-equivalent age is approximately equal to the quenching age, i.e. $T_{spec} \approx T_Q$.

For the purposes of this paper, the burst and strangulation scenarios make identical predictions, since they have the same mapping from quenching time to spectroscopic age, and we can treat them together, despite their different physical motivations. The scenarios given above are clearly simplistic, and unlikely to be consistent with other constraints on star formation histories, but serve to bracket a wide range of intermediate possibilities, e.g. abrupt quenching with superposed bursts. On physical grounds, we might expect the stripping scenario to be associated with a high halo-mass threshold criterion, since ram-pressure stripping of dense cold gas requires the high intra-cluster densities and velocities associated with massive clusters. The concentration of post-starbust spectra in the densest parts of the Coma Supercluster, noted by Gavazzi et al. (2010) and Mahajan, Haines & Raychaudhury (2010), suggests that the burst model may also be relevant in high-mass haloes. However, for completeness, we test all combinations of quenching criterion and quenching scenario in the analysis which follows.

A final ingredient required for the models is the *efficiency* of environmental quenching, f_{env}. For maximally efficient models, with $f_{env} = 1$, *all* galaxies are quenched when they meet the chosen environmental criterion. Reducing f_{env} corresponds to assuming that some fraction of galaxies are quenched through processes which are not driven by (or even correlated with) the environmental history.

Having chosen a quenching criterion and a quenching scenario, we can now predict the spectroscopic ages of all galaxies in the simulated clusters: for each galaxy we use the simulation data to identify its quenching age T_Q, i.e. the time at which it met the chosen quenching threshold. Then we convert this to a spectroscopic age using $T_{spec} = T_Q$ (in the burst or strangulation scenario) or $\log T_{spec} = \sqrt{1.087 \log T_Q + 0.190} - 0.07$ (for the stripping scenario). When $f_{env} < 1$, we model non-environmental quenching

Table 2. Age gradients predicted from quenching models applied to environmental histories of cluster galaxies in the *Millennium Simulation*. The units are decades in age per Mpc in cluster-centric distance. The predicted slopes are quoted for the maximally efficient case ($f_{env} = 1$) in which *all* galaxies are environmentally quenched. The quoted errors are half the difference between steepest and shallowest gradients obtained from the five $z = 0$ haloes, and hence reflect expected cluster-to-cluster scatter.

Quenching criterion	Burst or strangulation scenario		Stripping scenario	
	$R_{proj} < 1\,\mathrm{Mpc}$	$R_{proj} < 2.5\,\mathrm{Mpc}$	$R_{proj} < 1\,\mathrm{Mpc}$	$R_{proj} < 2.5\,\mathrm{Mpc}$
M_{halo} exceeds $10^{14}\,M_{\odot}$	-0.15 ± 0.06	-0.14 ± 0.09	-0.08 ± 0.03	-0.06 ± 0.03
M_{halo} exceeds $10^{13}\,M_{\odot}$	-0.15 ± 0.01	-0.11 ± 0.05	-0.08 ± 0.01	-0.05 ± 0.02
M_{halo} exceeds $10^{12}\,M_{\odot}$	-0.16 ± 0.02	-0.11 ± 0.04	-0.06 ± 0.01	-0.04 ± 0.02
Galaxy becomes a satellite for the first time	-0.15 ± 0.04	-0.12 ± 0.04	-0.06 ± 0.01	-0.05 ± 0.02
Galaxy last ceases to be a central	-0.16 ± 0.02	-0.11 ± 0.04	-0.07 ± 0.01	-0.05 ± 0.02

randomly reassigning ages for a fraction $1 - f_{env}$ of galaxies, drawing the new values from the distribution of all predicted T_{spec}.

As the point of contact with the observational results, we focus on the slope of the cluster-centric trends in spectroscopic age.[3] The trends are fitted with respect to projected, rather than physical, cluster-centric distance, to reproduce the observed quantities. We fit the trends within 2.5 Mpc and within 1.0 Mpc, to match the fits made to the data in Section 2.2. Because the aim is to compare to observations within a single cluster, the relevant error on the prediction is the halo-to-halo scatter in the gradient, rather than the formal error on the fit to all haloes. Specifically, we adopt half the range between the maximum and minimum slope among the five haloes.[4]

3.3 Results and comparison to observed trends

Table 2 summarizes the SSP-equivalent age gradients predicted by our quenching models, for the $f_{env} = 1$ case (i.e. all galaxies quenched by environment-related events).

