

Environmental vs. intrinsic quenching at cosmic noon: Predictions from cosmological hydrodynamical simulations for VLT-MOONRISE

Paul H. Goubert,^{1*}, Asa F. L. Bluck,^{1†}, Joanna M. Piotrowska², Paul Torrey^{3,4,5}, Roberto Maiolino^{6,7,8}, Thomas Pinto Franco¹, Camilo Casimiro¹ and Nicolas Cea¹

¹Stocker AstroScience Center, Dept. of Physics, Florida International University, 11200 SW 8th Street, Miami, 33199, Florida, USA

²Department of Astronomy, California Institute of Technology, 1200 East California Boulevard, Pasadena, California 91125, USA

³Department of Astronomy, University of Virginia, 530 McCormick Road, Charlottesville, VA 22904, USA

⁴Virginia Institute for Theoretical Astronomy, University of Virginia, Charlottesville, VA 22904, USA

⁵The NSF-Simons AI Institute for Cosmic Origins, USA

⁶Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

⁷Cavendish Laboratory, University of Cambridge, 19 J. J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom

⁸Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We present an investigation into the quenching of simulated galaxies across cosmic time, honing in on the role played by both intrinsic and environmental mechanisms at different epochs. In anticipation of VLT-MOONRISE, the first wide-field spectroscopic galaxy survey to target cosmic noon, this work provides clear predictions to compare to the future observations. We investigate the quenching of centrals, high-mass satellites, and low-mass satellites from two cosmological hydrodynamical simulations: IllustrisTNG and EAGLE. Satellites are split according to bespoke mass thresholds, designed to separate environmental and intrinsic quenching mechanisms. To determine the best parameter for predicting quiescence, we apply a Random Forest classification analysis for each galaxy class at each epoch. The Random Forest classification determines supermassive black hole mass as the best predictor of quiescence in centrals and high-mass satellites. Alternatively, the quenching of low-mass satellites is best predicted by group halo mass, at all epochs. Additionally, we investigate the evolution in the dependence of the quenched fraction with various parameters, revealing a more complex picture. There is strong evidence for the rejuvenation of star formation from $z = 2$ to $z = 0$ in EAGLE, but not in IllustrisTNG. The starker discrepancy between simulations rests in the mass threshold analysis. While IllustrisTNG predicts the existence of environmentally quenched satellites visible within the survey limits of MOONRISE, EAGLE does not. Hence, MOONRISE will provide critical data that is needed to evaluate current models, and constrain future models, of quenching processes.

Key words: Galaxies: evolution – Galaxies: formation – Galaxies: star formation

1 INTRODUCTION

Galaxies, both in the local Universe and at higher redshifts, demonstrate a strong bimodality in numerous parameters, including color, morphology, and star formation rate (SFR) (e.g., Strateva et al. 2001; Brinchmann et al. 2004; Williams et al. 2009; Peng et al. 2010, 2012). Indeed, when analyzing the distribution of galaxies along the star formation rate (SFR) - stellar mass (M_*) (Brinchmann et al. 2004), color - magnitude (Strateva et al. 2001; Williams et al. 2009), and color - density (Woo et al. 2013) planes, a clear segregation becomes apparent.

Accordingly, galaxies may then be split into two populations. These are colloquially referred to as the ‘blue cloud’, or main sequence, and the ‘red sequence’, or quenched population. The former

is comprised of galaxies with strong ongoing star formation, while the latter consists of galaxies whose star formation rates fall far below the main sequence at a fixed stellar mass and redshift (e.g., Peng et al. 2010, 2012). The process a galaxy undergoes to transition from the former to the latter is known as quenching. Studying this critical period in galaxies’ lifetimes plays a pivotal part in understanding galaxy formation and evolution in general.

Quenching has been shown to come in one of two main forms, intrinsic and environmental. Intrinsic, or ‘mass’, quenching is connected to properties such as stellar mass, black hole mass, central velocity dispersion, and bulge mass (Peng et al. 2010; Bluck et al. 2014; Bluck et al. 2016; Piotrowska et al. 2022; Brownson et al. 2022). Environmental quenching links quiescence to a galaxy’s surroundings, through parameters such as halo mass (Woo et al. 2013; Barsanti et al. 2018; Goubert et al. 2024) and local galaxy overdensity (van Dokkum & Quadri 2007; Peng et al. 2010, 2012; Woo et al. 2013; Contini et al. 2020; Bluck et al. 2020b). It is important to

* E-mail: pgoubert@fiu.edu

† E-mail: abluck@fiu.edu

note that certain intrinsic and environmental parameters are correlated with another, such as black hole mass with halo mass and local density. As such, for this study environmental quenching refers solely to effects due to the environment itself (e.g., as a consequence of ram pressure stripping, tidal stripping or galaxy-galaxy interactions).

Intrinsic quenching connects the quiescence of galaxies to mechanisms which are inherent to the system. Possible theoretical routes linked to intrinsic quenching include, supernova feedback (Bower et al. 2008; Schaye et al. 2015), virial shock heating (Dekel & Birnboim 2006; Woo et al. 2015; Gabor & Davé 2014), and active galactic nuclei (AGN) feedback (Croton et al. 2006; Sijacki et al. 2007; Fabian 2012; Maiolino et al. 2012; Weinberger et al. 2017). However, it has previously been shown that supernova feedback does not provide the energy necessary to fully quench more massive systems (Croton et al. 2006; Bower et al. 2006, 2008; Smith et al. 2018; Henriques et al. 2019). On the other hand, AGN feedback has been shown to be the most likely culprit in quenching massive galaxies, both in simulations and observations, both in the local Universe (Weinberger et al. 2018; Zinger et al. 2020; Bluck et al. 2016, 2020a; Piotrowska et al. 2022), and at higher redshifts (Bluck et al. 2022, 2023, 2024).

While intrinsic processes are efficient in quenching high-mass galaxies, the same can not be said for low-mass systems (Wang et al. 2018; Contini et al. 2020; Bluck et al. 2020b; Goubert et al. 2024). The presence of quiescent low-mass galaxies in the local Universe therefore indicates the need for a separate path to achieve quiescence. Causal links between environmental parameters, such as local over-density and halo mass, and the star formation rate (SFR), or color, of galaxies are well documented for the local Universe (see, e.g., Peng et al. 2010, 2012; Woo et al. 2013, 2015; Woo et al. 2017). Furthermore, quenched galaxies tend to be preferentially found in denser environments, and more massive groups, when compared to their star forming counterparts. As such, the environment of galaxies must be at least partially responsible for the quenching process, and is likely crucial at lower stellar masses. However, the exact mechanisms at play in environmental quenching remain unknown. There are numerous candidates, including ram pressure stripping (e.g., Gunn & Gott 1972; Kapferer et al. 2009; Cramer et al. 2019; Lotz et al. 2019; Rhee et al. 2024), galaxy-galaxy harassment (Moore et al. 1996, 1998; Marasco et al. 2016), tidal stripping (Moore et al. 1999; Fang et al. 2016), and other forms of starvation (see van den Bosch et al. 2008; Cortese & Hughes 2009; van de Voort et al. 2017; Xie et al. 2020; Wright et al. 2022).

The arrival of JWST (Finkelstein et al. 2023; Kocevski et al. 2023; Kartaltepe et al. 2023; Eisenstein et al. 2023a,b), has enabled the study of intrinsic quenching at intermediate and high redshifts (e.g., Bluck et al. 2024; Carnall et al. 2023, 2024). As in the local Universe, the quiescence of massive galaxies (primarily centrals) at higher redshifts has been linked to the mass of their supermassive black hole, or a well motivated proxy thereof (Bluck et al. 2023, 2024). Therefore, even in the early Universe, the probable role of intrinsic quenching operating via AGN feedback has been clearly demonstrated.

On the other hand, while there exists strong evidence for environmentally driven quenching at low redshifts (see, e.g., Bundy et al. 2006; Ji et al. 2018; Foltz et al. 2018; Goubert et al. 2024; Hamadouche et al. 2024), establishing its importance beyond the local Universe has proven challenging. The difficulty lies in the lack of wide-field spectroscopic observational surveys at intermediate and high redshifts. The dearth of spectroscopic redshift values is particularly detrimental to investigations into environment, as a 2D+ z observer-space is further reduced to a purely 2D phase space (with highly limited radial distance constraints from photometric redshifts). Wide spread group contamination is largely unavoidable, impacting

central - satellite classification, as well as the estimation of group parameters, especially halo mass.

Whilst some spectroscopic measurements do exist at these epochs (for up to ~ 1000 's of galaxies), these are highly limited in area (Lilly et al. 2007; Garilli, B. et al. 2014; Pentericci, L. et al. 2018; D'Eugenio et al. 2025). Consequently, measurements of environment at cosmic noon and above are highly incomplete, lacking constraints on the full dynamical range in the cosmic web. Fortunately, the upcoming arrival of the VLT-MOONRISE survey (Maiolino et al. 2020; Cirasuolo et al. 2020) will largely resolve this issue by probing environment across volumes approaching the local homogeneity scale. This represents the ideal scenario for investigating environmental quenching at, and around, cosmic noon.

VLT-MOONRISE is anticipated to gather data on half a million galaxies within a redshift range of $0.9 < z < 2.6$. This will allow us to study the effects of environment on galaxies, including their star formation rates, star formation histories, quiescent fractions, morphologies, and many other properties. In particular, the MOONRISE survey is predicted to probe a density range spanning over four orders of magnitude, exploring from voids to dense clusters. These observations, alongside robust estimates for halo masses and central - satellite classification (see Bluck et al. 2025, in press), will prove indispensable for analyzing the environment of galaxies and its connection to galaxy evolution. Furthermore, the spectroscopic data obtained will aid in constraining SFR, stellar mass, and even black hole mass (via calibrations, see Bluck et al. 2023, 2024), whilst also collecting valuable information on optical AGN. Consequently, VLT-MOONRISE is set to become a revolutionary survey for understanding the quenching of galaxies within the full environmental context at cosmic noon.

In the meantime, cosmological hydrodynamical simulations present a unique opportunity for developing testable predictions for what is expected to be found at cosmic noon in VLT-MOONRISE. A clear set of predictions will be invaluable for enabling observers to immediately test and constrain theoretical expectations. We choose to utilize two leading cosmological hydrodynamical simulations in this work (IllustrisTNG & EAGLE), enabling comparison both to future observations and to each other. Yet, to be effective, one must first carefully account for the observational, and theoretical, limitations in the comparisons.

Due to the resolution limits of hydrodynamical simulations, much of the physics of galaxy evolution is modeled subgrid. This includes, star formation, stellar evolution, stellar feedback, supermassive black hole formation & growth, and AGN feedback (among many other processes). The environment of galaxies is allowed to evolve according to the prescribed laws of physics (particularly those of hydrodynamics and gravitation), with a more modest impact from subgrid physics than on galaxy scales. Indeed, processes such as the hierarchical assembly of groups and clusters, the clustering of dark matter, and mechanisms such as ram pressure stripping, tidal stripping, and galaxy-galaxy harassment, occur essentially organically within simulations, albeit with limitations imposed through the resolution in volume and mass (see, Bahé et al. 2017; De Lucia et al. 2025). Hence, cosmological hydrodynamical simulations offer powerful insights into the impact of both environment and intrinsic properties on galaxy evolution.

The binary nature of quenching (i.e., that galaxies can be either star forming or quenched) leads to its analysis being ideally suited to classification. Classification aims to determine the fundamental discrepancies between objects of different groups, in order to best segregate them accurately. Therefore, we choose to employ a Random Forest classifier as our primary method for analysis. Random

Forest (RF) classification is a supervised machine learning classification algorithm, focused on identifying the key differences between objects of separate classes, so as to optimally sort them. This is achieved by determining which parameter, and threshold, is best at separating galaxies into either class. As the RF classifier excels at parsing through nuisance parameters to establish causal links between features and classes (see, Bluck et al. 2022), it is particularly well suited for analyzing quenching, due to its complex nature and the inter-correlation of parameters. Moreover, classification via deep learning (e.g., Domínguez Sánchez et al. 2018; Busca & Balland 2018; Pérez-Carrasco et al. 2019; de Diego, José A. et al. 2020) is not suitable for our purposes as it offers limited interpretability (see Bluck et al. 2022 for a discussion).

The RF method has previously successfully established causal links between the quenching of central galaxies and AGN feedback in simulations and observations of local Universe with the SDSS (see Bluck et al. 2020b; Piotrowska et al. 2022), and at higher redshifts (Bluck et al. 2022, 2023, 2024), comparing simulations to the HST-CANDELS & JWST-CEERS surveys. Furthermore, the work done in Piotrowska et al. (2022) and in Bluck et al. (2022) demonstrates the potential of Random Forest classification for identifying the true, causal connections between parameters and quenching, via reverse engineering contemporary cosmological simulations. Through RF classification analysis, the integrated AGN feedback (probed by supermassive black hole mass) is identified to be the dominant predictor for quenching of central galaxies, even when accounting for the instantaneous AGN output, stellar mass of galaxies, and halo mass of galaxy groups.

In this paper, we expand upon our previous quenching works (Bluck et al. 2022, 2023, 2024; Piotrowska et al. 2022; Brownson et al. 2022; Goubert et al. 2024) by investigating the mechanisms at play in the quenching of simulated galaxies at $z = 0, 1, 2$, from the EAGLE (Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016) and IllustrisTNG (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Springel et al. 2018; Pillepich et al. 2018b) hydrodynamical simulations. Of particular relevance to the present work, RF classification has previously been used in Goubert et al. (2024) to identify the features most predictive of quenching for central galaxies, high-mass satellites, and low-mass satellites at $z = 0$ for simulations, as well as in SDSS observations. Through this analysis we determine that, for both simulations and observations, the quenching process of centrals and high-mass satellites is remarkably similar at low redshift, directly related to intrinsic AGN feedback via black hole mass. On the other hand, low-mass satellites in the local universe are determined to quench solely via environmental mechanisms, particularly those connected to halo mass. In this work we expand our prior analysis to cosmic noon.

Furthermore, while previous work had determined a central - satellite classification to be sufficient in separating intrinsic and environmental quenching (Peng et al. 2012), more recent work has shown an additional mass cut to be necessary (Goubert et al. 2024; Hamadouche et al. 2024; Lim et al. 2025). Therefore, to best examine each quenching avenue we split galaxies into one of three groups, as motivated in Goubert et al. (2024). Explicitly, we separate galaxies into: (i) centrals (which only quench at relatively high masses in significant numbers), (ii) high-mass satellites (which present as either star forming or quenched in high numbers), and (iii) low-mass satellites (which also present as either star forming or quenched in high numbers). Applying RF classification, along with several other statistical tests, we aim to simultaneously study intrinsic and environmental quenching from the local Universe to cosmic noon. Furthermore, in anticipation of VLT-MOONRISE, we leverage the

known physics of the simulations, and our findings from each, to form specific predictions regarding quenching for upcoming observations. Through this analysis we gain insight in understanding the change in the impact of environment on quenching galaxies across cosmic time, and how it differs between simulations.

This work is structured as follows. In Sec. 2 we detail the simulations and observational data used in this work. In Sec. 3 we first outline the ‘MOONRISE-Like’ data set, which aims to apply conditions similar to those expected for VLT-MOONRISE to the simulations, and describe how we compute the nearest neighbor densities. We then give an overview of our Random Forest classification methodology, and finally present our quenched - star forming segregation criteria. In Sec. 4, we first present the stellar mass functions of different galaxy populations at each redshift from both simulations, and then define our mass thresholds which are used to segregate high and low-mass satellites. We include the results for the Random Forest classification and quenched fraction analyses. Finally, in Sec. 5 we summarize our findings. In Appendix B we include additional information regarding the parameters for our Schechter fits from Sec. 4.1, as well details of our RF classification tests, including hyper-parameter values, seen in Appendix C.

2 DATA

2.1 Hydrodynamical Simulations

In this paper we make use of two cosmological hydrodynamical simulations, IllustrisTNG and EAGLE. For each of these simulations we obtain data at three redshift snapshots, $z = 0, 1, 2$.

2.1.1 IllustrisTNG

From IllustrisTNG¹ (hereafter TNG) we use the TNG100-1 simulation suite, for multiple epochs ($z = 0, 1, 2$), contained within a simulation volume of $\sim (100 \text{ cMpc})^3$ (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Springel et al. 2018; Pillepich et al. 2018b). The simulation assumes the cosmological parameters from the Planck Collaboration (Planck Collaboration et al. 2016) assuming a spatially flat ΛCDM universe.

The TNG simulation builds upon the original Illustris simulation (Vogelsberger et al. 2014; Genel et al. 2014; Sijacki et al. 2015) and Illustris model (Vogelsberger et al. 2014; Torrey et al. 2014). TNG utilizes the moving mesh code AREPO (Springel et al. 2018) to model gravitation, hydrodynamics, and magnetic fields, and implements an updated subgrid galaxy formation model (Pillepich et al. 2018a; Weinberger et al. 2017). This allows for the modeling of galaxies and their environment, as well their evolution across cosmic time (see Pillepich et al. 2018a). Importantly, the interactions between galaxies and their group haloes, and centrals, are directly captured, although environmental effects on galaxies remain dependent on certain subgrid models, especially the feedback implementations. This is of particular interest as we aim to understand how galaxy environment impacts galaxy evolution at multiple epochs.

Due to the limited resolution of the simulation, some processes must be accounted for via subgrid models. These include the physics behind star formation as well certain feedback mechanisms, namely supernova and AGN feedback. In TNG, black holes are seeded with mass $M_{\text{BH}} = 8 \times 10^5 \text{ M}_\odot h^{-1}$, once a halo reaches the threshold mass of $M_{\text{Halo}} = 5 \times 10^{10} \text{ M}_\odot h^{-1}$. Their accretion is regulated

¹ IllustrisTNG Data Access: www.tng-project.org/

via Eddington-limited, Bondi-Hoyle-Lyttleton accretion (see, Bondi & Hoyle 1944; Hoyle & Lyttleton 1939). The AGN feedback prescription allows for two main feedback modes, a high accretion rate (quasar) mode, and a low accretion rate (radio/ kinetic) mode (Weinberger et al. 2017, 2018). The quasar mode models AGN driven outflows during periods of high accretion rate. The kinetic mode attempts to account for the effect of jets by injecting momentum kicks imparted in a random direction, into the interstellar medium (ISM). The kinetic ‘jets’ disturb the ISM, slowing star formation, and heat the circum-galactic medium (CGM), preventing cooling and accretion of gas into the system, enabling long-term quiescence. The AGN feedback mode at any given time is determined by the ratio of the black hole’s accretion rate and the Eddington accretion rate, f_{Edd} . Specifically, to switch from ‘quasar’ to ‘kinetic’ mode, the Eddington ratio of an AGN must be below the black hole mass dependent threshold χ , expressed as:

$$\chi = \min \left[0.002 \frac{M_{\text{BH}}}{10^8 M_{\odot}}, 0.1 \right]. \quad (1)$$

Therefore, the kinetic mode is usually only active for black holes with $M_{\text{BH}} \geq 10^8 M_{\odot}$ (Weinberger et al. 2017).

Additionally, the star formation rate of galaxies in IllustrisTNG is modeled following the Springel & Hernquist (2003) prescription (Pillepich et al. 2018a). For each individual gas cell within a subhalo, the gas becomes star forming once it reaches a threshold density of $n_{\text{Halo}} = 0.1 \text{ cm}^{-3}$. The gas cells are then attributed a star formation rate, dependent upon their density, so as to comply to the empirically determined Kennicutt-Schmidt relation (Schaye & Dalla Vecchia 2008), assuming a Chabrier stellar initial mass function (Chabrier 2003; Pillepich et al. 2018a). Star particles are then stochastically created in accordance with the star formation rate of the gas cell. The total SFR of a galaxy is taken to be the sum of the SFRs for all gas cells within the subhalo. It is possible for certain galaxies to have no gas that meets the criteria for star formation, and thus are assigned a SFR value of zero. Since we work in logarithmic units (where this would yield $-\infty$), for all galaxies with $\text{SFR} = 0$, we set $\log(\text{SFR}) = -5$, an arbitrary low value which is well below the star forming main sequence (and hence leads to no confusion with our star forming - quenched classification).

2.1.2 EAGLE

We also utilize EAGLE² (Evolution and Assembly of GaLaxies and their Environment; Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016), a hydrodynamical simulation run with a modified version of GADGET-3 smoothed particle hydrodynamics solver (Springel 2005). We use the EAGLE-RefL0100N1504 suite for three snapshots, at $z = 0, 1, 2$. EAGLE utilizes the Plank Collaboration I (Planck Collaboration et al. 2014) cosmological parameters and assumes a Λ CDM universe. Similar to IllustrisTNG, while the physics and interactions on large scales are simulated directly, the resolution limits necessitate the implementation of subgrid models at small scales, $\sim 1 \text{ kpc}$, and for non-gravitational and non-hydrodynamical physics (see Schaye et al. 2015, for full details).

In EAGLE, black holes are seeded with $M_{\text{BH}} = 10^5 M_{\odot} h^{-1}$ once a halo reaches $M_{\text{Halo}} = 10^{10} M_{\odot} h^{-1}$, and then accrete matter according to Eddington-limited, Bondi-Hoyle-Lyttleton accretion (see Bondi & Hoyle 1944; Hoyle & Lyttleton 1939). However, unlike

TNG, the AGN feedback prescription for EAGLE comes in only one mode. It is most similar to the quasar, or high accretion rate, mode of TNG, operating by injecting stored energy into neighboring gas cells, with the injection rate given by $\dot{E}_{\text{feedback}} = \epsilon_f \epsilon_r M_{\text{BH}} c^2$, where ϵ_f is the fraction of energy coupled to the ISM and ϵ_r is the radiative efficiency of the accretion disk.

The quasar mode leads to a strong ejective AGN feedback process, which works by blowing large amounts of gas out from the ISM, and into the CGM, of the galaxy. Each black hole has its own energy reservoir, which increases by $\dot{E}_{\text{feedback}} \delta t$ for each time step, δt . The temperature of the surrounding gas will be increased once enough energy is stored to heat the said gas by some amount, ΔT_{AGN} , set to $\Delta T_{\text{AGN}} = 10^{8.5} K$ for this specific simulation suite. Importantly, as opposed to TNG, EAGLE has no bespoke preventative AGN feedback mode. As such, EAGLE has no purpose-built mechanism in place to prevent galaxies from re-accreting previously expelled gas (for a detailed description of the black hole subgrid model we refer the reader to, Booth & Schaye 2009; Schaye et al. 2015). Such a preventative mode may not be necessary in the real Universe, and therefore the lack of its inclusion may well be correct. Nevertheless, this is a very interesting difference between the feedback - quenching connection in the two simulations, which in principle may lead to observable differences.

Similarly to the IllustrisTNG simulation, the star formation rate of the EAGLE simulation follows a subgrid model. Specifically, the star formation prescription follows the pressure law of Schaye & Dalla Vecchia (2008), and makes use of a simplified Kennicutt-Schmidt star formation law (Kennicutt 1998). Explicitly, in its pressure dependent form, this is shown to be:

$$\dot{m}_{\star} = m_g A (1 M_{\odot} \text{ pc}^{-2})^{-n} \left(\frac{\gamma}{G} f_g P \right)^{(n-1)/2}, \quad (2)$$

where m_g is the mass of the gas particle, γ is the ratio of specific heats, f_g is the mass fraction in gas, G is the gravitational constant, and P is the total pressure. The values of the free parameters A and n are set by observational measurements, with $A = 1.515 \times 10^{-4} \dot{m}_{\star} \text{ yr}^{-1} \text{ kpc}^{-2}$ and $n = 1.4$, and $\gamma = 5/3$ (Crain et al. 2015).

This allows for the subgrid model to be calibrated by parameters specified by observations. The EAGLE simulation suite employs a metallicity-dependent density threshold, $n_H^{\star}(Z)$, defined as:

$$n_H^{\star}(Z) = 0.1 \text{ cm}^{-3} \left(\frac{Z}{0.002} \right)^{-0.64} \quad (3)$$

where Z is the gas metallicity. This accounts for the increased cooling efficiency of metal-rich gas in the subgrid prescription.

2.2 Sample Selection

For all data, we make use of a stellar mass selection criteria of $M_{\star} > 10^9 M_{\odot}$ and a halo mass selection criteria of $M_{200} > 10^{11} M_{\odot}$. This is done to ensure galaxies are well above the resolution limits of the simulations.

As we desire to probe both environmental and intrinsic quenching at numerous redshifts, we first subdivide galaxies as either centrals or satellites (see Peng et al. 2010, 2012; Bluck et al. 2016; Goubert et al. 2024). A central galaxy is defined as the most massive galaxy of its group, while satellites are simply all of the remaining galaxies within each group. As previously seen, satellites are subject to quenching via both intrinsic and environmental means, dependent upon their stellar mass (Goubert et al. 2024; Lim et al. 2025). Therefore, we

² Eagle Data Access: <http://icc.dur.ac.uk/Eagle/>

Table 1. The number of galaxies which make up each galaxy class at every redshift for both TNG and EAGLE. We use the acronyms ‘LMS’ and ‘HMS’ for low and high-mass satellites respectively, while ‘Q’ refers to quenched galaxies.

Simulation	N_{Cen}	N_{LMS}	$N_{\text{LMS};\text{Q}}$	N_{HMS}	$N_{\text{HMS};\text{Q}}$
TNG; $z = 0$	11726	8152	3801	1024	775
TNG; $z = 1$	11415	5551	1264	651	330
TNG; $z = 2$	8956	2792	353	277	102
EAGLE; $z = 0$	7125	4676	2299	422	211
EAGLE; $z = 1$	7796	3105	515	740	223
EAGLE; $z = 2$	5894	1138	98	464	111

Table 2. The number of galaxies which make up each galaxy class at every redshift for the ‘MOONRISE-Like’ of both TNG and EAGLE.

Simulation	N_{Cen}	N_{LMS}	$N_{\text{LMS};\text{Q}}$	N_{HMS}	$N_{\text{HMS};\text{Q}}$
TNG; $z = 0$	5365	2013	915	647	507
TNG; $z = 1$	5074	1307	298	379	191
TNG; $z = 2$	3588	530	51	183	58
EAGLE; $z = 0$	3143	1269	578	256	132
EAGLE; $z = 1$	3158	406	45	431	134
EAGLE; $z = 2$	2050	0	0	257	58

further categorize satellite galaxies as either low or high-mass, split according to the mass thresholds determined in Section 4.1 for the simulations.

Additionally, future analyses are performed for up to three separate data sets for each parent sample. These are: (i) a complete data set (for which the only restrictions are the stellar mass selection criteria of $M_* > 10^9 M_\odot$ and the halo mass selection criteria of $M_{\text{Halo}} > 10^{11} M_\odot$); (ii) a 2D+ z complete data set which follows the same mass criteria but operates in 3D ‘observer’ space (2D coordinate space plus observation-like spectroscopic redshift), as opposed to the full 6D phase-space of the simulations; and (iii) a ‘MOONRISE-Like’ data set (described in Section 3.1). We show the number of galaxies belonging to each class at each redshift in Tables 1 & 2, for the complete and ‘MOONRISE-Like’ data sets respectively.

3 METHODS

In this section we first detail the construction of the ‘MOONRISE-Like’ data set, which aims to form samples under similar conditions to what is expected to be observable in VLT-MOONRISE. We then describe how local galaxy over-density is measured, and give an overview of the Random Forest classification method. Finally in this section, we present our methodology for determining quiescence at each redshift, for both TNG and EAGLE.

3.1 MOONRISE-Like Data

VLT-MOONRISE is anticipated to become the largest spectroscopic survey at intermediate-to-high redshifts, targeting up to half a million galaxies (Maiolino et al. 2020; Cirasuolo et al. 2020). The survey will target the environment and behavior of galaxies at (and around) cosmic noon ($z \sim 1.5 - 2$). In this paper, we leverage the IllustrisTNG and EAGLE simulations to pinpoint the important processes, and pa-

rameters, with respect to quenching across cosmic time, specifically with the goal to test this with VLT-MOONRISE in the near future.

In order to make accurate predictions for observational data it is important to have comparable simulated data sets. Therefore, we choose to form ‘MOONRISE-Like’ data sets for the simulations, applying restrictions similar to those expected for VLT-MOONRISE. These include a stellar mass selection criteria of $M_* > 10^{9.5} M_\odot$ (above which MOONRISE will be approximately mass complete); operating within a 3D observer-space (2-coordinate space, X-Z in our case, & observational-like redshift) rather than the full 6D phase space (3-coordinate and 3-velocity space) of simulations³; and randomly removing 30% of the data to mimic the anticipated 70% completeness expected for VLT-MOONRISE.

The observational-like redshift is computed as the sum of a cosmological redshift, z_{cos} , and a peculiar redshift, z_{pec} , dependent on the peculiar velocity of each galaxy. As the simulation boxes are each at a specific redshift, the cosmological redshift is computed as:

$$z_{\text{cos}} = z(D_c(z_{\text{snapshot}}) + Y_i) \quad (4)$$

where $D_c(z_{\text{snapshot}})$ is the comoving distance to the redshift of a specific snapshot ($z = 0, 1, 2$) and Y_i is the comoving distance in the Y axis, the coordinate perpendicular to the ‘2D plane’ (X-Z plane here). The redshift at this total comoving distance can now be calculated from:

$$D_c(z) = \frac{c}{H_0} \int_0^z \frac{dz'}{E(z')} \quad (5)$$

with $E(z)$ is given by:

$$E(z) = \sqrt{\Omega_M, 0(1+z)^3 + \Omega_k, 0(1+z)^2 + \Omega_\Lambda, 0}, \quad (6)$$

where Ω_M is the total matter density, Ω_k describes the curvature of the Universe, and Ω_Λ is the fraction of the critical density attributed to dark energy (or the cosmological constant Λ). Therefore, z_{cos} represents the cosmological redshift for each galaxy, accounting for both the snapshot redshift and the positioning of the galaxy within the simulation box.

A peculiar redshift is then evaluated for each galaxy, using their peculiar velocities in the Y axis, V_y . Specifically, z_{pec} is given by:

$$z_{\text{pec}} = \frac{V_y(1+z_{\text{snapshot}})}{c} \text{ for } V_y \ll c. \quad (7)$$

It is then possible to assign an observational-like redshift to every galaxy, computed as:

$$z_{\text{obs}} = z_{\text{cos}} + z_{\text{pec}}. \quad (8)$$

As is done for observational data, we apply velocity restrictions prior to estimating distance measurements. This assures objects are likely to be physically connected. Specifically we require:

$$\Delta V_{\text{los}} = \frac{c|\Delta z|}{(1+\bar{z})} \leq 1500 \text{ km/s} \quad (9)$$

³ While we choose the Y-axis as our line of sight (LoS) for all data sets, we find our results to be consistent independent of the axis upon which the LoS is projected.

where ΔV_{los} is the difference in line of sight velocities for two galaxies, Δz the difference in observer like redshift, and \bar{z} is the mean redshift of the two galaxies.

For a more detailed description of the methodology used here, we refer the reader to Bluck et al. 2025 (in press).

We present the amount of galaxies belonging to each class at each redshift in Table 2, for the ‘MOONRISE-Like’ data sets. From the information in Table 2, the impact of the imposed criteria, particularly the mass completeness limit threshold of $M_* > 10^{9.5} M_\odot$, on each population is made clear. Most notable is the depletion of low-mass satellites for EAGLE beyond $z \geq 1$. We will return to this important difference between EAGLE and TNG throughout Section 4.

3.2 Nearest Neighbor Density

In this work we define a galaxy’s local over-density as:

$$\delta_n = \log(\rho_n) - \langle \log(\rho_n) \rangle; \text{ where, } \rho_n = \left(\frac{3n}{4\pi R^3} \right) \text{ for 6D,} \quad (10)$$

and

$$\delta_n = \log(\Sigma_n) - \langle \log(\Sigma_n) \rangle; \text{ where, } \Sigma_n = \left(\frac{n}{\pi R^2} \right) \text{ for 2D+z.} \quad (11)$$

Here n denotes the nearest neighbor (third, fifth, tenth, etc...) and R is the distance to the n^{th} nearest neighbor. For all data sets, we determine the local density to the third, fifth, and tenth nearest neighbors, yielding δ_3 , δ_5 , and δ_{10} . We present the nearest neighbor density as an over-density, normalized by the geometric mean in log space. For the complete, 6D phase space, sample sets we utilize equation (10), while for the ‘2D complete’ and ‘MOONRISE-Like’ data sets we employ equation (11).

To properly account for the periodic boundary conditions of the simulation boxes, we create duplicates of the original box, and arrange one to each of its sides and corners (see Goubert et al. 2024, for more detail). For the relatively small scales of clusters (compared to the simulation box) this poses no concern of duplications via multiple box crossing on the scales of interest in this work.

3.3 Random Forest Classification

In this work we utilize a Random Forest classifier, constructed from the SCIKIT-LEARN PYTHON package (Pedregosa et al. 2011a), to determine the parameters most important for predicting quiescence. Random Forest classification is a supervised machine learning classifier that works by sorting data into types given a set of input features. A Random Forest is composed of multiple decision trees, each of which is given a subset of randomly bootstrapped data samples. In this way, Random Forest classification is highly efficient at determining whether parameters are causally linked, or simply inter-correlated nuisance parameters (see, Bluck et al. 2022).

For all trees within the Random Forest, the input data serves as a ‘root’ node, and is repeatedly split into ‘parent’ and ‘daughter’ nodes. The ‘parent’ nodes are split so as to generate increasingly more homogeneous ‘daughter’ nodes (better segregated according to class). There are multiple techniques that may be employed to yield an optimal split. For our purposes, we choose to focus on the Gini impurity. The Gini impurity measures the probability of selecting a misclassified object in a random node, given by:

$$G(n) = 1 - \sum_i^c p_i(n^2), \quad (12)$$

where the summation is performed over all classes, c , and where, $p(n)$ indicates the probability of randomly selecting class, i , at node, n . At each node of each tree, the input feature and threshold which best minimizes the Gini impurity is chosen to split the remaining data.

As opposed to some machine learning techniques, such as artificial neural networks, Random Forest classification is not a black box ML algorithm. Explicitly, this enables the extraction of the performance of each parameter in classifying the data individually (rather than in aggregate in more complex methodologies). The ‘performance’ of a parameter is then interpreted as the relative importance of a given input feature when classifying. Computing the relative importance for each input feature is done via the following equation:

$$I_R(k) = \frac{1}{N_{\text{trees}}} \sum_{\text{trees}} \left\{ \frac{\sum_{n_k} N(n_k) \Delta G(n_k)}{\sum_n N(n) \Delta G(n)} \right\}, \quad (13)$$

where $I_R(k)$ is the relative importance of input feature k (Pedregosa et al. 2011a). The numerator is a sum over all nodes that split with feature k , and the denominator is summed over all of the nodes. Here, ΔG corresponds to the improvement of the Gini impurity from parent to daughter node, weighed by the number of input items to be classified which reach the parent node. The mean relative importance of each given feature, k , is computed by averaging $I_R(k)$ over the totality of the Random Forest. In our case, this returns the importance to quenching of each input feature, controlling for all other available features simultaneously. For deeper insight on Random Forest classification, we refer the reader to Bluck et al. (2022), where a number of detailed tests on mock and simulation data are performed in their Appendix B.

For our purposes, galaxies are taken to be in one of two classes, either quenched or star forming. The input features for each RF test are some combination of the following parameters: black hole mass (M_{BH}), halo mass (M_{Halo}), stellar mass (M_*), nearest neighbor density (δ), the distance to the central galaxy (D_{cen}), and a random number in the range 0 - 1 (Rand.). The relative performance of δ is computed as the sum of the relative importances of δ_3 , δ_5 , and δ_{10} . For all future Random Forest classification results, the relative importance presented represents the mean relative importances of a given parameter averaged over 100 Random Forest classification runs. The uncertainty in relative importance is given as the standard deviation across the 100 training and testing runs (as in Goubert et al. 2024).

3.4 Determining Quiescence

Since the SFR - M_* relation evolves with redshift, so must the threshold at which a galaxy is to be deemed quenched (e.g., Madau & Dickinson 2014; Ilbert, O. et al. 2015; Girelli, Giacomo et al. 2019; Bluck et al. 2024). As such, it is important to establish a method to determine this threshold in a consistent manner throughout varying epochs and simulations. In this work, we accomplish this by leveraging the characteristics of the sSFR ($\equiv \text{SFR}/M_*$) distributions at each relevant redshift of the TNG and EAGLE simulations. We present the sSFR distributions in Figs. 1a & 1b, for TNG and EAGLE, respectively.

A Gaussian curve is fit around the peak of the star forming region of each distribution, from which the mean and standard deviation are extracted. We exclude obviously quenched galaxies prior to fitting (explicitly, those at $\log(\text{sSFR}/\text{yr}^{-1}) < -11.5$, which is significantly offset from the star forming peak at all epochs in both simulations).

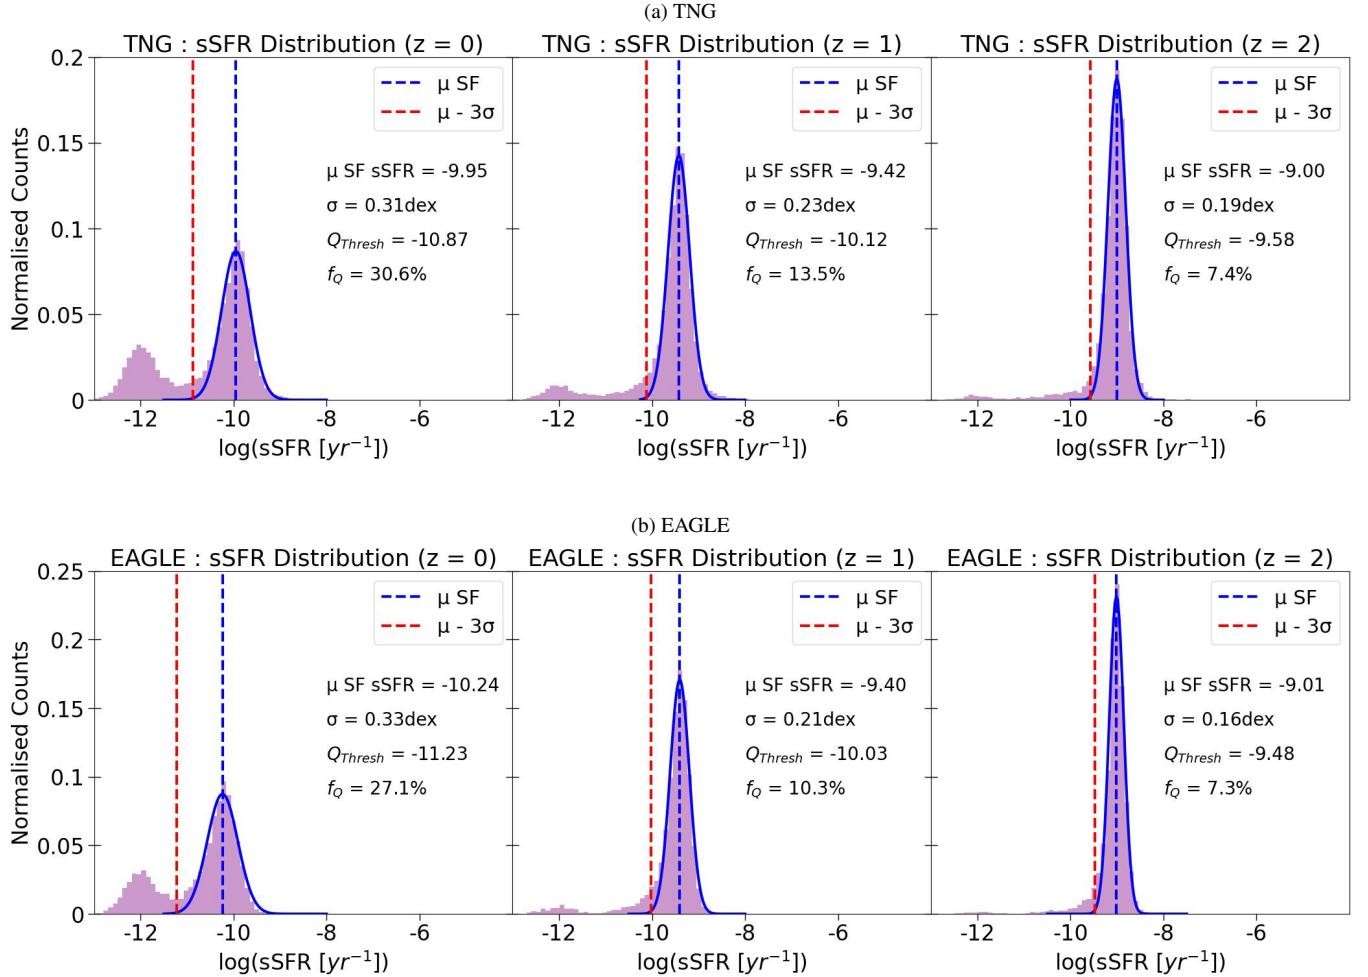


Figure 1. Illustration of our method to determine the quenched threshold for galaxies from TNG (top panels) and EAGLE (bottom panels) at $z = 0, 1, 2$ (along each row). The panels show the sSFR distribution at each redshift for each simulation. Additionally, a Gaussian fit to the star forming population is overlaid (removing obviously quiescent systems at $\log(\text{sSFR}/\text{yr}^{-1}) < -11.5$). On each panel we present the mean sSFR of star forming systems (indicated by a blue dashed line) and the determined quenching threshold (shown as a dashed red line), see equation (14). Note that we redistribute very low sSFR values to a normal distribution centered around $\log(\text{sSFR}/\text{yr}^{-1}) = -12$, in order to represent all galaxies (which may extend to $\log(\text{sSFR}/\text{yr}^{-1}) = -\infty$) on each panel. The mean sSFR increases with redshift and the width of the distribution decreases with redshift in both simulations. The fraction of quenched galaxies expected from this methodology at each epoch is displayed on the panels. In both simulations, there is a clear reduction in the quenched fraction as redshift increases, from $f_Q \sim 30\%$ at $z = 0$, to $f_Q \sim 7\%$ at $z = 2$. Clearly, the majority of quenching takes place between $z = 2$ and $z = 0$.