The predicted trends within 2.5 Mpc are in the range 0.04–0.14 dex Mpc^{-1}, while those within the central 1 Mpc are 0.06–0.16 dex Mpc^{-1}. In general, the slopes are fairly insensitive to the choice of quenching *criterion*, which follows from the similar slopes obtained for the various environmental ages in Fig. 5. (Note that for the burst or strangulation scenario, the slope of the predicted ages is identical to the slope quoted in the figure.) For the halo mass threshold quenching criteria, slightly steeper slopes are recovered for higher threshold mass. The sense of this effect is as expected from pre-processing: for lower threshold masses, more galaxies are quenched in groups prior to accretion into the final cluster, and hence the slopes are diluted. The predictions depend much more sensitively on the quenching *scenario*, with the abrupt stripping case leading to predicted trends that are flatter by a factor of about 2. This effect is due to the different mapping from quenching time, T_Q (which is the same as in the burst or strangulation scenarios) to SSP-equivalent age, T_{spec}. In the centre of the halo, where most galaxies quenched early, this mapping has little effect, since an early

short burst looks much like an SSP (e.g. $T_Q = 10$ Gyr corresponds to $T_{spec} = 11.5$ Gyr in the stripping scenario). However, in the outskirts, where most galaxies quenched fairly recently, the difference is much greater (e.g. $T_Q = 5$ Gyr corresponds to $T_{spec} = 8$ Gyr). Hence, the stripping model tends to compress the range of SSP-equivalent age, leading to shallower predicted gradients.

The predictions in Table 2 can be directly compared to the observed gradients in Table 1. As a first-order result, we note that the range of observed slopes is almost precisely spanned by the range of predictions from the quenching models. On a general level, this confirms that the observed slopes can be reproduced by reasonable environment-driven quenching processes. In greater detail, we can tentatively attempt to discriminate between the various models based on their agreement with the observed trends. Since the effect of reducing f_{env} is to dilute the predicted gradient, and tests show that the effect is roughly linear, i.e. setting $f_{env} = 0.5$ approximately halves the cluster-centric slope, a model which predicts too *steep* a slope in Table 2 can be reconciled with the data by reducing f_{env} accordingly. Only a model which predicts a slope that is too *shallow* compared to the data can be rejected.

For the dwarf galaxies, which show strong gradients (0.13 dex Mpc^{-1} within 2.5 Mpc and 0.23 dex Mpc^{-1} within 1 Mpc), the burst or strangulation scenarios can reproduce the measured trends within the errors, so long as the efficiency is close to maximal. By contrast the slopes predicted by the models with abrupt stripping (0.04–0.06 dex Mpc^{-1} within 2.5 Mpc and 0.06–0.08 dex Mpc^{-1} within 1 Mpc) are shallower than observed, even for $f_{env} = 1$. For the $R_{proj} < 2.5$ Mpc sample, the mismatch is at the $\gtrsim 2\sigma$ level (for the inner sample, the observational errors are larger, and the discrepancy less significant). For more massive galaxies, the observed trends are much shallower (0.04–0.05 dex Mpc^{-1}), and can be matched either by the stripping scenario with $f_{env} > 0.5$, or by the burst/strangulation scenarios with $f_{env} \approx 0.3$–0.5. Within either the burst/strangulation or the stripping scenario, the predicted gradients do not depend substantially on the environmental criterion adopted. Hence, using this test alone we cannot discriminate between models in which quenching is effective in group-mass haloes, and those in which massive clusters are required.

3.4 Caveats and limitations

We highlight here a few limitations of the analysis presented above.

(i) *Biases in the semi-analytic galaxy selection:* the models are based on semi-analytic galaxies selected at $z = 0$, and hence it is implicitly assumed that these objects have environmental histories that are statistically similar to those of real cluster galaxies. From

---

[3] The absolute age calibration of the stellar population models used to extract ages from the spectra is uncertain at (at least) the 20 per cent level due to errors in the stellar effective temperature scale (Percival & Salaris 2009). Hence comparing relative shifts in the ages is preferred over the comparison of the ages themselves.