Given that the large majority of a Gaussian distribution is contained within 3σ of the mean, few star forming galaxies are expected to lie beyond this threshold. Therefore, we choose a robust quenching threshold of:

$$\text{sSFR}_{QT}(z) \equiv \mu_{SF}(z) - 3\sigma_{SF}(z), \quad (14)$$

where $\mu_{SF}(z)$ is the mean of the star forming Gaussian distribution, and $\sigma_{SF}(z)$ its standard deviation, at a given redshift z . This results in a bespoke threshold at each redshift for each simulation. Explicitly, galaxies are defined as quenched if:

$$\text{sSFR}_i(z_i) \leq \text{sSFR}_{QT}(z_i). \quad (15)$$

If the above criterion is not met, we define galaxies to be star forming.

We illustrate the process of determining the quenched threshold in Fig. 1. On each panel we display the mean sSFR of the star forming population, the standard deviation of the Gaussian fit, and the asso-

ciated quenching thresholds at $z = 0 - 2$ (across each row) for TNG (upper panels) and EAGLE (lower panels). In both simulations, the mean sSFR (and hence quenching threshold) increases with redshift, while the star forming Gaussian distributions become progressively tighter in log-space. The location of the quenched threshold is displayed as a red dashed line on each panel, with the location of the mean of the star forming distribution shown as a dashed blue line on each panel. Galaxies with sSFR below the quenched threshold form our quenched population at each epoch, with galaxies above this threshold forming our star forming population at each epoch.

It is important to note that the estimated quenching thresholds are likely to lead to some contamination of the quenched sample with ‘Green Valley’ (GV) galaxies (those residing between the star forming and quenched peaks). However, these galaxies must still exhibit properties of the quiescent sample as they transition from the Main Sequence. Therefore, the ‘GV contamination’ of the quenched sample may lend additional insight to the early stages of quenching,

particularly at higher redshift, where there is little-to-no established quiescent population.

From Fig. 1, it is clear that the fraction of quenched galaxies, indicated within each panel as f_Q , decreases as redshift increases, in both simulations. Indeed, at $z = 0$ we find $f_Q \sim 30\%$, while at $z = 2$ this is reduced to $<10\%$, for both TNG and EAGLE. As such, a large portion of quenching occurs within this redshift range ($z = 2 - 0$).

4 RESULTS

4.1 Identifying Quenching Mass Thresholds

In Fig. 2, we present the stellar mass functions for different galaxy populations, for both simulations. Explicitly, we show: all galaxies (in black), quenched galaxies (in red), star forming galaxies (in blue), quenched satellites (in purple), and quenched centrals (in orange). The top and bottom rows correspond to TNG and EAGLE simulations (respectively), while the panels are organized by redshift, from $z = 0 - 2$ (along each row). A Schechter (Schechter 1976), or double Schechter, functional fit is applied to each stellar mass function, with the following functional forms:

Single:

$$\phi(M)dM = \ln(10)\phi^* \left(10^{(M-M^*)(\alpha+1)} \right) \exp \left\{ -10^{M-M^*} \right\} dM \quad (16)$$

Double:

$$\begin{aligned} \phi(M)dM = \ln(10) & \left[\phi_1^* \left(10^{(M-M^*)(\alpha_1+1)} \right) + \phi_2^* \left(10^{(M-M^*)(\alpha_2+1)} \right) \right] \\ & \times \exp \left\{ -10^{M-M^*} \right\} dM \end{aligned} \quad (17)$$

where, $M \equiv \log(M_*/M_\odot)$, ϕ^* is the normalization constant, α is the slope of the mass function, and M^* is the characteristic mass, associated with the ‘knee’ of the stellar mass function. The values of each parameter for each fit are presented in Appendix B, in Tables B1 & B2, for TNG and EAGLE respectively. Additionally, included within each panel is the percentage of the population quenched via environmental means, Quench_E , and via intrinsic means, Quench_I , demonstrating their respective importance to quiescence at each redshift. The values of Quench_E and Quench_I are determined by evaluating the integral of the quenched Schechter function for the low and high-mass regimes (respectively), normalized by the total area under the quenched stellar mass function.

Quenching in the low-mass regime is expected to be dominated by environmental processes (e.g., Contini et al. 2020; Mao et al. 2022; Goubert et al. 2024; Hamadouche et al. 2024), while quenching in the high-mass regime is linked to intrinsic mechanisms (e.g., Bluck et al. 2022, 2023, 2024; Piotrowska et al. 2022; Goubert et al. 2024). However, the stellar mass at which this turnover is expected, and its evolution with redshift, remains unclear.

Considering the stellar mass functions of quenched galaxies for both simulations reveals a bimodality at all redshifts, though this is quite weak for EAGLE at $z = 0$. Furthermore, while the quenched satellite stellar mass function retains this bimodality, it is largely absent for quenched centrals, whereby only higher mass centrals are able to quench in significant numbers. This strongly implies that

there must be a quenching mechanism available only to low-mass satellites, unrelated to intrinsic processes. High-mass satellites, on the other hand, are likely subject to quenching similar to centrals, due to the similarity of their mass functions (see also Bluck et al. 2019, 2020b; Goubert et al. 2024; Hamadouche et al. 2024; Lim et al. 2025). Therefore, there is clear evidence for the presence of separate quenching avenues at low and high masses, which cannot be isolated by a simple central - satellite split.

Therefore, to best segregate these two quenching routes we develop a method to determine where to bisect the low and high mass galaxy regimes. This is done by leveraging the stellar mass functions of the quenched satellite population at each epoch (although similar thresholds can be obtained analyzing the full quenched population stellar mass functions). First, we ascertain the position of lowest gradient along the Schechter fit of the quenched satellite stellar mass function, which identifies the local minimum between the bimodal peaks. We present its location as a dotted black line traversing each panel in Figs. 2a & 2b. While this offers a possible stellar mass at which to separate the populations, one would expect environmental and intrinsic mechanisms to share a role in quenching galaxies near this limit. To remedy this potential issue, we choose to introduce a buffer zone of ± 0.2 dex about this point, indicated by the area shaded in green on each panel, to reduce contamination of environmental quenching in the intrinsic sample, and vice versa. We are then left with bespoke mass thresholds for TNG and EAGLE at all redshifts studied in this work, which are presented in Table 3.

However, as shown in Kukstas et al. (2022), there is evidence of over-quenching of low-mass satellites in simulations when compared to observational data, particularly in clusters. Importantly, this over-quenching could result in a non-physical (seen solely in simulations) upturn of the quenched (total and satellite) stellar mass function at low masses in simulations. It follows then that the mass threshold methodology implemented in this study to separate quenching avenues would not necessarily map onto observational data, as the quenched stellar mass functions may not present with such strong bimodality.

Nonetheless, the low mass upturn in the quenched stellar mass function is real in simulations, and its presence, or absence, in VLT-MOONRISE will provide insight on the validity (or otherwise) of their quenching prescriptions at cosmic noon. Furthermore, a similar method leveraging the bimodality of the quenched stellar mass function has been used in Hamadouche et al. (2024) and was shown to be efficient in isolating environmental and intrinsic quenching in observations, up to $z \sim 2$. It should be noted that in Hamadouche et al. (2024) galaxies are split at $M_* = 10^{10} M_\odot$ at all redshifts. Therefore, whether the upturn of the quiescent stellar mass function at low stellar masses is present in both simulations and observations, or solely in simulations due to a low-mass satellite over-quenching, is yet to be fully determined. In fact, as VLT-MOONRISE will be the first wide field spectroscopic survey at intermediate and high redshifts, it will provide the ideal testing data to resolve this ambiguity.

When comparing the evolution of the TNG and EAGLE mass thresholds with redshift, shown in Table 3, a distinct difference is apparent. The TNG thresholds evolve very little from $z = 0 - 2$ (decreasing by only 0.1 dex in total). On the other hand, the mass thresholds determined for the EAGLE simulation decrease significantly with redshift, first by 0.45 dex from $z = 0 - 1$, then by a further 0.3 dex from $z = 1 - 2$. This may indicate a weakening impact on quenching carried by environment, as well as the presence of an intrinsic quenching mechanism more effective at higher redshift (see later for further discussion). Crucially, these predicted differences in the stellar mass functions of quiescent objects can be tested directly

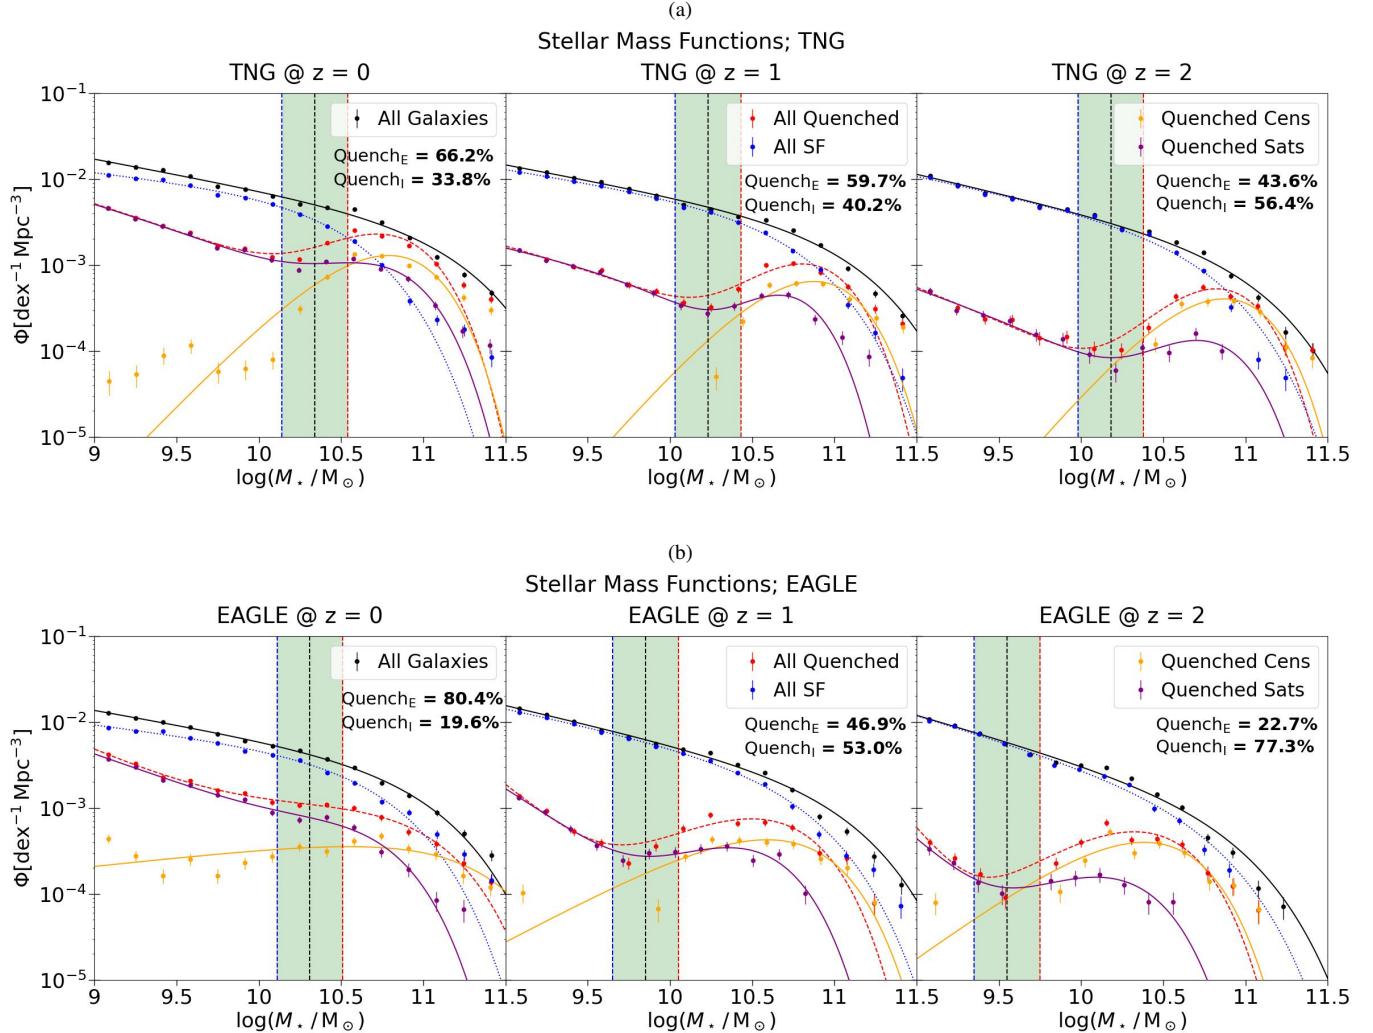


Figure 2. The stellar mass functions of various galaxy populations for the TNG (Fig. 2a) and EAGLE (Fig. 2b) simulations. Explicitly, we show stellar mass functions for: (i) all galaxies (in black); (ii) quenched galaxies (in red); (iii) star forming galaxies (in blue); (iv) quenched satellites (in purple); and (v) quenched centrals (in orange). The rows correspond to a given simulation (TNG in the top row and EAGLE in the bottom row), while the columns display redshift snapshots ($z = 0, 1, 2$, from left to right). In most cases (all redshifts for both simulations except for EAGLE at $z = 0$), the stellar mass functions of the full quenched and quenched satellite populations are clearly bimodal. We leverage this bimodality, specifically of the quenched satellite population, for each simulation to determine a ‘high-mass/low-mass’ separation in the quenched population of satellites. This is achieved by locating the local minimum of the gradient of its Schechter fit (displayed as a vertical dotted black line on each panel). We introduce a buffer zone of ± 0.2 dex between the low and high-mass thresholds (displayed as a green shaded region in each panel). This enables a separation of the quenching population on the basis of stellar mass. Note that while the quenched satellite population is bimodal, similarly to the full quenched population, the quenched central population is largely uni-modal. This indicates that centrals quench predominantly at high-masses, yet satellites may quench throughout the full stellar mass range probed. The corresponding stellar mass thresholds for each simulation and redshift are provided in Table 3.

in VLT-MOONRISE. This will offer a powerful test on which, if either, of the simulations is more effective at describing quenching in nature.

Furthermore, the sharp decrease in the low-mass threshold in EAGLE directly impacts the ‘MOONRISE-Like’ data sets for the EAGLE simulation. In fact, the low-mass thresholds at $z = 1 \& 2$ (i.e., $\log(M_*/M_\odot) \leq 9.65$ and $\log(M_*/M_\odot) \leq 9.35$, respectively) lie near, or even below, the MOONRISE survey mass completeness limit. Therefore, the ‘MOONRISE-Like’ sample sets at $z \geq 1$ for EAGLE are essentially devoid of ‘low-mass’ satellites, i.e., satellites which are expected to quench solely via environmental means. Conversely, since the mass thresholds for TNG shift only slightly with redshift, they do not have the same impact on the ‘MOONRISE-Like’ data

sets. This leads to important testable differences between the two simulations: TNG predicts that MOONRISE will observe low-mass satellite quenching via environmental processes, but EAGLE does not.

4.2 Evolution of Environment

To explore the evolution of environment with cosmic time, we present the distributions of the tenth nearest neighbor over-density ($\log(\delta_{10})$) at $z = 0, 1, 2$, for TNG (Fig. 3) and EAGLE (Fig. 4). We choose to show δ_{10} , as it presents with the strongest bimodality of all nearest neighbor densities, although similar trends are observed for all, as can be seen in Appendix A. We show these results separately for the complete and

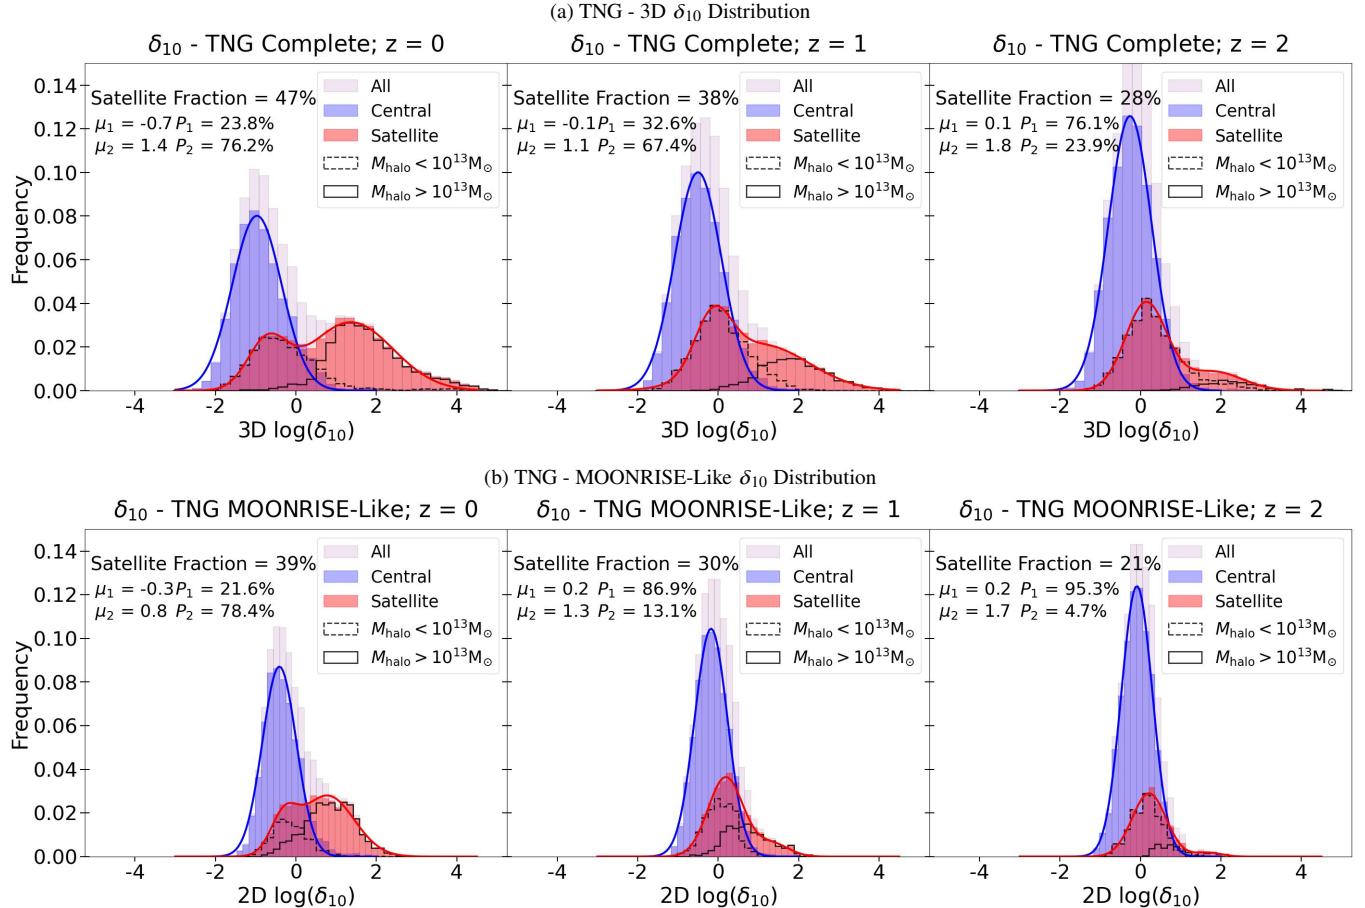


Figure 3. The distributions of $\log(\delta_{10})$ for the entire sample (shown in light grey), centrals (shown in blue), and satellites (shown in red) for the complete TNG data set (top panels) and the ‘MOONRISE-Like’ data set (bottom panels). We fit a single Gaussian to the uni-modal central distribution, and a double Gaussian to the bimodal satellite distribution. The mean of each Gaussian (μ_1 & μ_2) of the satellites distributions are displayed on each panel. Additionally, we present the fraction of satellites belonging to both the low and high density regimes of the bimodal distribution (P_1 & P_2 , respectively). The distribution in δ_{10} of satellites in high and low-mass groups ($M_{\text{Halo}} \geq 10^{13} M_{\odot}$ and $M_{\text{Halo}} \leq 10^{13} M_{\odot}$) are included as solid and dashed open histograms respectively. There is a rising fraction of satellites found in the high density regime, and corresponding decreasing fraction found in the low density peak, as a function of cosmic time. This, in combination with a growing satellite fraction (indicated within each panel), indicates a clear growth in environment as large groups and galaxy clusters form. Furthermore, it is clear that splitting satellites into high and low-mass groups is effective in recovering the bimodality in density present for satellites. As such, local over-density is a strong tracer of halo mass, or at least efficient in segregating satellites into high and low-mass groups. Whilst the results are qualitatively similar in the ‘MOONRISE-Like’ data sets, the range of δ_{10} is much more curtailed and shows much less prominent bimodality, than for the complete data set. The group mass split is also much less effective at reproducing the density bimodality. This highlights a significant loss of information between 6D and 2D+z density measurements.

Table 3. The high and low-mass thresholds determined via the stellar mass functions of the quenched satellite population for TNG and EAGLE.

Simulation	bisecting-mass; $\log(M_*/M_{\odot})$
TNG; $z = 0$	10.30 ± 0.2
TNG; $z = 1$	10.25 ± 0.2
TNG; $z = 2$	10.15 ± 0.2
EAGLE; $z = 0$	10.30 ± 0.2
EAGLE; $z = 1$	9.85 ± 0.2
EAGLE; $z = 2$	9.55 ± 0.2

‘MOONRISE-Like’ data sets (top and bottom rows in each figure). This serves to highlight the differences in the distributions expected to occur as data moves from a simulated 6D phase space to an observational-like 2D+z space. Each panel ($z = 0 - 2$, from left to right) contains the distributions for the total sample (shown in faint

gray), centrals (shown in blue), and satellites (shown in red). The distributions of satellites in high and low-mass groups ($M_{\text{Halo}} \geq 10^{13} M_{\odot}$ and $M_{\text{Halo}} \leq 10^{13} M_{\odot}$, respectively) are included as solid and dashed open histograms, respectively.

The central galaxy population is clearly uni-modal and Gaussian-like in both simulations and at all epochs probed. Conversely, the satellite population is bimodal, most pronouncedly at $z = 0$. As such, we fit a single Gaussian to the central population within each panel, and a double Gaussian to the satellite populations. This enables us to determine the characteristics of the low and high density populations in satellites, and their evolution with redshift. The mean of each Gaussian of the bimodal fits are included within each panel. To quantify the change in the environment of satellites with redshift, we display the total fraction of satellites, as well as the fraction of satellites which belong to each Gaussian of the bimodal fit, expressed as P_1 (under-dense region) and P_2 (over-dense region) in the panels.

As redshift increases, the population fraction of the over-dense

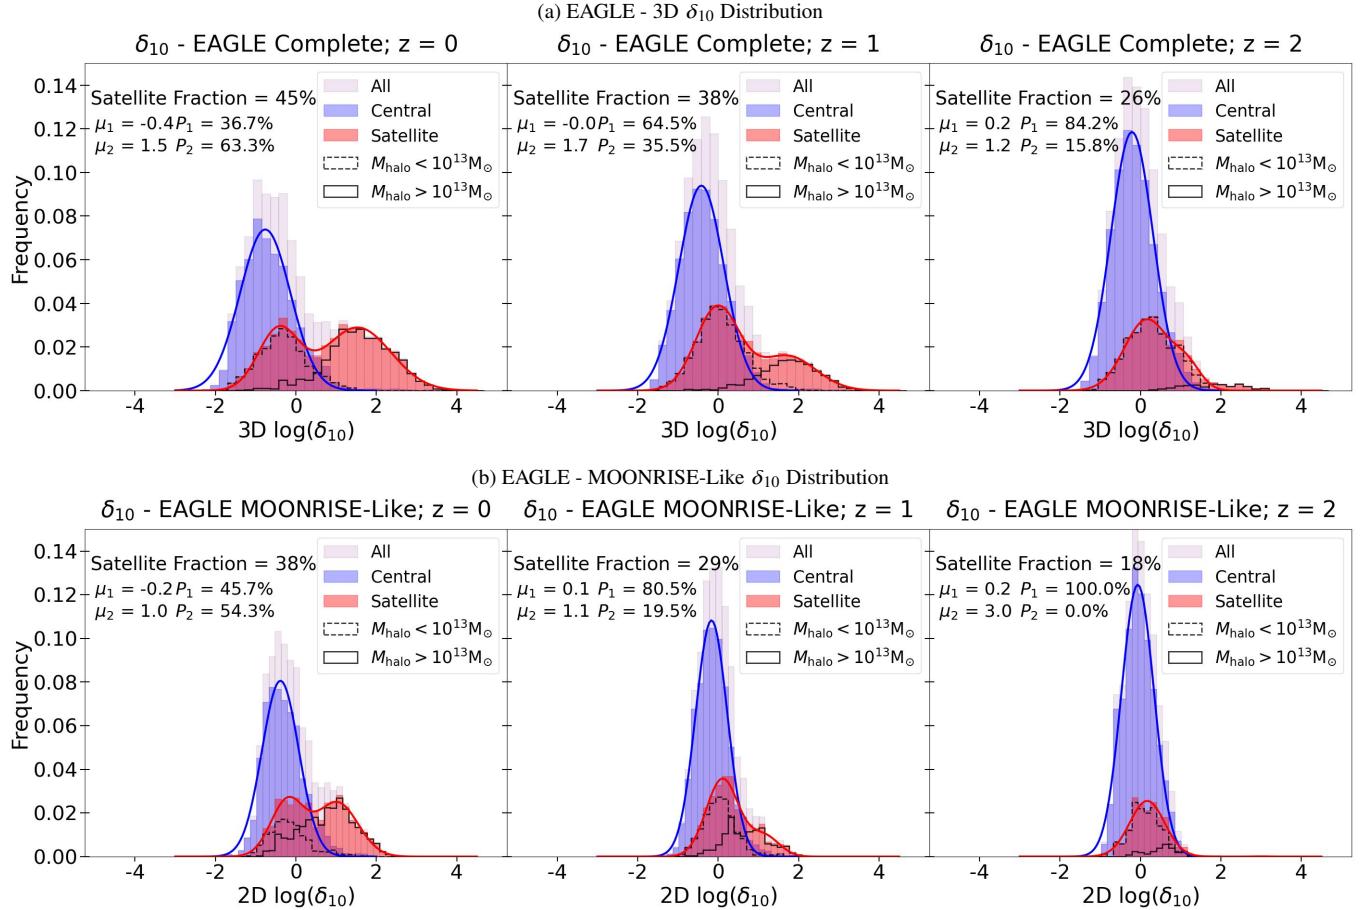


Figure 4. Identical in structure and method to Fig. 3, but here presenting the evolution of the distribution in galaxy over-density for the EAGLE simulation. The high density regime satellite population, and total satellite fraction, increase with cosmic time, indicating the growth of environment (as seen before for TNG). Furthermore, the distribution becomes increasingly bimodal, becoming a better tracer to identify differences in the environment. This is further highlighted by the distributions of satellites in high and low-mass groups, which become increasingly separate as redshift decreases. As seen for TNG in Fig. 3, the distributions of δ_{10} for the ‘MOONRISE-Like’ data sets are less bimodal and more curtailed in their density range than for the complete data set. This again indicates information loss going from 6D phase space to an observation like 2D+z phase space. Nonetheless, it is still possible to see a growth in the high density population as redshift decreases.

peak decreases, and the distributions begin to concentrate toward the under-dense peak, independent of simulation. At $z = 2$, there is a nearly uni-modal distribution, with very few satellites left in high density regions. Furthermore, there is a clear anti-correlation between the total satellite fraction and redshift, which grows from 28% to 47% from $z = 2$ to $z = 0$ in TNG, and grows from 26% to 45% over the same time period for EAGLE.

The evolution of the over-dense region of the bimodal distribution with redshift, in combination with a growing satellite fraction, points towards an increase in the abundance of high-mass groups and clusters, as groups merge, over cosmic time. This can also be seen by the evolution in the distribution of satellites in low and high-mass groups. There is a clear growth in the fraction of satellites in high-mass groups as redshift decreases. Additionally, the low and high-mass group split performs well in recovering the bimodality in δ_{10} present for all satellites. This is true at all redshifts, but is increasingly apparent as redshift decreases. This points towards density as an effective tracer of halo mass. Furthermore, this hierarchical assembly will convert past centrals into current satellite galaxies. As such, it is interesting to probe the effect of an ever changing environment on the quenching of satellites across redshift.

The general trends in the evolution of local galaxy density are

stable to the ‘MOONRISE-Like’ sample selection (lower panel sets in Figs. 3 & 4). However, there is an evident reduction in the density range, and, to a lesser degree, in the bimodality of the distributions. It can also be seen that splitting satellites by halo mass is less effective at separating the low and high density satellite populations when compared to the complete data set. This indicates a significant loss of information when moving from a 6D phase space to 2D+z observer space. As such, we will evaluate all quenching trends for both the base simulation (revealing the inner workings of the model) and for the ‘MOONRISE-Like’ sample (revealing what ought to be seen in MOONRISE, if the simulation is an accurate representation of nature).

4.3 Random Forest Classification

We present the results from a series of Random Forest quenching classification analyses of the TNG and EAGLE simulations in Figs. 5 & 6, respectively. Each bar represents the mean of the relative importances of a given parameter from a series of 100 independent Random Forest classification runs. The errors are given by the standard deviation across the 100 runs. Every row corresponds to a galaxy class (from top to bottom: centrals, high-mass satellites, low-

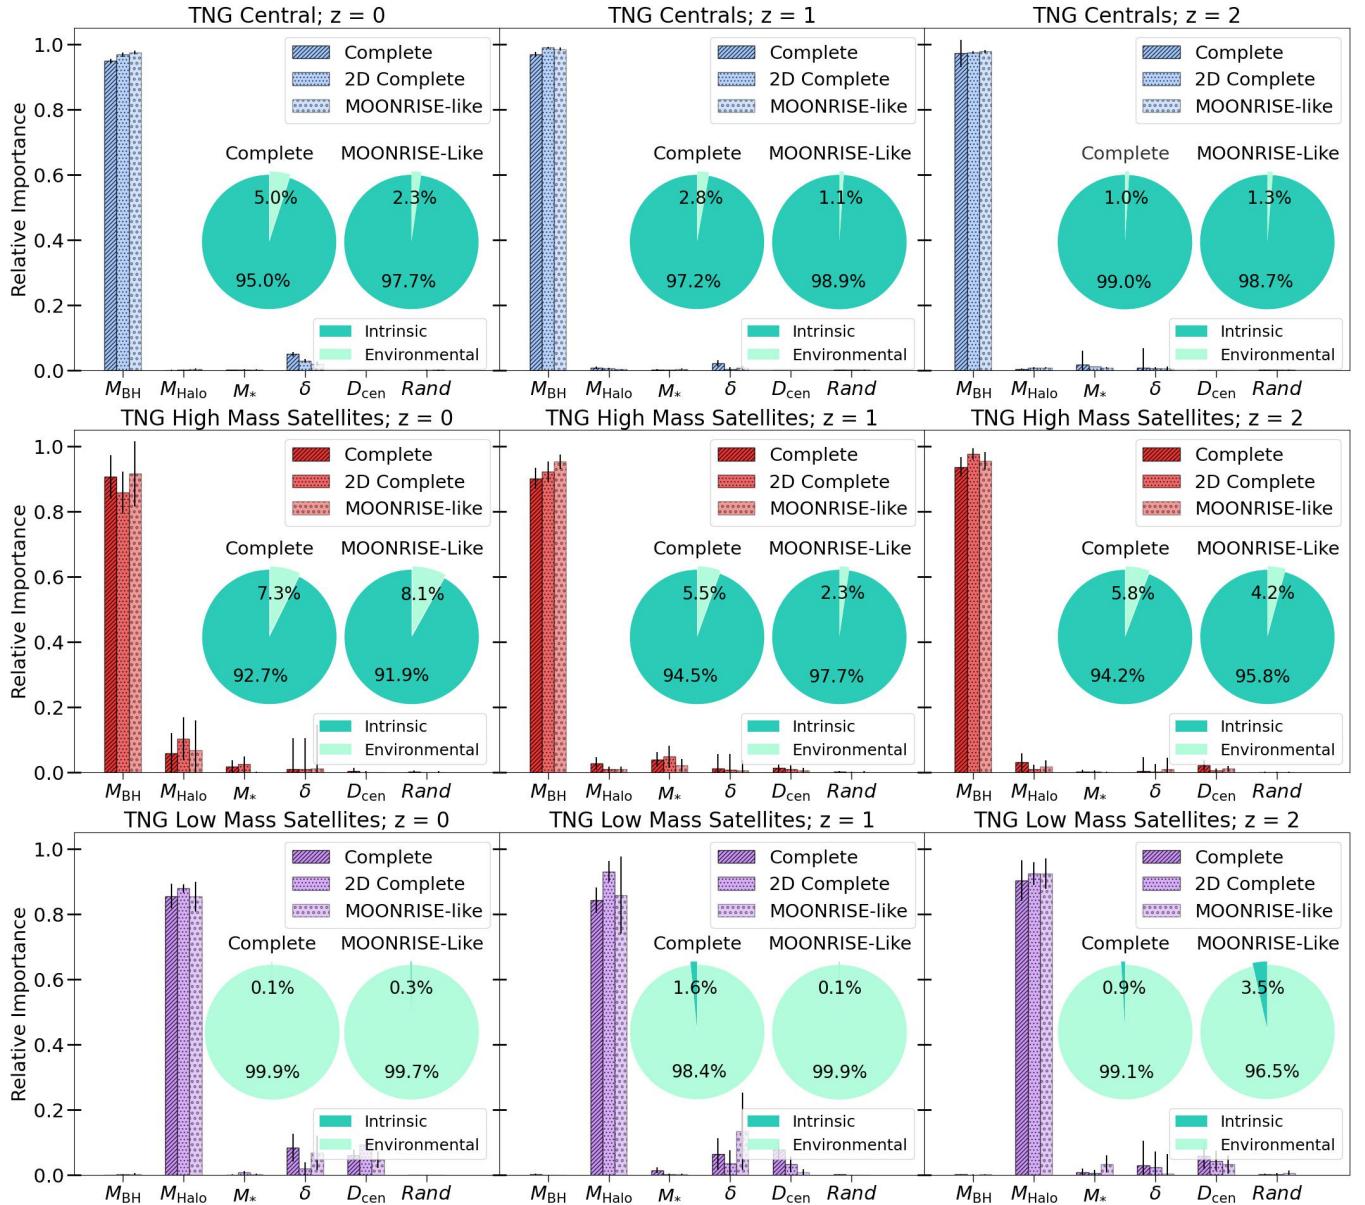


Figure 5. Random Forest quenching classification analyses for TNG. Each row corresponds to a galaxy class (from top to bottom: centrals, high-mass satellites, and low-mass satellites), and each column to a given redshift ($z = 0 - 2$, from left to right). The input features tested are listed along the x -axis, with their relative quenching importances indicated by the y -axis bar heights. Uncertainties are given by the variance of 100 independent classification runs. Separate bars are shown for the complete, the 2D complete, and ‘MOONRISE-Like’ datasets, indicated by different hatches on the bars (see legends). A pair of pie charts are displayed on each panel, indicating the total importance to quenching of intrinsic and environmental parameters for the complete (left) and ‘MOONRISE-Like’ (right) data sets. Across all redshifts and data samples, we find that centrals and high-mass satellites quench predominantly by intrinsic processes, with M_{BH} being by far the best predictor of quiescence. Conversely, low-mass satellites are determined to quench through environmental processes, with M_{Halo} found to be the dominant parameter. This indicates that in the TNG simulation both centrals and high-mass satellites quench predominantly via AGN feedback, yet low-mass satellites quench via environmental mechanisms at all epochs from $z = 0 - 2$. This provides a clear, testable prediction for the VLT-MOONRISE survey.

mass satellites), and each column to a redshift ($z = 0, 1, 2$; from left to right). Within each panel, results are shown separately for three data sets: ‘Complete’, ‘2D Complete’, and ‘MOONRISE-Like’, as indicated by different hatches on the corresponding bars (see legends). A pair of pie charts are displayed on each panel, indicating the total importance to quenching of intrinsic and environmental parameters for the complete (left) and ‘MOONRISE-Like’ (right) data sets. We present the hyperparameters used for each RF classification test presented in this section in Appendix B.

The input features tested are: M_{BH} , M_* , M_{Halo} , δ , D_{cen} , where δ is the sum of the relative importances of the third, fifth, and tenth nearest neighbor densities. Additionally, a number drawn from a random distribution between 0 and 1 is added to gauge the impact of random uncertainty in this method. The panels are accompanied by a pair of pie charts, showing the total relative importance of intrinsic (M_{BH} and M_*) and environmental (M_{Halo} , δ , D_{cen}) parameters for the complete (left pie plot) and ‘MOONRISE-Like’ (right pie plot) data sets. This is to facilitate the comparison in importance to quenching

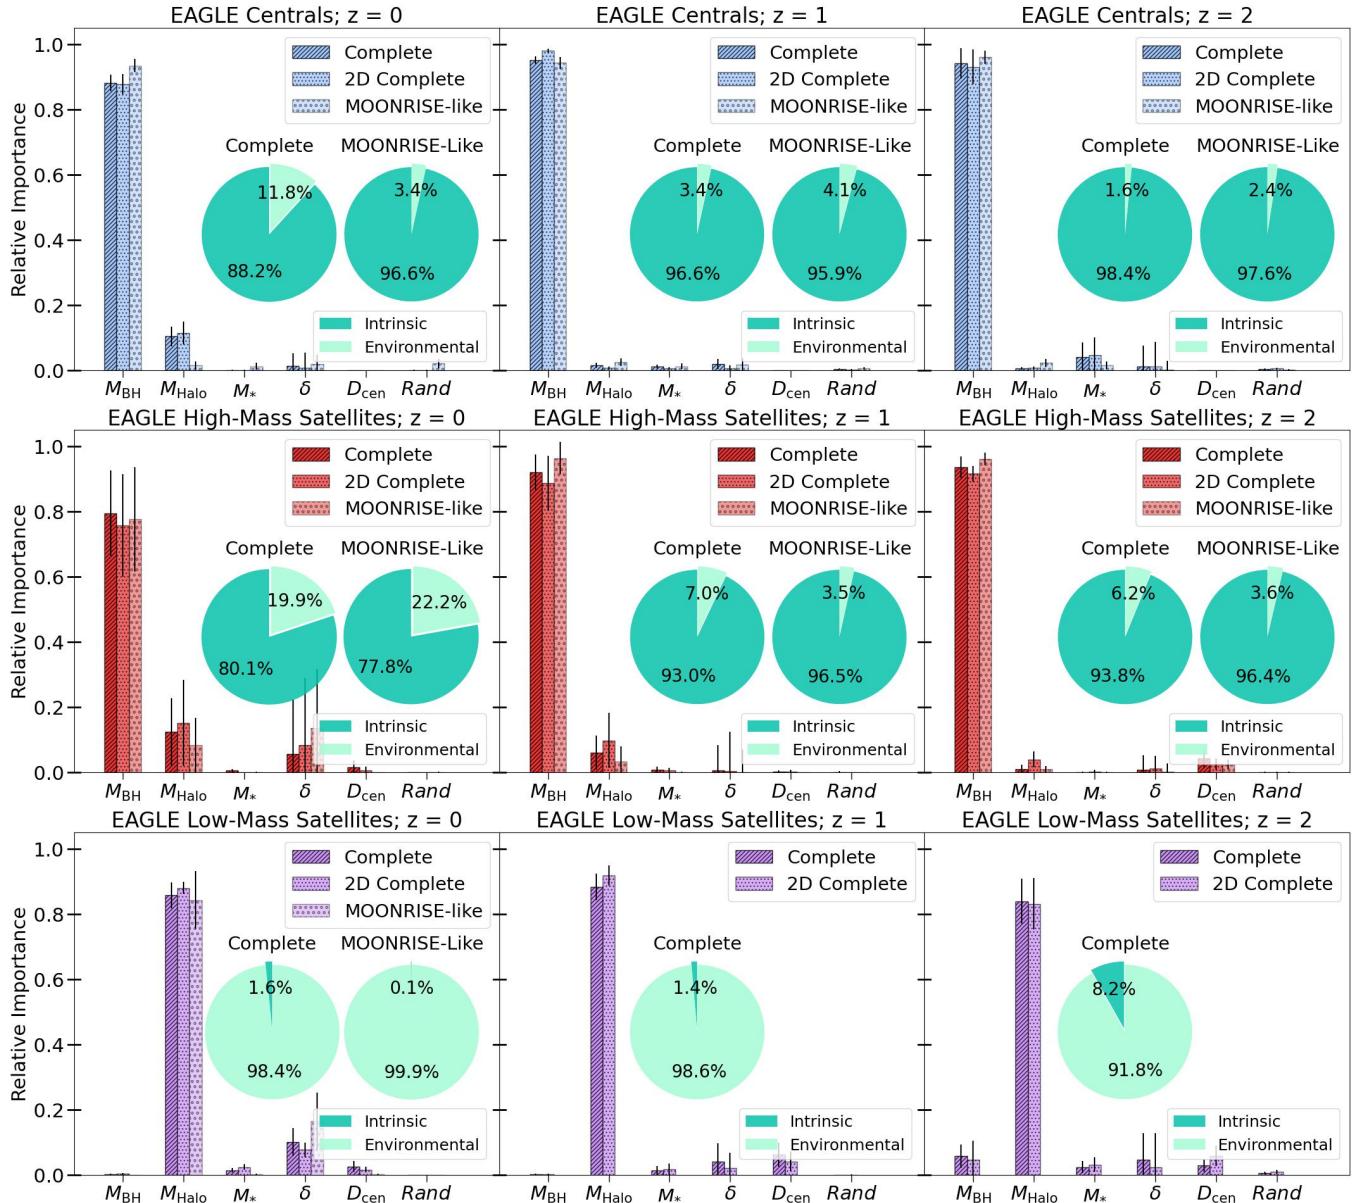


Figure 6. Random Forest quenching classification analyses for EAGLE. This figure is structured exactly like Fig. 5. As seen for TNG, at all redshifts and independent of data sample, the quiescence of centrals and high-mass satellites is best predicted by intrinsic parameters, with M_{BH} being the dominant feature. For the complete and 2D complete data sets, low-mass satellites are determined to quench through environmental processes at all redshifts, with M_{Halo} found to be the dominant parameter. However, as the mass completeness threshold of the ‘MOONRISE-Like’ data set is near, or greater than, the low-mass thresholds at $z = 1$ and $z = 2$, virtually no low-mass satellites are present in these samples. This implies that environmental quenching would not be detectable in MOONRISE, if the EAGLE model is correct. The results for the low-mass satellites of the ‘MOONRISE-Like’ data set at $z = 0$ are in agreement with the complete and 2D complete results, with environment dominating. Taken together, this figure provides numerous clear predictions for quenching across cosmic time.

between intrinsic and environmental processes as a function of galaxy type, as well as the evolution with redshift.