[4] For a normally distributed random variable, this quantity estimated from five samples is on average 16 per cent larger than the standard deviation and is within a factor of 2 from the standard deviation in 90 per cent of realizations.

comparison to galaxy clustering statistics, there is evidence for an excess of low-mass satellite galaxies in semi-analytic models, compared to the real universe (e.g. Li et al. 2007; Kim et al. 2009). If the 'missing' satellites in real clusters have been disrupted by tidal interactions (Kim et al. 2009; Henriques & Thomas 2010), then it is conceivable that the excess semi-analytic satellites are preferentially those accreted at early epochs, and having orbits which take them close to the cluster centre. The expected sense of this bias would be to steepen the gradients predicted by the models, by preserving too many old galaxies in the cluster core.

(ii) *Limited range of quenching criteria:* we have not explored all possible environmental criteria for quenching. In particular, it is unclear that 'membership' of a given halo, i.e. passing within its virial radius, should be sufficient to trigger quenching. It would be interesting to consider criteria defined by passage closer to the halo centre, where intra-cluster gas densities are much higher. Unfortunately, the limited number of output time-steps from the *Millennium Simulation* precludes an accurate analysis based on such an event, since the rapid passage of a galaxy through the halo core is not temporally resolved. It seems plausible that a quenching criterion based on a close pericentre passage would yield steeper slopes than our current models.

(iii) *Limited range of quenching scenarios:* likewise, the star formation histories assumed in our quenching scenarios are very restricted and indeed unrealistic. It would be desirable to use pre-quenching star formation recipes motivated by observational constraints such as the 'main sequence' of star formation (Noeske et al. 2007). A wider range of post-quenching behaviour should also be explored. In particular, we note that using a longer post-quenching e-folding time in the strangulation models would likely result in steeper predicted gradients for a given f_{env}.

(iv) *Inclusion of star-forming satellites:* the observational results are limited explicitly to non-star-forming galaxies, to avoid emission contamination of the age-sensitive Balmer lines. By contrast the simulation data implicitly include all galaxies, including those which may still be forming stars. As a crude test, we can remove all simulated galaxies within 1 Gyr of their quenching time, at $z = 0$, and re-fit the slopes. We find that the slopes computed over all bins are unchanged by excluding the youngest galaxies, while the slopes within 1 Mpc are flatter by ∼20 per cent. Hence it seems unlikely that the inclusion of star-forming galaxies in the models produces a serious bias.

A full exploration of these issues is beyond the scope of this paper, but should be addressed in future work.

4 DISCUSSION AND CONCLUSIONS

We have analysed the SSP-equivalent stellar ages of red-sequence galaxies in the Coma cluster, focusing on the systematic correlation of age with projected distance from the cluster centre.

Combining data for dwarf and giant galaxies, from the core to the virial radius, and over a wide angular extent, has helped to clarify many of the outstanding questions posed by S09, and previous observational work in this area. In particular it is now clear that strong radial trends in the ages of dwarf galaxies are true, global, cluster-centric effects. The merger of the NGC 4839 group in the south-west is not the sole cause of the trends seen in previous work, although there is still a suggestion that the trends are somewhat stronger in this direction than elsewhere, at least for the dwarfs. This is qualitatively consistent with the work of Caldwell & Rose (1997), who extended the original Caldwell et al. (1993) study to

a field in the north-west of Coma and found a lower, but non-zero incidence of star-burst and post-star-burst galaxies as compared to the south-west. If the age–radius trend in Coma is not localized to a particular merging event, then we should expect to observe similar behaviour in other massive clusters. Few other systems have been studied so intensively, in particular in the dwarf regime, where the trends are strongest. None the less, work in Abell 496 (Chilingarian et al. 2008) and Virgo (Michielsen et al. 2008) tend to support our conclusions, albeit at lower significance due to their smaller samples and/or restricted radial coverage. Our work also resolves the discrepancy between the strong cluster-centric trends seen in S09 for the Coma dwarfs, and the weaker dependence found by Smith et al. (2006) for more massive galaxies in an ensemble of clusters. By analysing a combined sample that spans a wide range in mass within the same cluster, we have demonstrated a transition between two regimes on the red sequence: giants for which star formation histories are governed chiefly by their 'mass' (though with some residual radius dependence) and dwarfs whose evolution depends primarily on environment.