Focusing first on central galaxies, we find that M_{BH} dominates quenching at all redshifts, independent of the sample set, for both TNG and EAGLE. This implies an intrinsic quenching mechanism coupled to black hole growth, explicitly AGN feedback driven quenching (as seen in Piotrowska et al. 2022; Bluck et al. 2023, 2024). Furthermore, the pie plots clearly identify quenching as being intrinsic in origin, with environment attributed $\lesssim 15\%$ of importance across redshifts.

The second row of plots presents the classification results for

high-mass satellites. As with the central galaxies, for both simulations we find black hole mass to be the best predictor of quiescence at all redshifts. This points to an AGN feedback driven quenching process for high-mass satellites, in both simulations. However, a secondary dependence on environment (specifically halo mass) is clearly present. The ‘quenching importance’ pie plots confirm and highlight this intrinsic quenching dominance in both the complete and ‘MOONRISE-Like’ data sets. Environmental parameters make up less than 10% and 20% of the pie at all redshifts for the complete data sets of TNG and EAGLE, respectively. This is particularly inter-

esting because historically many authors have considered all satellites to be environmentally quenched (see, e.g., Peng et al. 2012).

Overall, the results of the Random Forest classification analyses for the TNG and EAGLE simulations are in strong agreement for centrals and high-mass satellites. We find their quiescence to be best predicted by black hole mass, which is known to be intricately linked to AGN feedback in these simulations. This is consistent across redshifts and data sets. However, it seems the AGN feedback mechanism may be more effective at higher redshifts, as the relative importance in M_{BH} increases with redshift. This is most apparent for the EAGLE simulation, particularly regarding high-mass satellites, which lacks the preventative AGN feedback mode present in TNG. We further explore this later in the paper, specifically in Sec. 4.4.

The results of the RF classification analyses for low-mass satellites are shown in the bottom rows of Figs. 5 & 6. In both TNG and EAGLE, it is clear that this population is quenched via a purely environmental mechanism at all redshifts, and for all data sets. Specifically, we find M_{Halo} to be the dominant parameter, with weak secondary dependencies on δ and D_{cen} . These lesser importances indicate preferential quenching in the highest density environments, denoted by the importance of δ , and towards the center of haloes, suggested by the importance of D_{cen} (see Sec. 4.5.2).

However, a glaring difference between the simulations appears when focusing on the ‘MOONRISE-Like’, low-mass satellite populations. While TNG shows consistent results across redshifts, the EAGLE simulation suffers from a dearth of low-mass satellites at $z = 1$ and $z = 2$, due to increasingly low mass thresholds, as seen in Table 3. Therefore, as opposed to TNG, EAGLE predicts that environmental quenching will be undetectable in VLT-MOONRISE. This is a fundamental difference in the predictions formed from the simulations, and will prove useful in determining which model is more accurate.

4.3.1 MOONRISE-Like with limited observational parameters

Stellar mass and local over-density, rather than black hole mass and halo mass, are often used as parameters to investigate intrinsic and environmental quenching when analyzing observational data (Peng et al. 2010, 2012). Therefore, to obtain results comparable to other studies, we conduct the Random Forest classification tests for the ‘MOONRISE-Like’ data sets, absent both black hole mass and halo mass, so as to highlight the performances of stellar mass and local over-density. The input features tested for here are therefore: M_* , δ , D_{cen} , and a randomly generated number between 0 and 1. We present the hyperparameters used for each RF classification test in this section in Appendix B.

The results for these classification tests are shown in Figs. 7 & 8, for TNG and EAGLE (respectively). Importantly, as VLT-MOONRISE will be a wide field spectroscopic survey it will be possible to conduct accurate central - satellite classification. As such, the panels are organized by galaxy class, i.e., centrals, high-mass satellites, and low-mass satellites (from left to right). Within each panel, we present the results at $z = 0$, 1, and 2, indicated by the different colored bars (see legends).

The results of the Random Forest classification tests for centrals and low-mass satellites of the TNG ‘MOONRISE-Like’ and EAGLE ‘MOONRISE-Like’ data sets, absent M_{BH} and M_{Halo} , are broadly consistent with the results presented in Figs. 5 & 6. The quiescence of central galaxies is best predicted by intrinsic parameters, here stellar mass. However, the underlying mechanism of AGN feedback could be missed if the black hole mass is absent as a parameter. This result establishes that mass quenching (e.g., as in Peng et al.

2010, 2012) is ultimately a result of black hole mass quenching (at least in these simulations), and hence ultimately AGN feedback (see also Terrazas et al. 2016, 2017; Piotrowska et al. 2022; Bluck et al. 2016, 2020a,b, 2022, 2023, 2024, for further evidence of this in both simulations and observation). On the other hand, quenching in low-mass satellites is determined by the local over-density parameter, indicating environmental quenching, at all redshifts (absent EAGLE at $z = 1$ & 2, as nearly no quenched low-mass satellites are present). However, crucially, halo mass is predicted to be the superior environmental parameter regulating quenching in both simulations (and at all epochs), when available.

Comparing the results of the classification tests for high-mass satellites, absent black hole and halo mass (Figs. 7 & 8) to those seen in Figs. 5 & 6, a notable difference in the performance of intrinsic parameters is apparent. Absent M_{BH} , δ is found to best predict quiescence at $z = 0$ for both simulations. For TNG, the dependence on local over-density wains in strength as redshift increases, as M_* carries a similar importance as δ at $z = 1$, and dominates at $z = 2$. However, for EAGLE, δ remains the dominant parameter at $z = 1$. At $z = 2$, M_* becomes the best predictor of quiescence, although a strong secondary dependence on environmental parameters remains. We find these results to be consistent with Peng et al. (2010, 2012), wherein the quenching of satellites in the local Universe is determined to operate primarily environmentally (regardless of mass).

However, ultimately the quiescence of high-mass satellites in the local Universe is directly linked to their supermassive black hole (as demonstrated in Section 4.3). This disagreement can be resolved by linking a galaxy’s local density to its black hole mass. According to the morphology-density relation (e.g., Postman & Geller 1984; Goto et al. 2003; Teklu et al. 2017), galaxies become progressively more spheroidal in higher density environments. Additionally, supermassive black holes tend to be more massive in spheroidal galaxies than disk galaxies, at a fixed stellar mass (see, e.g., Saglia et al. 2016; Piotrowska et al. 2022), establishing a link between a galaxy’s local over-density and the mass of its black hole.

Therefore, in the case of high-mass satellite quenching, δ is likely correlated with M_{BH} (at a fixed stellar mass). Hence, when black hole mass is absent, and the range of stellar masses is reduced (as it is for this sub-sample), the local density may trace black hole mass, or at least its link to quiescence, better than any other available parameter, especially for $z < 2$. Consequently, it is crucial to obtain as accurate as possible spectroscopic measurements to best estimate values of black hole mass in observational surveys (e.g., VLT-MOONRISE). Failure to do so may result in erroneously attributing intrinsic satellite quenching to environment.

4.4 Quenched Fraction Evolution

In this sub-section we explore the evolution with redshift of the quenched fraction of central and satellite galaxies, for the ‘MOONRISE-Like’ data sets of both TNG and EAGLE. In Figs. 9 (TNG) & 10 (EAGLE) we present the quenched fractions of centrals and satellites as a function of M_{BH} . Black hole mass represents the most predictive quenching parameter for central galaxies, enabling one to assess how other quenching routes impact galaxies once this channel is accounted for. The quenched fraction of central galaxies is shown in blue, while the quenched fraction for satellite galaxies is presented in violet. We compute the uncertainty as the Poisson statistics uncertainty in each bin.

We first focus on the results for central galaxies from the TNG and EAGLE simulations. Across all redshifts, the quenched fraction as a function of M_{BH} sits at nearly zero for low black hole mass for both

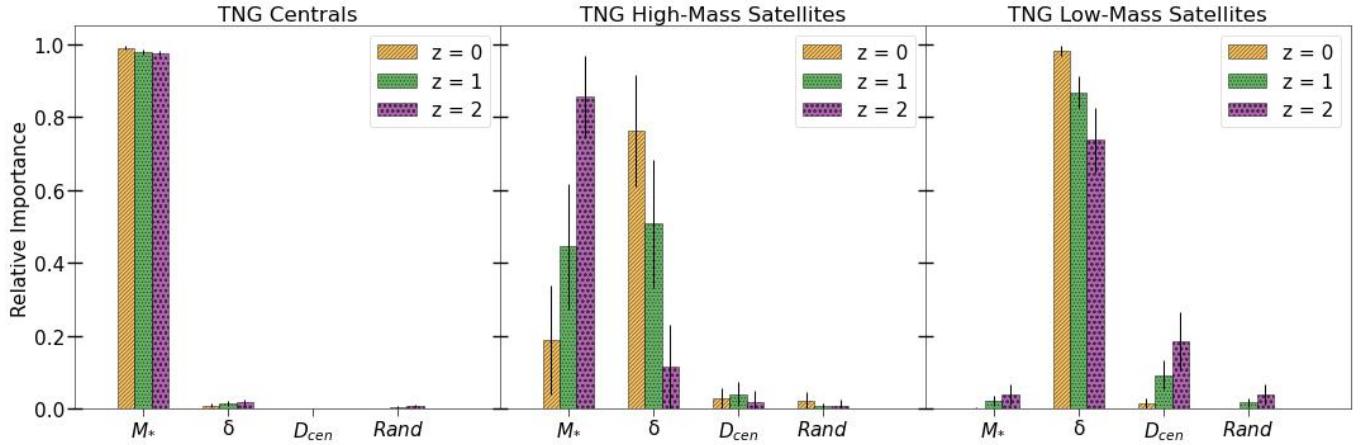


Figure 7. Random Forest quenching classification analyses for the TNG ‘MOONRISE-Like’ data set, absent M_{BH} and M_{Halo} . Each panel corresponds to a galaxy class (centrals, high-mass satellites, low-mass satellites), with the results at each redshift presented with different colored bars on each panel. The parameters used for training the Random Forest are listed along the x -axis, with their relative quenching importances indicated by the y -axis bar heights. Uncertainties are given by the variance of 100 independent classification runs. We find the quiescence of centrals to be unanimously predicted by stellar mass at all redshifts. In contrast, the quenching of high-mass satellites is best predicted by δ at $z = 0$, though a secondary dependence on M_* is present. As redshift increases, so does the relative importance of M_* . At $z = 1$, stellar mass and δ share similar importance (50-40%), whereas M_* dominates at $z = 2$. Finally, we find that δ dominates at all redshifts for low-mass satellites, although a small contribution from D_{cen} is present at $z = 1$ & 2.

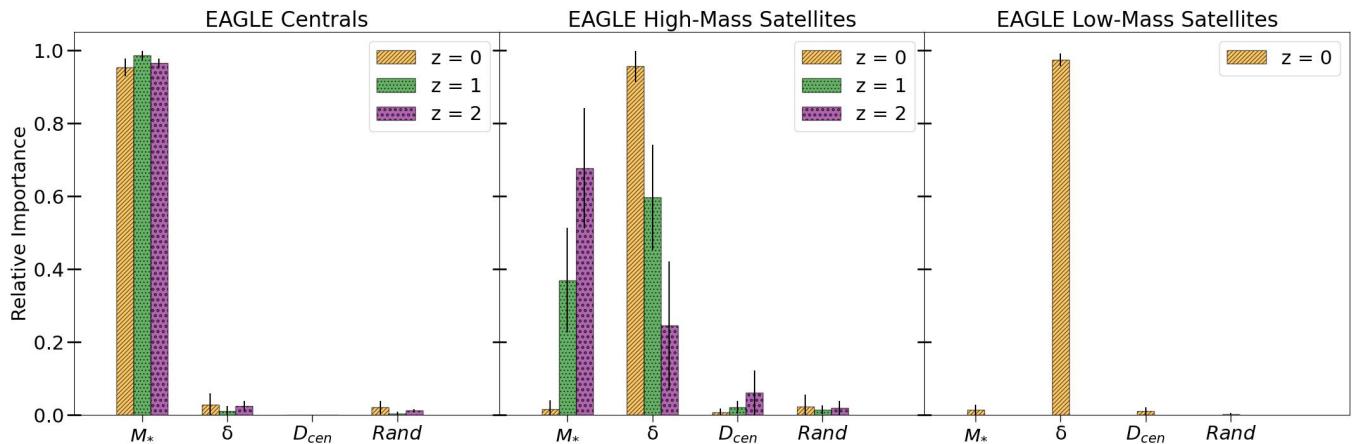


Figure 8. Random Forest quenching classification analyses for the EAGLE ‘MOONRISE-Like’ data set, absent M_{BH} and M_{Halo} . This figure is structured identically to Fig. 7. We find that central galaxies are predicted to quench intrinsically, with M_* dominating at all redshifts, as for TNG. High-mass satellites present with δ dominating at $z = 0$, and it remains the best predictor at $z = 1$. At $z = 2$ stellar mass is the best predictor of quiescence. Finally, at $z = 0$, the quenching of low-mass satellites is best predicted by δ . The results for low-mass satellites are absent at $z = 1$ and $z = 2$, due to the mass completeness limit of the ‘MOONRISE-Like’ data sets and the relatively low mass threshold for environmental quenching in EAGLE.

simulations ($\log(M_{\text{BH}}/\text{M}_\odot) \leq 8$ for TNG, and $\log(M_{\text{BH}}/\text{M}_\odot) \leq 7$ for EAGLE). Concentrating on TNG centrals, the quenched fraction behaves similarly at all redshifts, rapidly increasing once beyond the low mass regime, then plateauing around $f_Q = 0.9$. On the other hand, the quenched fraction of centrals from the EAGLE simulation differs greatly from redshift to redshift. At $z = 2$, f_Q steadily rises, resulting in a near fully quenched population at the highest black hole masses. At $z = 1$, however, the quenched fraction first rises then quickly plateaus, reaching a maximal value of $f_Q \simeq 0.6$. There is even a downward trend once $\log(M_{\text{BH}}/\text{M}_\odot) \geq 8$. At $z = 0$, f_Q reaches a maximum at the relatively low value of $f_Q = 0.5$, and there is further evidence of a downwards trend at high black hole masses, as seen at $z = 1$. Indeed, EAGLE explicitly predicts that the higher M_{BH} central galaxies will be star forming at least $\sim 50\%$ of the time, which is in stark contrast to observations (see, e.g., Peng et al. 2010; Bluck et al.

2016; Terrazas et al. 2016, 2017; Piotrowska et al. 2022; Bluck et al. 2022).

Whilst the results for centrals from TNG and EAGLE agree that low black hole mass centrals are unable to quench, they exhibit a very different evolution with redshift for higher black hole masses. The evolution of the quenched fraction of centrals for the EAGLE simulation demonstrates clear signs of rejuvenation - i.e., formally quenched galaxies restart star formation at late cosmic times. This is likely due to the absence of an explicitly preventative AGN feedback mode in the EAGLE simulation (Booth & Schaye 2009; Schaye et al. 2015; Crain et al. 2015).

The EAGLE simulations makes use of an effective thermal feedback mode which drives outflows and enables the rapid quenching of galaxies by expelling gas at high redshifts (where the most powerful AGN reside, (Aird et al. 2010; Fabian et al. 2014; Aird et al. 2018).

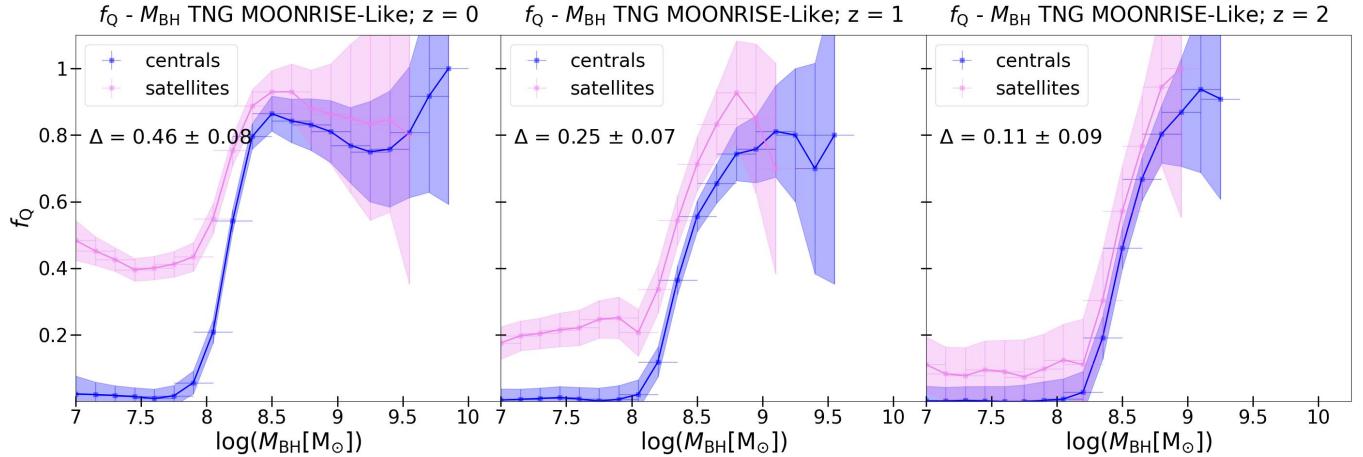


Figure 9. The quenched fraction of TNG central (shown in blue) and satellite (shown in violet) galaxies as a function of black hole mass for the ‘MOONRISE-Like’ data set. The results are shown for different redshifts from left to right along each row (as labeled by the panel titles). At $z = 0$, there is a clear separation of the two populations in the low-mass regime, wherein satellites are more likely to be quenched than centrals. This is a strong indicator of environment quenching, operating at fixed black hole mass. The magnitude of the offset, indicated by Δ , decreases strongly with redshift. Hence, we conclude that environmental quenching is more effective at late cosmic times. On the other hand, at high masses, and all redshifts, the satellite quenched fraction resembles the central quenched fraction. As such, one would expect quenching at high mass to be operated by the same mechanism (AGN feedback), regardless of central - satellite classification.