Many authors have recently addressed mass and environment contributions to galaxy quenching using large surveys such as the SDSS (e.g. Kauffmann et al. 2004; Haines et al. 2007; Peng et al. 2010), with recent advances being made through application of group/cluster catalogues to examine the dependence on host halo properties (e.g. Weinmann et al. 2006). In general, these analyses employ cruder, 'instantaneous', measures of galaxy properties (e.g. simply whether a galaxy has been quenched or not), and are limited to more luminous galaxies than we probe in our study. However, such studies are able to average over large numbers of clusters and groups over a wide mass range, and so provide a useful complement to our results from a single massive halo. Recently among such work, Wetzel, Tinker & Conroy (2011) find trends of 'quenched fraction' with mass and environment in the SDSS, which seem to be analogous to our results. They find that a greater fraction of satellite galaxies are quenched compared to central galaxies in haloes of all masses studied ($\sim 10^{12-15}\,M_\odot$), and infer that there is no particular minimum halo mass responsible for satellite quenching. They find that the quenched fraction among satellites decreases at greater halo-centric distance in haloes of all masses, but with a slope which depends on stellar mass: low-mass satellites show a much steeper dependence than high-mass satellites. (Note that their lowest bin is centred at $M_{stel} = 8 \times 10^9\,M_\odot$, substantially more massive than our 'dwarfs'.) In a separate study using the SDSS, Peng et al. (2011) highlight the apparently separable effects of 'mass quenching' and 'environment quenching'. The latter, which they argue is driven solely by quenching of satellite galaxies, is apparently dependent on halo-centric radius rather than halo mass. Although these results are qualitatively similar to our conclusions in this paper, there is a key difference in that our stellar population ages reflect, albeit imperfectly, the time *since* the red satellite galaxies were quenched. In principle, this should provide additional power to constrain the physical mechanisms responsible for driving galaxy evolution in clusters and groups.

By directly comparing our results against predictions from simulated cluster assembly histories, we have explicitly accounted for the hierarchical growth of clusters, and the associated complication of pre-processing within groups prior to the final infall. Our analysis shows how the typical environmental history experienced by galaxies is correlated with their location within the cluster at $z = 0$. In turn, this allowed us to test simple environmental quenching models in which star formation is shut down in a galaxy when it enters a halo above a given mass threshold, or becomes a satellite galaxy.

The predicted gradients are quite insensitive to what environmental 'event' is associated with quenching the star formation. In particular, pre-processing in groups does not prevent age gradients from being established in the final cluster. In models with quenching at the group halo-mass scale, the radial dependence arises in the models partly because around half of the galaxies fall into the cluster as central galaxies within small haloes, and hence do not undergo any pre-processing. However, even among those galaxies that are accreted through massive groups, their history of accretion *into* those groups still shows some correlation with present-day location: location in the cluster core at $z = 0$ effectively selects galaxies which passed through all stages of the environmental history, including pre-processing, earlier than those which today reside in the cluster outskirts.

The strength of the predicted gradients does, however, depend on how the quenching event is reflected in the star formation history. Models with very abrupt quenching lead to shallower age trends than those in which quenching is followed by a more gradual decline in the SFR. Abrupt stripping models are marginally ruled out for the dwarfs, but can match the shallower gradients followed by the more massive galaxies, if they operate at nearly maximum efficiency (i.e. if *all* galaxies are quenched this way, without any contribution from non-environmental causes). The strong trends seen for the dwarfs can be reproduced by models with a dominant burst, or an exponentially declining SFR after quenching, though again a very high environmental quenching efficiency is required. These models can also match the shallower trends in the more massive galaxies, if the environmental efficiency is set lower, i.e. if many of the giants were quenched by processes uncorrelated with halo assembly (cf. the 'mass quenching' of Peng et al. 2010, 2011). This result may be related to the fact that more massive galaxies on average have a higher fraction of their mass in bulges. Hudson et al. (2010) showed that the colours of galaxy discs become bluer at greater cluster-centric distance, but the colours of bulges are uncorrelated with R_{proj}. A disc-specific quenching mechanism, as suggested by this result, would lead to lower f_{env} for bulge-dominated giants than for the on average more discy dwarfs.