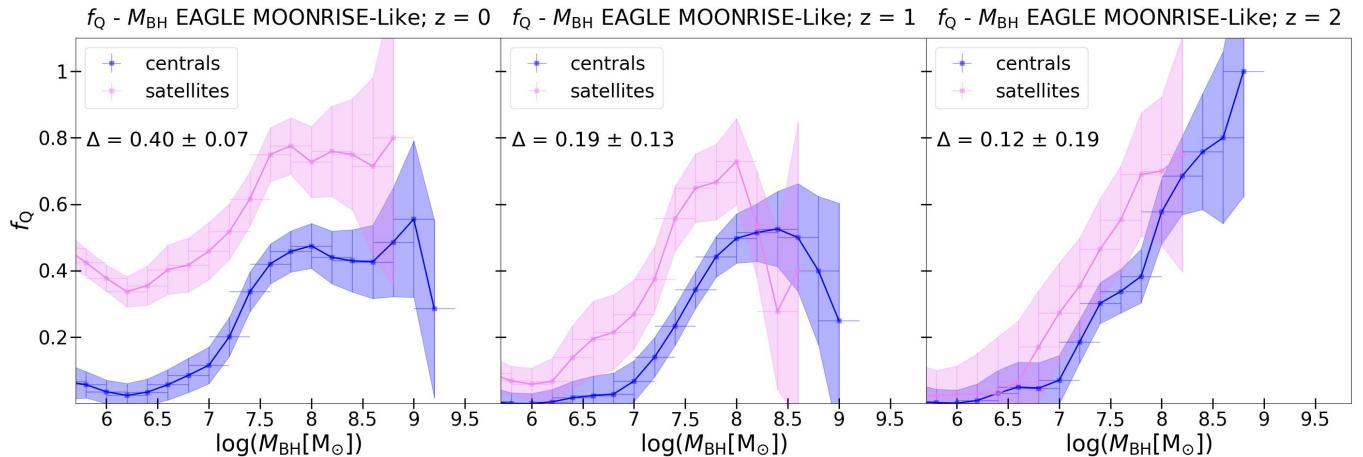


Figure 10. This figure is identical in structure to Fig. 9, but here showing results from the EAGLE simulation. In comparison to TNG, most trends are similar, with one important exception. The quenched fraction of central galaxies plateaus at a much lower value at $z = 0$ & 1. This can also be seen for satellite galaxies at $z = 1$, though it is less apparent at $z = 0$. This is indicative of the EAGLE quenching mechanism failing to keep galaxies quenched over long cosmic times, leading to star formation rejuvenation. There remains a noticeable gap between the quenched fractions of satellites and centrals at low black hole mass, especially at $z = 0$. This gap points towards an environmentally dominated quenching process for low-mass satellites, which decreases as redshift increases. Additionally, while the satellite and central quenched fractions are similar in trends at high masses, the satellite quenched fraction is elevated in comparison to the central quenched fraction. Therefore, while this does imply a similar causal origin to the quenching of high mass satellites to that of centrals, it is clear other factors are aiding in the quenching of satellites, even at higher masses.

However, the feedback prescription employed in EAGLE weakens as low Eddington ratios dominate. Therefore, as f_{Edd} decreases with redshift, so does the potential impact incurred by AGN on their galaxies and surrounding haloes. Furthermore, there is no explicit mechanism to prevent the galaxy from accreting gas from the hot halo via cooling at late cosmic times. These effects in combination may lead to the rejuvenation of star formation. This situation is in stark contrast with TNG, which shows no sign of late-time rejuvenation. This is almost certainly a result of TNG employing a ‘kinetic mode’ of AGN feedback, which acts as a preventative feedback mode, keeping the halo gas hot. Crucially, the kinetic mode is most efficient at low Eddington ratios, which are more common at late cosmic times. As the gas of the CGM is unable to cool (due to continued

heating), collapse into the galaxy at late cosmic times is suppressed, enabling long-term quiescence (see, [Weinberger et al. 2017](#); [Nelson et al. 2018](#); [Pillepich et al. 2018a](#); [Zinger et al. 2020](#)).

We turn our attention now to satellite galaxies. For TNG, at higher black hole masses the quenched fractions of satellites and centrals are remarkably similar at all redshifts. However, for EAGLE, while the satellite quenched fraction is similar in trend to the central quenched fraction, it is markedly more elevated than the latter at high mass, especially at $z = 0$. On the other hand, while rejuvenation is best seen in the quenched fraction of centrals, there is also clear evidence of it for high black hole mass satellites as well, particularly at $z = 1$. This demonstrates the similarity in the quenching mechanisms of high-mass satellites and central galaxies, driven by their supermas-

sive black hole (as confirmed by the RF analyses in the previous subsection). However, given the elevated quenched fraction of high black hole mass satellites at $z = 0$, there appears to be some processes, likely environmental, which aids quenching in, or prevents the rejuvenation of, high-mass satellites.

At low black hole mass, the satellite and central quenched fractions differ to varying degrees at different redshifts. To quantify the offset between the central and satellite quenched fractions in the low M_{BH} (or low M_*) regime, we define Δ , the maximum difference between the satellite and central quenched fractions in said region, indicated within each panel. The error in Δ is determined by the sum in quadrature of the Poisson uncertainties for the associated bins of the satellite and central quenched fractions. As such, Δ serves as an indicator for the impact of environment on quenching at each redshift. For both TNG and EAGLE at $z = 2$, the satellite quenched fraction is only slightly raised when compared to that of centrals. However, as redshift decreases, Δ increases. For TNG, the offset increases from $\Delta = 0.11$ at $z = 2$, to $\Delta = 0.25$ at $z = 1$, and $\Delta \simeq 0.45$ at $z = 0$. Focusing on the EAGLE simulation, there is a similar trend as Δ grows from about $\Delta \simeq 0.12$ at $z = 2$ to $\Delta \simeq 0.4$ at $z = 0$. The growth in separation of the quenched fractions of each population indicates an increase in the impact of environmental quenching for low-mass satellites as redshift decreases. This is likely due to the formation and growth of high-mass cluster environments, and therefore high density regions, through hierarchical assembly over cosmic time (as is also evident in Figs. 3 & 4).

It is important to note that a similar analysis, showing the quenched fraction as a function of stellar mass, is performed for EAGLE in Schaye et al. (2015) at $z = 0.1$, albeit with a different definition for the quenching threshold. Comparing our findings for EAGLE at $z = 0$ to those of the aforementioned study, there is good agreement at low mass but a stark difference at higher stellar, and therefore black hole, masses ($M_* \sim 10^{11} M_\odot$). Indeed, while we find the quenched fraction, particularly of central galaxies, to plateau and even show a downturn at high black hole mass, this is not present in Schaye et al. (2015). However, due to the sparsity of data at these stellar masses in EAGLE, in Schaye et al. (2015) the quenched fraction beyond $M_* \sim 10^{11} M_\odot$ is said to be ‘unreliable’, and this region is presented as an extrapolation (shown as a dotted line). The results shown in Fig. 10 are supported by several previous works (Bahé et al. 2017; Donnari et al. 2021; Piotrowska et al. 2022; Bluck et al. 2023). Therefore, while there may appear to be disagreement between the results presented and previous analyses, this is due to a difference in methodology. The actual quenched fraction at high masses in EAGLE is lower than presented in Schaye et al. (2015). Whilst it remains possible that this is partly due to a lack of high mass galaxies simulated, the fact that this trend is absent in TNG (which simulates a similar volume) leads us to conclude that this is more likely a result of the EAGLE AGN feedback mechanism not being effective at maintaining long-term quiescence in high-mass galaxies and haloes.

4.5 Satellite quenching as a function of environment

In the following two sub-sections we present the quenched fraction of satellites as function of several environmental parameters, for the ‘MOONRISE-Like’ data sets of TNG and EAGLE, at $z = 0, 1, 2$. To ensure solely environmental effects are probed, we compute and analyze: Δf_Q . This statistic is defined as the difference in the quenched state of a given satellite to a control sample of centrals with similar black hole masses ($|\log(M_{\text{BH}_i}) - \log(M_{\text{BH}_{\text{cen}}})| \leq 0.2$). Given intrinsic quenching (specifically AGN feedback) is the sole avenue to quench centrals in both simulations, Δf_Q probes the effect

of additional quenching channels (e.g., through environment) on the satellite population in different scenarios.

In Sec. 4.5.1 we compute and compare Δf_Q as a function of stellar mass for two satellite populations, those in high-mass groups, defined as $M_{\text{Halo}} \geq 10^{13} M_\odot$, and low-mass groups, $M_{\text{Halo}} \leq 10^{13} M_\odot$. We then explore how a galaxy’s location within its group or cluster affects quenching, presented in Sec. 4.5.2. To do so we evaluate Δf_Q as a function of D_{cen} , the distance of a satellite to its central, normalized by the virial radius of the halo.

4.5.1 Quenched Fraction of satellites in clusters and groups as a function of stellar mass

In Figs. 11 (TNG) & 12 (EAGLE) we present Δf_Q for satellites in high-mass groups (presented in violet), and low-mass groups (presented in blue), as a function of stellar mass for the ‘MOONRISE-Like’ data sets. This enables direct comparison to the upcoming VLT-MOONRISE galaxy survey. The panels are ordered with increasing redshift from left to right. To quantify the maximal impact environment may carry on quenching, we define ζ , which is computed by the Δf_Q value at the lowest stellar mass probed (where environmental effects dominate). Furthermore, we perform an OLS fit to Δf_Q for satellites residing in high-mass groups, where environmental effects are significant. From these fits we extract the gradient, ∇ , quantifying the (anti) correlation between Δf_Q and stellar mass. We include the value of each of these parameters within each panel. In essence, this analysis reveals the stabilizing effect of increasing stellar mass on the impact of environment on star formation within satellite galaxies.

We start by analyzing the results for the TNG simulation, seen in Fig. 11. It is clear that satellites preferentially quench in high-mass groups at low stellar masses, while even low-mass satellites are unlikely to quench in low-mass groups. Furthermore, these features remains constant across redshifts. Notably, the quenching of satellites in halos is most efficient at $z = 0$, likely due to an increased abundance of the highest-mass groups. Most interesting is an anti-correlation of Δf_Q with stellar mass at all redshifts, as demonstrated by ∇ .

Honing in on $z = 0$ for the TNG simulation, $\Delta f_Q \sim 70\%$ at the lowest stellar mass, and then steadily decreases as stellar mass increases, reaching $\Delta f_Q = 0$ when $\log(M_*/M_\odot) \simeq 11$. While the specific values differ slightly across redshift ranges, the trend of an anti-correlation with stellar mass, as indicated by the negative values of ∇ , is constant across cosmic time.

The strong anti-correlation with stellar mass points towards an increased dependence on intrinsic mechanisms for satellite quenching as stellar mass increases. Therefore, it is not the central - satellite class that truly matters when comparing quenching mechanisms, but rather a low/ high-mass class. Low-mass satellite galaxies are subject to environmental quenching, but all high-mass galaxies are primarily subject to intrinsic (i.e., AGN feedback) quenching. This provides further evidence for the similar conclusions reached in Section 4.3, via the Random Forest analyses.

When comparing the results for the TNG and EAGLE simulations, a clear difference appears in the impact environment carries on quenching as redshift increases. It is evident in EAGLE that as redshift increases, environment plays a lesser role in quenching satellites. This is demonstrated by both a decreasing value of ζ , from 0.67 at $z = 0$ to 0.26 at $z = 2$, and a flattening of the decline in gradient of Δf_Q with increasing stellar mass (see panels).

This decline in the impact environment has on quenching is in line with the evolution of the Quench_E and Quench_I parameters shown in Fig. 2. While it is clear both TNG and EAGLE show an increasing

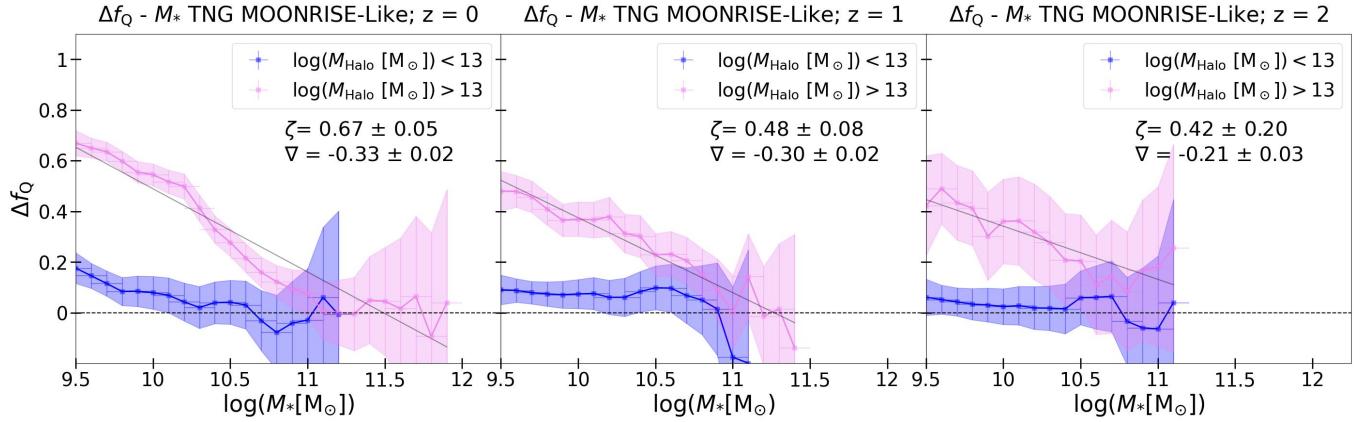


Figure 11. The quenching of satellites in TNG as a function stellar mass, separated by the mass of their parent halo (see legends). We present the delta quenched fraction (Δf_Q) as a function of stellar mass for the ‘MOONRISE-Like’ data sets. This statistic quantifies the enhancement in quenching of satellites over centrals, evaluated by an offset in quenched fraction from a central control sample with similar central black hole masses. The panels correspond to a given redshift, from $z = 0$ to $z = 2$ (left to right). Within each panel we include the values of the ζ (maximum offset) and ∇ (gradient) parameters, for the high mass group sample. Low mass groups experience weak environmental enhancement of quenching throughout the full range in stellar mass. At all redshifts, and independent of data set, low-mass satellites in high-mass groups/clusters undergo stronger quenching than centrals with a similar black hole mass. Meanwhile, Δf_Q has a clear anti-correlation with satellite stellar mass, demonstrated by the negative values of ∇ steadily declining as M_* increases, which is a strong indication that intrinsic quenching grows in relative importance with stellar mass. Low-mass satellites quench preferentially in high-mass groups, whereas the highest mass satellites are no more frequently quenched than centrals in either halo mass population at any epoch.

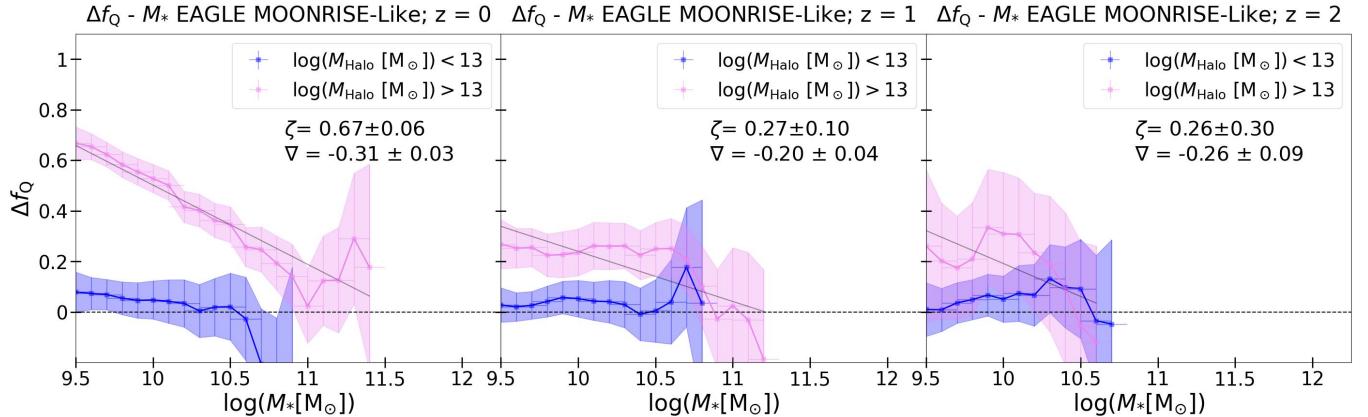


Figure 12. Identical in structure to Fig. 11, but here presenting results for EAGLE. The effect of environment on quenching is strongest at $z = 0$ and weakens significantly as redshift increases. This is demonstrated by the decrease of Δf_Q for the least massive satellites in high-mass groups as redshift increases, as quantified by ζ (displayed on panels). This is a strong indication that the impact of environment is weak at higher redshifts for the EAGLE simulation. On the other hand, the anti-correlation of Δf_Q with stellar mass, indicated by ∇ , demonstrates a growing similarity in quenching between satellites and centrals as black hole and stellar masses increase, as is also seen in TNG.

intrinsic quenching fraction with redshift, the effect is more drastic in EAGLE. In the TNG simulation, Quench_I shifts from 40% at $z = 0$ to 60% at $z = 2$, while for EAGLE Quench_I moves from 20% to 80% over the same redshift interval. This points to less effective environmental quenching at higher redshifts, especially for EAGLE, presumably due to a lack of high-mass halos and clusters at these very early cosmic time. Hence, the early Universe is predicted to be progressively more dominated by intrinsic quenching (as also seen in Kawinwanichakij et al. 2017; Xie et al. 2024).

Therefore, the simulations predict significantly different environmental impact on quenching at higher redshifts. The VLT MOONRISE survey will provide ideal observational data to test these differing predictions in the coming years.

4.5.2 Satellite quenching as a function of location within the halo

In Fig. 13, we present Δf_Q as a function of satellite location within a group, represented by the distance to the central galaxy (in units of virial radii), for two halo mass ranges (see legends) for the TNG ‘MOONRISE-Like’ data sets. We perform this analysis solely for low-mass satellites, as the quenching of high-mass satellites is largely independent of environment. As there are few quenched low-mass satellites present in the EAGLE ‘MOONRISE-Like’ data sets at $z = 1$, and none at $z = 2$, we present the results for TNG only. However, the results for EAGLE at $z = 0$ are consistent with those obtained for TNG.

At all redshifts, low-mass satellites struggle to quench when in a low-mass halo ($\log(M_{\text{Halo}}/M_{\odot}) < 10^{13}$). However, Δf_Q is slightly elevated for satellites near the center of their group, and gradually decreases as D_{cen} increases.

Analyzing Δf_Q in the most massive groups and clusters, the

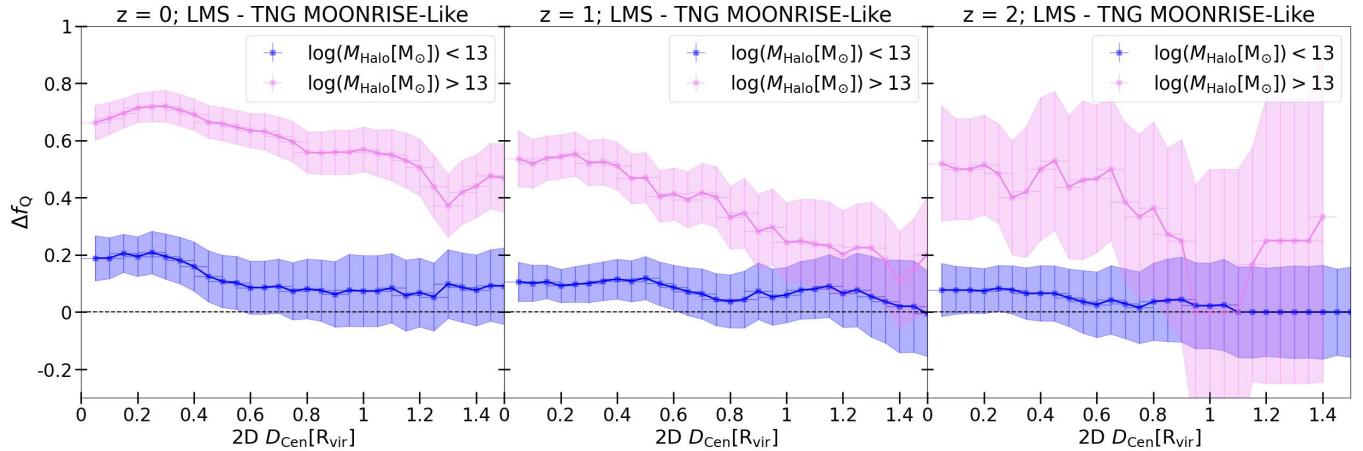


Figure 13. The quenching of low-mass satellites (‘LMS’ in titles) as a function of location within their groups or clusters for TNG. We present the delta quenched fraction (Δf_Q) as a function of the distance to the nearest central galaxy (D_{cen}) separated into bins of halo mass (as indicated on the legends). The bins are formed from objects that are within a defined D_{cen} range of $\pm 0.15 R_{\text{vir}}$, while the shaded regions are the margins of error for each bin, computed via Poisson statistics. We present these plots for the ‘MOONRISE-Like’ data sets. As previously, Δf_Q measures the offset in quenched fraction for each satellite population relative to a control sample of central galaxies with similar black hole masses. Low-mass satellites are strongly offset to higher quenched fractions at close distances to their centrals, especially within high-mass groups/clusters, at all redshifts. Furthermore, there is a clear anti-correlation between Δf_Q and D_{cen} , as Δf_Q decreases as D_{cen} increases. In low-mass groups however, even low-mass satellites exhibit difficulties in achieving quiescence.

quenching of low-mass satellites occurs most efficiently near the center of these systems. Similar to what is seen for lower mass halos, Δf_Q is clearly inversely correlated with D_{cen} , indicating an increasing quenching efficiency as satellites approach the center of their groups. Furthermore, at a fixed D_{cen} , Δf_Q is much greater for satellites of the most massive groups than those with $\log(M_{\text{Halo}}/M_{\odot}) < 13$. Therefore, we conclude that while the mass of a satellites’ group is most important to its quenching process (as revealed by the random forest analysis of Section 4.3), its location in the group additionally factors in reducing star formation. We find the results for low-mass satellites for EAGLE at $z = 0$ to be consistent with those of TNG.

In summary, the location of satellites within their group or cluster haloes results in a significant secondary quenching trend for low-mass systems within high mass haloes. This operates such that quenching is enhanced as satellites approach closer to their centrals, and hence the center of their groups or clusters. However, for low mass satellites in low mass haloes, the location within the halo is of little impact to quenching. These results are seen primarily from TNG, because EAGLE lacks the low-mass quenching channel accessible to MOONRISE. Hence, there are clear predictions provided here to be tested by future observations at cosmic noon with VLT-MOONRISE.

5 SUMMARY

Over the course of this paper we analyze the quenching of different galaxy classes across cosmic time in two cosmological hydrodynamical simulations: IllustrisTNG and EAGLE. Furthermore, we employ these simulations to extract specific predictions for quenching at cosmic noon in anticipation of VLT-MOONRISE, the first wide-field spectroscopic galaxy survey at intermediate-to-high redshifts (see Cirasuolo et al. 2020; Maiolino et al. 2020, for full details).

We make use of sSFR distributions to determine bespoke quenching thresholds at each redshift for each simulation. Furthermore, with the aim of isolating intrinsic and environmental quenching avenues, we determine the optimal mass thresholds for splitting low and high-mass satellites in terms of quenching by leveraging the bimodality of the stellar mass function of quenched satellites. This approach is

vindicated by the results from the Random Forest analyses, whereby we show that this method is highly effective at separating intrinsically from environmentally quenched galaxies. The final mass thresholds for each simulation and redshift are presented in Table 3.

To investigate the evolution of different quenching avenues between $z = 0$ and $z = 2$, we split galaxies into three classes: (i) central galaxies, (ii) high-mass satellites, and (iii) low-mass satellites. We leverage the power of machine learning, specifically the Random Forest classification approach, to determine the dominant parameter in predicting quenching for each galaxy class at $z = 0 - 2$, for both simulations. This technique enables us to carefully control for nuisance variables before establishing the underlying causal relationships at play (for further discussion see, Bluck et al. 2022). Additionally, we perform more traditional experiments, including probing the quenched fraction of centrals and satellites in different situations as a function of various parameters.

Our primary results, and predictions for VLT-MOONRISE, are as follows:

(i) TNG predicts the existence of environmentally quenched low-mass satellites at cosmic noon, which will be visible to the survey limits of VLT-MOONRISE ($\log(M_*/M_{\odot}) \geq 9.5$), as seen in Fig. 2. Conversely, EAGLE predicts that low-mass environmental quenching will be nearly absent at this mass completeness limit. Hence, the identification and abundance of environmentally quenched objects in MOONRISE (or lack thereof) will provide an excellent comparative test of these two galaxy evolution models.

(ii) As demonstrated in Figs. 3 & 4, satellites exhibit a strong bimodality in over-density, for the complete, 6D phase space data sets of both simulations, especially at $z = 0$. While this bimodality weakens as redshift increases it can still be detected. Moreover, this bimodality in density clearly maps onto a separation in parent halo mass, whereby the high density peak corresponds to high-mass haloes and the low density peak corresponds to low-mass haloes. This highlights δ as an effective tracer for halo mass, at least in

simulation space.

(iii) However, as seen in the lower row of plots in Figs. 3 & 4, this bimodality is much weaker for the observation-like (2D+z observer space) ‘MOONRISE-Like’ data sets. In fact, for both simulations, the distribution appears nearly unimodal for $z \geq 1$. Furthermore, the range of over-density is greatly diminished when compared to that obtained for the complete data sets. This loss of information moving from the simulation space to an observation-like space is highly likely to impact the performance of δ as a tracer for environment in observational surveys, and serves to emphasize the need for accurate measurements of halo mass (e.g., Bluck et al. 2025, in press).

(iv) Both TNG and EAGLE predict that central galaxies quench via intrinsic mechanisms at all redshifts, specifically via a clear dependence on supermassive black hole mass. This is clearly indicated by the dominant predictive power of black hole mass in the Random Forest classification analyses for this galaxy class, presented in Figs. 5 & 6. Hence, there is complete agreement between TNG and EAGLE as to the underlying mechanism of central galaxy quenching at cosmic noon (see also Bluck et al. 2022, 2023, 2024, for similar results). This is a clear testable prediction for VLT-MOONRISE, in which black hole mass may be estimated via the $M_{BH} - \sigma$ relation.

(v) However, an analysis of the evolution of the quenched fraction across cosmic time reveals EAGLE to have a much more effective AGN feedback prescription at higher redshifts than at lower redshifts, whereby high-mass galaxies are more frequently quenched at early cosmic times than at later epochs, see Fig. 10. This is a clear sign of rejuvenation, likely a consequence of the absence of a bespoke AGN ‘maintenance’ mode in EAGLE. Conversely, TNG is equally effective at quenching high-mass systems across all cosmic times, presented in Fig. 9. This is expected due to its ‘kinetic/radio’ mode feedback, which is most effective at low Eddington ratios, and hence within high mass black holes at late cosmic times. Ultimately, these discrepancies between TNG and EAGLE yield a key method to test which simulation better approximates the real Universe. If there is significant rejuvenation, EAGLE performs better; yet, if there is no (or little) rejuvenation, TNG performs better.

(vi) Both TNG and EAGLE predict high-mass satellites to quench via intrinsic mechanisms across all redshifts, specifically connected with black hole mass, with a weak secondary dependence on environmental parameters (which is completely absent for centrals), seen in Figs. 5 & 6. In both simulations, this secondary dependence weakens with redshift. This clearly identifies AGN feedback as the driving quenching process for high-mass satellites in both simulations, aided by environmental effects such as starvation in some cases. Again this is a clear testable prediction for MOONRISE at intermediate-to-high redshifts.

(vii) Both TNG and EAGLE predict low-mass satellites to quench due to environment, with halo mass found to be the best parameter to determine quiescence at all redshifts, as presented in Figs. 5 & 6. Of note, secondary dependencies on D_{cen} and δ are found when performing the Random Forest analysis.

(viii) While several studies have supported a central - satellite decomposition as sufficient for separating intrinsic and environmental quenching, the Random Forest classification results (seen in Figs. 5 & 6) alongside the numerous quenched fraction analyses, point to stellar mass as a better tool to isolate quenching avenues.

This can also be seen from the total quenched population stellar mass functions, which present with a strong bimodality for both simulations at all epochs, shown in Figs. 2a & 2b. This bimodality is present for the quenched satellite population but absent from the quenched central population. As such, a low-mass quenching mechanism must be accessible to satellites, but not centrals. Crucially, both avenues are open to high-mass satellites, but intrinsic quenching dominates.

For both simulations, the low-mass satellite quenched fraction as a function of location within a halo, seen in Fig. 13, and as a function of stellar mass for low and high-mass groups, presented in 11 & 12, reinforce the secondary dependencies on D_{cen} and δ found from the Random Forest analyses. Taken together, this points towards low-mass satellite quenching being most effective towards the center of high mass groups and clusters. This indicates that mechanisms such as ram pressure stripping, galaxy-galaxy interaction, and starvation are the likely candidates responsible for this low-mass quenching. However, while this trend remains essentially constant across redshift for TNG, the impact of environment weakens significantly with increasing redshift in EAGLE.

Conversely, high-mass satellites are not more frequently quiescent as one moves towards the center of even the most massive groups and clusters. This finding underscores the intrinsic nature of the quenching of high mass satellites. As the first wide-field spectroscopic survey at intermediate and high redshift, VLT-MOONRISE will be able to test the impact of environment on quenching from $z \approx 1$ to 2.5, discovering whether these simulations accurately predict the quenching process, or not.

(ix) In the absence of M_{BH} and M_{Halo} as parameters (which are often lacking in observational analyses), the RF analysis of the ‘MOONRISE-Like’ data point to δ as the best predictor of quiescence for high-mass satellites at $z = 0$ (and $z = 1$ for EAGLE), see Figs. 7 & 8. This results in a misidentification of environmental quenching in these cases, since the true parameter regulating quenching is actually supermassive black hole mass in these galaxies. This may be explained by the morphology - density relation, whereby at a fixed stellar mass galaxies in higher density environments tend to be more bulge-dominated and hence are likely to host more massive central black holes (e.g., Saglia et al. 2016; Bluck et al. 2022). Ultimately, this highlights the need for supermassive black hole mass estimates, particularly at $z < 2$, when investigating galaxy quenching in both central and satellite galaxies. Fortunately, this will be possible with VLT-MOONRISE, enabling rigorous tests to these predictions from simulations.

In summary, we find EAGLE and TNG to agree on the quenching mechanisms of all galaxy classes from cosmic noon to the present epoch, which points towards a theoretical consensus. Centrals and high-mass satellites are predicted to quench via AGN feedback, although environment plays a small secondary role in the latter class. Conversely, low-mass satellites are predicted to quench via purely environmental means, with ram pressure stripping, galaxy-galaxy interactions, and starvation working in tandem as the leading candidates for the mechanism. These are all testable predictions which can be verified via VLT-MOONRISE.

In addition to the important similarities between EAGLE and TNG outlined above, there are also important differences between the two simulations with respect to quenching. Most notably, the mass thresholds determined for environmental quenching of satellites are very different between the two simulations. In fact, EAGLE predicts that

the mass completeness limit of MOONRISE is too high to detect environmentally quenched satellites. Alternatively, TNG predicts that this quenching avenue should be clearly visible in this survey. As such, the abundance (or even the presence or absence) of low-mass environmentally quenched satellites in the MOONRISE survey will provide strong evidence for or against the TNG and EAGLE models.

In conclusion, while TNG and EAGLE agree in large part on the fundamental quenching mechanisms responsible in different galaxy classes, they disagree on the details. In combination with the VLT-MOONRISE survey, this will enable the astrophysics community to rigorously test these galaxy evolution models in an unprecedented parameter space, clearly identifying the strengths and weaknesses of each approach. From this, we will be in an excellent position to improve on the current status quo of the theory of galaxy evolution.

6 ACKNOWLEDGMENTS

We are very grateful to the anonymous reviewer for many insightful and helpful comments on this work, which have helped to significantly improve the presentation of our results. AFLB gratefully acknowledges support from an NSF research grant: NSF-AST 2408009, in addition to the ORAU Ralph E. Powe Junior Faculty Enhancement Award in Physical Sciences, and research start-up funds from FIU. PG also acknowledges support from NSF-AST 2408009. RM acknowledges support from the Science and Technology Facilities Council (STFC), ERC Advanced Grant 695671 “QUENCH”, and by the UKRI Frontier Research grant RISEandFALL. RM also acknowledges funding from a research professorship from the Royal Society. PT acknowledges support from NSF-AST 2346977 and the NSF-Simons AI Institute for Cosmic Origins which is supported by the National Science Foundation under Cooperative Agreement 2421782 and the Simons Foundation award MPS-AI-00010515.

7 DATA AVAILABILITY

All data used in this study have been previously published and are available at the following online locations:

- EAGLE: <http://icc.dur.ac.uk/Eagle/>
- TNG: <https://www.tng-project.org/>

REFERENCES

Aird J., et al., 2010, *MNRAS*, **401**, 2531
 Aird J., Coil A. L., Georgakakis A., 2018, *MNRAS*, **474**, 1225
 Bahé Y. M., et al., 2017, *MNRAS*, **470**, 4186
 Barsanti S., et al., 2018, *The Astrophysical Journal*, **857**, 71
 Bluck A. F. L., Ellison S. L., Patton D. R., Simard L., Mendel J. T., Teimoorinia H., Moreno J., Starkenburg E., 2014, *arXiv e-prints*, p. arXiv:1412.3862
 Bluck A. F. L., et al., 2016, *MNRAS*, **462**, 2559
 Bluck A. F. L., et al., 2019, *MNRAS*, **485**, 666
 Bluck A. F. L., Maiolino, Roberto Sánchez, Sebastian F Ellison, Sara L Thorp, Mallory D Piotrowska, Joanna M Teimoorinia, Hossen Bundy, Kevin A 2020a, *MNRAS*, **492**, 96
 Bluck A. F. L., et al., 2020b, *MNRAS*, **499**, 230
 Bluck A. F. L., Maiolino R., Brownson S., Conselice C. J., Ellison S. L., Piotrowska J. M., Thorp M. D., 2022, *A&A*, **659**, A160
 Bluck A. F. L., Piotrowska J. M., Maiolino R., 2023, *ApJ*, **944**, 108
 Bluck A. F. L., et al., 2024, Galaxy quenching at the high redshift frontier: A fundamental test of cosmological models in the early universe with JWST-CEERS, Submitted
 Bondi H., Hoyle F., 1944, *MNRAS*, **104**, 273
 Booth C. M., Schaye J., 2009, *Monthly Notices of the Royal Astronomical Society*, **398**, 53
 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
 Bower R. G., McCarthy I. G., Benson A. J., 2008, *MNRAS*, **390**, 1399
 Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, **351**, 1151
 Brownson S., Bluck A. F. L., Maiolino R., Jones G. C., 2022, *MNRAS*, **511**, 1913
 Bundy K., et al., 2006, *ApJ*, **651**, 120
 Busca N., Balland C., 2018, *arXiv e-prints*, p. arXiv:1808.09955
 Carnall A. C., et al., 2023, *Monthly Notices of the Royal Astronomical Society*, **520**, 3974
 Carnall A. C., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, **534**, 325
 Chabrier G., 2003, *PASP*, **115**, 763
 Cirasuolo M., et al., 2020, *The Messenger*, **180**, 10
 Contini E., Gu Q., Ge X., Rhee J., Yi S. K., Kang X., 2020, *The Astrophysical Journal*, **889**, 156
 Cortese L., Hughes T. M., 2009, *MNRAS*, **400**, 1225
 Crain R. A., et al., 2015, *MNRAS*, **450**, 1937
 Cramer W. J., Kenney J. D. P., Sun M., Crowl H., Yagi M., Jáchym P., Roediger E., Waldron W., 2019, *The Astrophysical Journal*, **870**, 63
 Croton D. J., et al., 2006, *MNRAS*, **365**, 11
 D’Eugenio F., et al., 2025, *ApJS*, **277**, 4
 De Lucia G., Fontanot F., Hirschmann M., Xie L., 2025, *arXiv e-prints*, p. arXiv:2502.01724
 Dekel A., Birnboim Y., 2006, *MNRAS*, **368**, 2
 Domínguez Sánchez H., Huertas-Company M., Bernardi M., Tuccillo D., Fischer J. L., 2018, *MNRAS*, **476**, 3661
 Donnari M., Pillepich A., Nelson D., Marinacci F., Vogelsberger M., Hernquist L., 2021, *Monthly Notices of the Royal Astronomical Society*, **506**, 4760
 Eisenstein D. J., et al., 2023a, *arXiv e-prints*, p. arXiv:2306.02465
 Eisenstein D. J., et al., 2023b, *arXiv e-prints*, p. arXiv:2310.12340
 Fabian A., 2012, *ARA&A*, **50**, 455
 Fabian A. C., Walker S. A., Celotti A., Ghisellini G., Mocz P., Blundell K. M., McMahon R. G., 2014, *MNRAS*, **442**, L81
 Fang Y., Clampitt J., Dalal N., Jain B., Rozo E., Moustakas J., Rykoff E., 2016, *MNRAS*, **463**, 1907
 Finkelstein S. L., et al., 2023, *ApJ*, **946**, L13
 Foltz R., et al., 2018, *The Astrophysical Journal*, **866**, 136
 Gabor J. M., Davé R., 2014, *Monthly Notices of the Royal Astronomical Society*, **447**, 374
 Garilli, B. et al., 2014, *A&A*, **562**, A23
 Genel S., et al., 2014, *MNRAS*, **445**, 175
 Girelli, Giacomo Bolzonella, Micol Cimatti, Andrea 2019, *A&A*, **632**, A80
 Goto T., Yamauchi C., Fujita Y., Okamura S., Sekiguchi M., Smail I., Bernardi M., Gomez P. L., 2003, *Monthly Notices of the Royal Astronomical Society*, **346**, 601
 Goubert P. H., Bluck A. F. L., Piotrowska J. M., Maiolino R., 2024, *Monthly Notices of the Royal Astronomical Society*, **528**, 4891
 Gunn J. E., Gott J. Richard I., 1972, *ApJ*, **176**, 1
 Hamadouche M. L., et al., 2024, *arXiv e-prints*, p. arXiv:2412.09592
 Henriques B. M. B., White S. D. M., Lilly S. J., Bell E. F., Bluck A. F. L., Terrazas B. A., 2019, *MNRAS*, **485**, 3446
 Hoyle F., Lyttleton R. A., 1939, *Proceedings of the Cambridge Philosophical Society*, **35**, 405
 Ilbert, O. et al., 2015, *A&A*, **579**, A2
 Ji Z., Giavalisco M., Williams C. C., Faber S. M., Ferguson H. C., Guo Y., Liu T., Lee B., 2018, *The Astrophysical Journal*, **862**, 135
 Kapferer W., Sluka C., Schindler S., Ferrari C., Ziegler B., 2009, *A&A*, **499**, 87
 Kartaltepe J. S., et al., 2023, *ApJ*, **946**, L15
 Kawinwanichakij L., et al., 2017, *ApJ*, **847**, 134
 Kennicutt Jr. R. C., 1998, *ARA&A*, **36**, 189
 Kocevski D. D., et al., 2023, *ApJ*, **946**, L14

Kukstas E., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 518, 4782

Lilly S. J., et al., 2007, *ApJS*, 172, 70

Lim S., Tacchella S., Maiolino R., Schaye J., Schaller M., 2025, *arXiv e-prints*, p. [arXiv:2504.02027](https://arxiv.org/abs/2504.02027)

Lotz M., Remus R.-S., Dolag K., Biviano A., Burkert A., 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 5370

Madau P., Dickinson M., 2014, *Annual Review of Astronomy and Astrophysics*, 52, 415

Maiolino R., et al., 2012, *MNRAS*, 425L, L66

Maiolino R., et al., 2020, *The Messenger*, 180, 24

Mao Z., Kodama T., Pérez-Martínez J. M., Suzuki T. L., Yamamoto N., Adachi K., 2022, *A&A*, 666, A141

Marasco A., Crain R. A., Schaye J., Bahé Y. M., van der Hulst T., Theuns T., Bower R. G., 2016, *MNRAS*, 461, 2630

Marinacci F., et al., 2018, *MNRAS*, 480, 5113

McAlpine S., et al., 2016, *A&C*, 15, 72

Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, *Nature*, 379, 613

Moore B., Governato F., Quinn T., Stadel J., Lake G., 1998, *ApJ*, 499, L5

Moore B., Lake G., Quinn T., Stadel J., 1999, *MNRAS*, 304, 465

Naiman J. P., et al., 2018, *MNRAS*, 477, 1206

Nelson D., et al., 2018, *MNRAS*, 475, 624

Pedregosa F., et al., 2011a, *Journal of Machine Learning Research*, 12, 2825

Pedregosa F., et al., 2011b, *Journal of Machine Learning Research*, 12, 2825

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

Peng Y., Lilly S. J., Renzini A., Carollo M., 2012, *ApJ*, 757, 4

Pentericci L., et al., 2018, *A&A*, 619, A147

Pillepich A., et al., 2018a, *MNRAS*, 473, 4077

Pillepich A., et al., 2018b, *MNRAS*, 475, 648

Piotrowska J. M., Bluck A. F. L., Maiolino R., Peng Y., 2022, *MNRAS*, 512, 1052

Planck Collaboration et al., 2014, *A&A*, 571, A1

Planck Collaboration et al., 2016, *A&A*, 594, A13

Postman M., Geller M. J., 1984, *ApJ*, 281, 95

Pérez-Carrasco M., Cabrera-Vives G., Martínez-Marin M., Cerulo P., Demarco R., Protopapas P., Godoy J., Huertas-Company M., 2019, *Publications of the Astronomical Society of the Pacific*, 131, 108002