In summary, our main conclusions are as follows.

(i) Significant spatial trends in the SSP-equivalent ages of red-sequence galaxies are observed in the Coma cluster, such that galaxies projected close to the cluster core are older on average than those *of similar mass* located further from the cluster centre.

(ii) The age trend appears to be a true global cluster-centric gradient. It is not driven by an excess of young galaxies in the south-west region, associated with the merging group centred on NGC 4839.

(iii) The observed radial age trends are stronger for dwarfs than for more massive galaxies, and strongest in the central parts of the cluster.

(iv) Intriguingly, the oldest average ages (after controlling for mass) are concentrated ∼200 kpc west-south-west of NGC 4974, in an area where significant diffuse stellar light has been detected.

(v) Analysing the assembly history of massive haloes in the *Millennium Simulation* shows that galaxies located at larger distance from the cluster centre not only entered the final halo later than those near the cluster core, but also passed through all key environmental thresholds at later epochs.

(vi) Simple environmental quenching models, overlaid on the accretion histories, can quantitatively reproduce the range of observed cluster-centric age trends.

(vii) The agreement between data and models is fairly generic and does not favour a particular halo mass threshold as responsi-

ble for stripping, nor necessarily favour any halo mass threshold model rather than a model in which quenching follows upon simply becoming a satellite in a halo of *any* mass.

(viii) The observed age trends for dwarfs can be reproduced by models with quenching via strangulation, but only if this acts with nearly maximal efficiency, i.e. almost all dwarfs were quenched by events tightly linked to their host-halo environment.

(ix) The weaker radial trends seen for more massive galaxies probably imply that internal processes, uncorrelated with environment, dilute the cluster-centric gradients for giants.

In this paper, we have limited our modelling analysis to interpret the cluster-centric age trends in Coma, but a more comprehensive treatment is clearly desirable. In addition to resolving some of the caveats noted in Section 3.4, future work should address the dependence of halo-centric age trends on halo mass and redshift. The predictions of such models could be extended to other observables such as the passive fraction and the post-star-burst galaxy fraction. Used in combination, these constraints should provide greater power to discriminate between candidate environmental quenching schemes.

REFERENCES

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

Allanson S. P., Hudson M. J., Smith R. J., Lucey J. R., 2009, ApJ, 702, 1275

Balogh M. L., Navarro J. F., Morris S. L., 2000, ApJ, 540, 113

Berrier J. C., Stewart K. R., Bullock J. S., Purcell C. W., Barton E. J., Wechsler R. H., 2009, ApJ, 690, 1292

Boselli A., Gavazzi G., 2006, PASP, 118, 517

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

Byrd G., Valtonen M., 1990, ApJ, 350, 89

Caldwell N., Rose J. A., 1997, AJ, 113, 492

Caldwell N., Rose J. A., Sharples R. M., Ellis R. S., Bower R. G., 1993, AJ, 106, 473

Caldwell N., Rose J. A., Concannon K. D., 2003, AJ, 125, 2891

Carter D. et al., 2002, ApJ, 567, 772

Chilingarian I. V., Cayatte V., Durret F., Adami C., Balkowski C., Chemin L., Laganá T. F., Prugniel P., 2008, A&A, 486, 85

Chung A., van Gorkom J. H., Kenney J. D. P., Vollmer B., 2007, ApJ, 659, L115

Cody A. M., Carter D., Bridges T. J., Mobasher B., Poggianti B. M., 2009, MNRAS, 396, 1647

Cressie N. A. C., 1993, Statistics for Spatial Data. Wiley, New York

Dressler A., 1980, ApJS, 42, 565

Fabricant D. et al., 2005, PASP, 117, 1411

Font A. S. et al., 2008, MNRAS, 389, 1619

Gao L., White S. D. M., Jenkins A., Stoehr F., Springel V., 2004, MNRAS, 355, 819

Gavazzi G., Fumagalli M., Cucciati O., Boselli A., 2010, A&A, 517, A73

Gregg M. D., West M. J., 1998, Nat, 396, 549

Gunn J. E., Gott J. R., 1972, ApJ, 176, 1

Guzmán R., Lucey J. R., Carter D., Terlevich R. J., 1992, MNRAS, 257, 187

Haines C. P., Gargiulo A., La Barbera F., Mercurio A., Merluzzi P., Busarello G., 2007, MNRAS, 381, 7

Henriques B. M. B., Thomas P. A., 2010, MNRAS, 403, 768

Hudson M. J., Stevenson J. B., Smith R. J., Wegner G. A., Lucey J. R., Simard L., 2010, MNRAS, 409, 405

Jørgensen I., 1999, MNRAS, 306, 607

Kauffmann G., White S. D. M., Heckman T. M., Ménard B., Brinchmann J., Charlot S., Tremonti C., Brinkmann J., 2004, MNRAS, 353, 713

Kim H.-S., Baugh C. M., Cole S., Frenk C. S., Benson A. J., 2009, MNRAS, 400, 1527

Kroupa P., 2001, MNRAS, 322, 231

Larson R. B., Tinsley B. M., Caldwell C. N., 1980, ApJ, 237, 692

Li C., Jing Y. P., Kauffmann G., Börner G., Kang X., Wang L., 2007, MNRAS, 376, 984

Łokas E. L., Mamon G. A., 2003, MNRAS, 343, 401

McGee S., Balogh M. L., Bower R. G., Font A. S., McCarthy I. G., 2009, MNRAS, 400, 937

Mahajan S., Haines C. P., Raychaudhury S., 2010, MNRAS, 412, 1098

Maraston C., 2005, MNRAS, 362, 799

Matković A., Guzmán R., 2005, MNRAS, 362, 289

Michielsen D. et al., 2008, MNRAS, 385, 1374

Moore S. A. W., Lucey J. R., Kuntschner H., Colless M. M., 2002, MNRAS, 336, 382

Nelan J. E., Smith R. J., Hudson M. J., Wegner G. A., Lucey J. R., Moore S. A. W., Quinney S. J., Suntzeff N. B., 2005, ApJ, 632, 137

Noeske K. G. et al., 2007, ApJ, 660, L43

Peng Y. et al., 2010, ApJ, 721, 193

Peng Y., Lilly S. J., Renzini A., Carollo M., 2011, ApJ, submitted (arXiv:1106.2546)

Percival S. M., Salaris M., 2009, ApJ, 703, 1123

Price J., Phillipps S., Huxor A., Smith R. J., Lucey J. R., 2011, MNRAS, 411, 2558

Quilis V., Moore B., Bower R., 2000, Sci, 288, 1617

Schiavon R. P., 2007, ApJS, 171, 146

Serra P., Trager S. C., 2007, MNRAS, 374, 769

Smith R. J. et al., 2004, AJ, 128, 1558

Smith R. J., Hudson M. J., Lucey J. R., Nelan J. E., Wegner G. A., 2006, MNRAS, 369, 1419

Smith R. J. et al., 2008, MNRAS, 386, L96

Smith R. J., Lucey J. R., Hudson M. J., Allanson S. P., Bridges T. J., Hornschemeier A. E., Marzke R. O., Miller N. A., 2009a, MNRAS, 392, 1265 (S09)

Smith R. J., Lucey J. R., Hudson M. J., 2009b, MNRAS, 400, 1690

Smith R. J. et al., 2010, MNRAS, 408, 1417

Springel V. et al., 2005, Nat, 435, 629

Sun M., Donahue M., Roediger E., Nulsen P. E. J., Voit G. M., Sarazin C., Forman W., Jones C., 2010, ApJ, 708, 946

Thomas D., Maraston C., Bender R., Mendes de Oliveira C., 2005, ApJ, 621, 673

Trager S. C., Faber S. M., Dressler A., 2008, MNRAS, 386, 715

Weinmann S. M., van den Bosch F. C., Yang X., Mo H. J., 2006, MNRAS, 366, 2

Weinmann S. M., van den Bosch F. C., Pasquali A., 2011, preprint (arXiv:1101.3244)

Welch G. A., Sastry G. N., 1971, ApJ, 169, L3

Wetzel A. R., Tinker J. L., Conroy C., 2011, preprint (arXiv:1107.5311)

Zabludoff A. I., Mulchaey J. S., 1998, ApJ, 496, 39

This paper has been typeset from a TEX/LATEX file prepared by the author.