Rhee J., et al., 2024, *The Astrophysical Journal*, 971, 111

Saglia R. P., et al., 2016, *ApJ*, 818, 47

Schaye J., Dalla Vecchia C., 2008, *Monthly Notices of the Royal Astronomical Society*, 383, 1210

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

Schechter P., 1976, *ApJ*, 203, 297

Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877

Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575

Smith M. C., Sijacki D., Shen S., 2018, *MNRAS*, 478, 302

Springel V., 2005, *MNRAS*, 364, 1105

Springel V., Hernquist L., 2003, *Monthly Notices of the Royal Astronomical Society*, 339, 289

Springel V., et al., 2018, *MNRAS*, 475, 676

Strateva I., et al., 2001, *ApJ*, 122, 1861

Teimoorinia H., Bluck A. F. L., Ellison S. L., 2016, *MNRAS*, 457, 2086

Teklu A. F., Remus R., Dolag K., Burkert A., 2017, *Monthly Notices of the Royal Astronomical Society*, 472, 4769

Terrazas B. A., Bell E. F., Henriques B. M. B., White S. D. M., Cattaneo A., Woo J., 2016, *The Astrophysical Journal Letters*, 830, L12

Terrazas B. A., Bell E. F., Woo J., Henriques B. M. B., 2017, *ApJ*, 844, 170

Torrey P., Vogelsberger M., Genel S., Sijacki D., Springel V., Hernquist L., 2014, *MNRAS*, 438, 1985

Vogelsberger M., et al., 2014, *MNRAS*, 444, 1518

Wang H., et al., 2018, *ApJ*, 852, 31

Weinberger R., et al., 2017, *MNRAS*, 465, 3291

Weinberger R., et al., 2018, *MNRAS*, 479, 4056

Williams R. J., Quadri R. F., Franx M., van Dokkum P., Labb   I., 2009, *ApJ*, 691, 1879

Woo J., et al., 2013, *MNRAS*, 428, 3306

Woo J., Dekel A., Faber S. M., Koo D. C., 2015, *MNRAS*, 448, 237

Woo J.-H., Son D., Bae H.-J., 2017, *ApJ*, 839, 120

Wright R. J., Lagos C. d. P., Power C., Stevens A. R. H., Cortese L., Poulton R. J. J., 2022, *Monthly Notices of the Royal Astronomical Society*, 516, 2891

Xie L., De Lucia G., Hirschmann M., Fontanot F., 2020, *MNRAS*, 498, 4327

Xie L., et al., 2024, *ApJ*, 966, L2

Zinger E., et al., 2020, *MNRAS*, 499, 768

de Diego, Jos   A. et al., 2020, *A&A*, 638, A134

van Dokkum P., Quadri R., 2007, in Metcalfe N., Shanks T., eds, *Astronomical Society of the Pacific Conference Series Vol. 379, Cosmic Frontiers*. p. 249 ([arXiv:0704.1760](https://arxiv.org/abs/0704.1760)), doi:10.48550/arXiv.0704.1760

van de Voort F., Bah   Y. M., Bower R. G., Correa C. A., Crain R. A., Schaye J., Theuns T., 2017, *MNRAS*, 466, 3460

van den Bosch F. C., Aquino D., Yang X., Mo H. J., Pasquali A., McIntosh D. H., Weinmann S. M., Kang X., 2008, *MNRAS*, 387, 79

APPENDIX B: SCHECHTER FIT COEFFICIENTS

We include in this appendix (see Tables B1 & B2) the values of the coefficients, along with their associated uncertainties, for the best fit Schechter functions for all populations and redshifts of both simulations, as seen in Fig. 2. The only uncertainties present for the simulations are from Poisson counting statistics, since stellar masses are known precisely. The tables' columns present the median of the following parameters: $\log(M^*)$; ϕ ; α ; ϕ_2 ; and α_2 . The uncertainties are presented as the total 1σ errors.

Here, $\log(M^*)$ corresponds to the ‘knee’ or characteristic mass of the Schechter function. This is to say the mass at which the stellar mass function turns over from a power law to an exponential function.

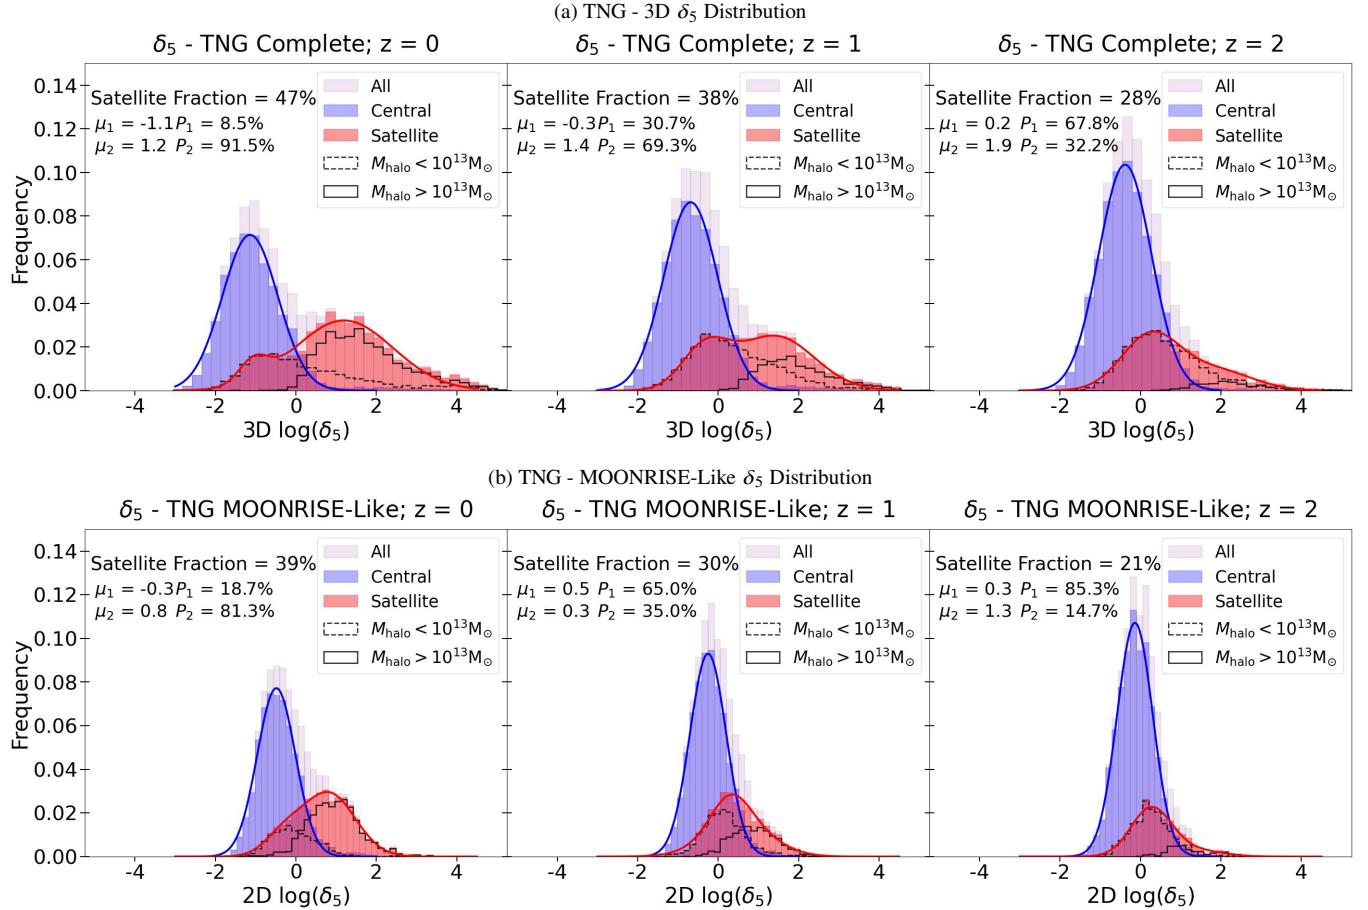


Figure A1. Identical in structure and method to Fig. 3, but here presenting the evolution of the distribution in galaxy over-density of the fifth nearest neighbor density, δ_5 for the TNG simulation.

The normalization, ϕ , corresponds to the number density at the characteristic mass, and α is defined as the slope of the power law which describes the stellar mass function prior to the characteristic mass. For double Schechter functions, the sum of two Schechter functions with the same characteristic mass, there are two additional parameters: ϕ_2 , the number density at the characteristic mass of the second Schechter function; and α_2 , the power law slope corresponding to the second Schechter function.

APPENDIX C: MACHINE LEARNING REPRODUCIBILITY

For all Random Forest classification analyses we make use of the SCIKIT-LEARN Random Forest Classifier architecture (Pedregosa et al. 2011b). In order to achieve optimal performance, the Random Forest classifier allows for the fine tuning of multiple hyper-parameters. We include below a brief description for each of the relevant hyper-parameters (which are listed in Tables C1, C3, C5, & C7 for centrals and Tables C2, C4, C6, & C8 for satellites).

Method: In all cases, we make use of the Random Forest Classifier from SCIKIT-LEARN, abbreviated to RFclass in the Tables.

Normalized: Indicates whether the data is normalized prior to going through the Random Forest classifier. In all cases, we subtract the median and normalize by the inter-quartile range. This ensures

the classifier will not be influenced by the different distributions or value ranges between features.

N_{int}: The number of times a random forest classification is run, altering the data randomly separated into the training and testing sample sets. In this case, we set all to, $N_{int} = 100$.

Balanced: Indicates whether the data is balanced to have an even amount of quiescent and star forming galaxies. While it is best to evenly sample the data prior to each Random Forest classification run, the dearth of quenched objects at higher redshift, especially quiescent satellites for $z = 2$, renders us unable to do so in some cases.

Train : Test: The training to testing ratio prior to running the random forest classifier. For sufficiently large data sets, an even 50:50 split is ideal. However, for smaller amounts of data it may be preferable to allow for a larger training than testing data set (see Bluck et al. 2024). As such, for certain parent sample sets, an uneven training:testing split is applied, prior to the Random Forest classification. In these cases, the training sample comprises up to 70% of the total data set, and the testing data set is therefore the remaining 30%.

Max Depth: The maximum amount of decision splits which can occur within any given decision tree. This is set to *Max Depth* = 200 for all of our Random Forest runs. This always exceeds the limit set

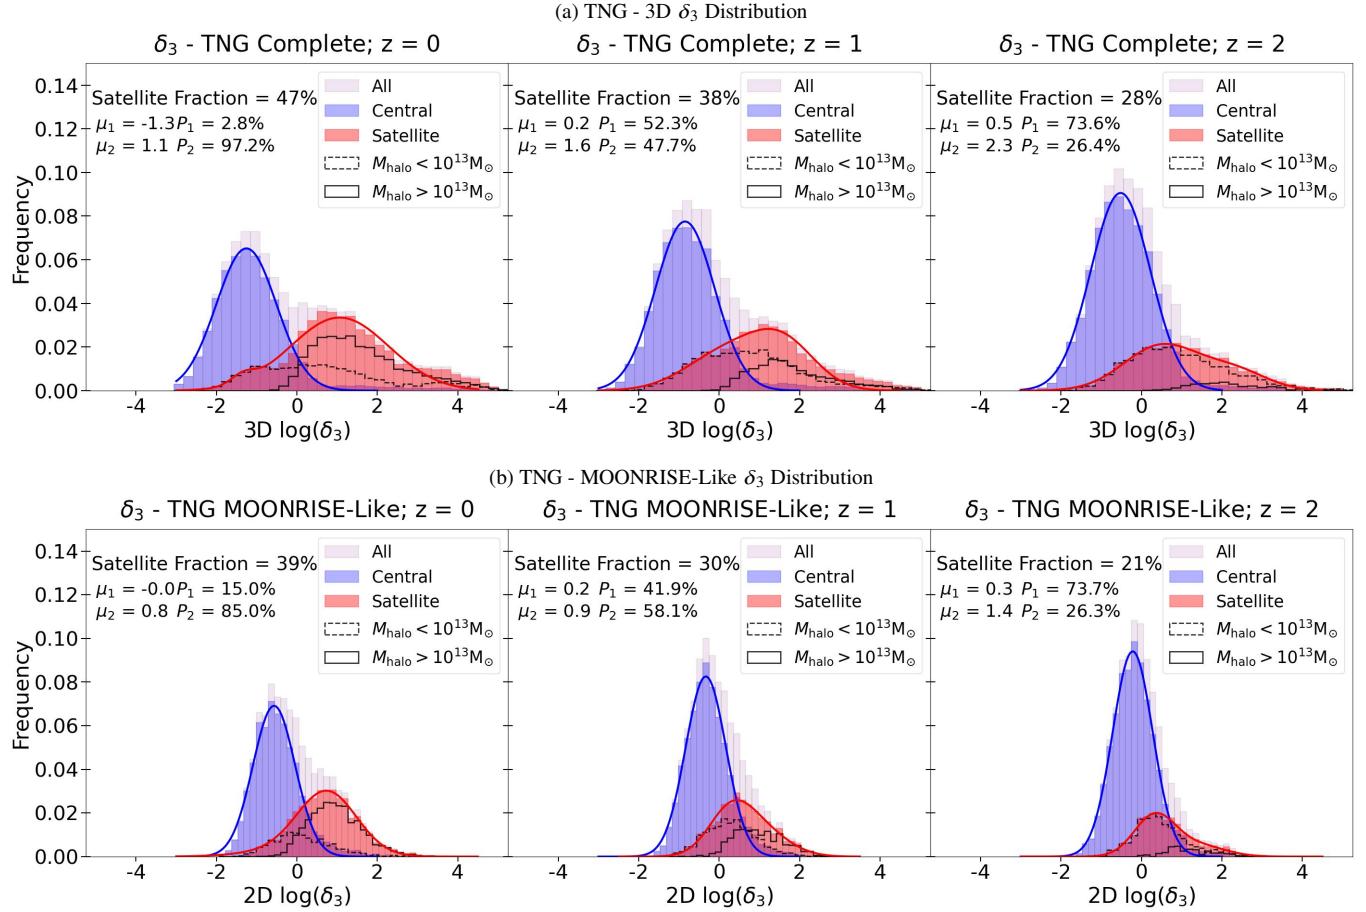


Figure A2. Identical in structure and method to Fig. 3, but here presenting the evolution of the distribution of galaxy over-density for the third nearest neighbor density, δ_3 , for the TNG simulation.

by the *min samples leaf* value (MSL, see below).

min-samples-leaf - MSL: The minimum amount of data remaining in a node needed to attempt further splits. Although setting the MSL to its logical minimum of two seems intuitive, this would in general lead to overfitting. That is, while the training data would be optimally classified, the Random Forest algorithm would likely learn pathological attributes within the training data which do not scale to unseen testing data. To prevent this, we fine tune the MSL value to simultaneously maximize the AUC (further discussed below), which quantifies the training performance, and minimize the ΔAUC (see below), which quantifies the difference between the performance of the RF on the training and unseen data.

Random State: The random state informs the random bootstrapping of data. Specifying the value of the random state assures the bootstrapping to be consistent through both iterations, and should enable other researchers to reproduce the results precisely, even given their dependence on ‘random’ variables.

AUC: The area under true positive - false positive receiver operator curve (see, e.g., Teimoorinia et al. 2016). The AUC reveals the performance of the Random Forest classification, with an AUC = 1 being a perfect classification and AUC = 0.5 indicating an entirely random performance. We include the mean training data AUC from

the 100 iterations as the AUC column within each table.

ΔAUC : Quantifies the difference in the performances of the Random Forest classifier on the training and testing data sets. We attempt to minimize the ΔAUC for all RF runs. While a threshold of $\Delta\text{AUC} \leq 0.02$ is optimal as according to Bluck et al. (2022), we accept all tests with $|\Delta\text{AUC}| \leq 0.07$ as a successful fit here. This change is due mainly to the sparsity of data available from certain parent sample sets.

To assure our RF classification results as reproducible as possible, we report information on said hyperparameters used for each RF test in Tables C1 - C8. Additionally, we set *max-features* = ‘None’ for all classifications, so as to insure that all features are evaluated at every node of all decision trees of the Random Forest. As shown in Bluck et al. (2022), this is the optimal method available in RF analyses for determining the most causally connected feature for classification, enabling controlling for all other available features.

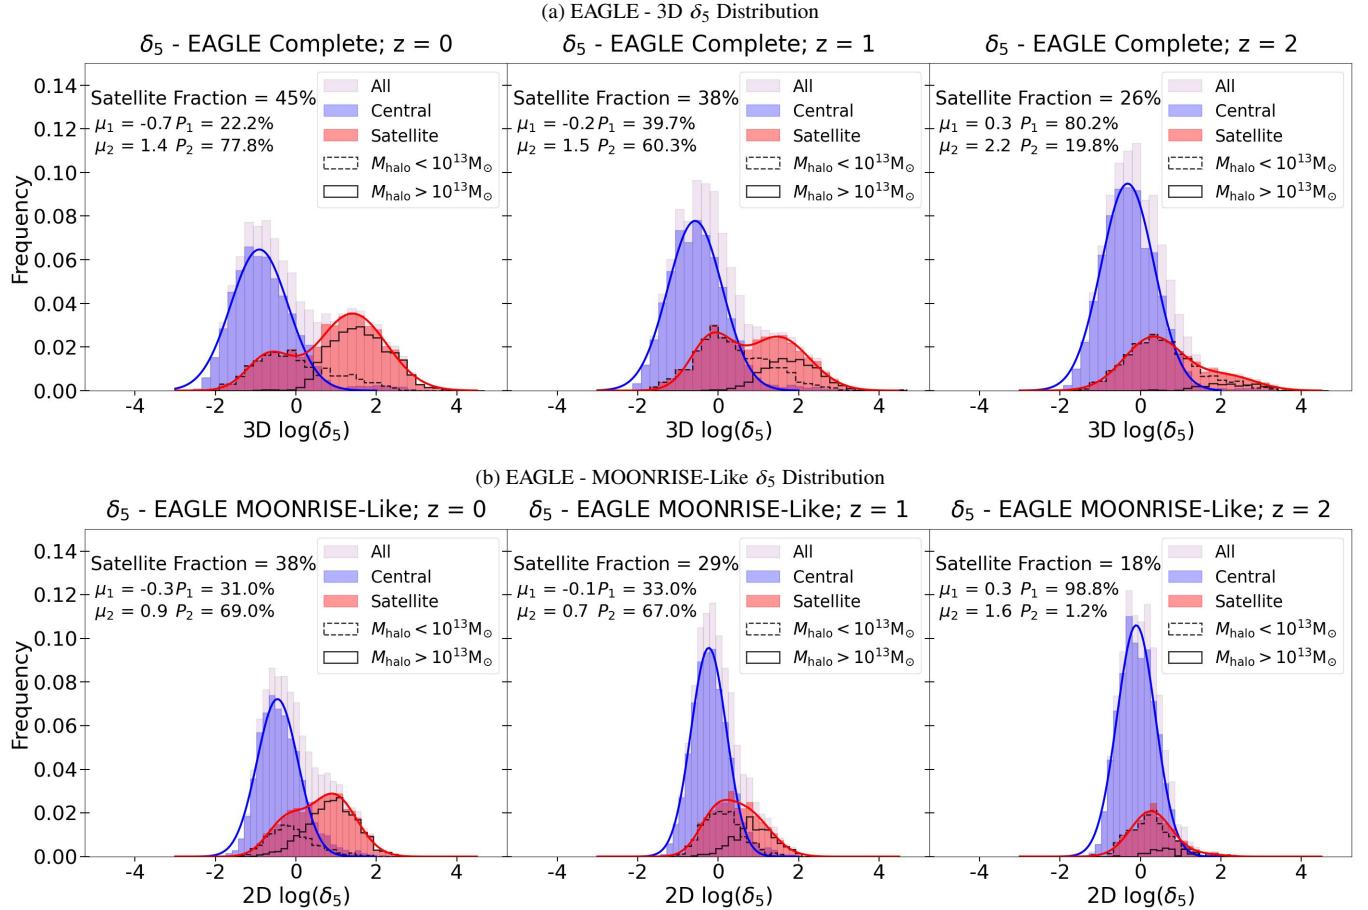


Figure A3. Identical in structure and method to Fig. 4, but here presenting the evolution of the distribution in galaxy over-density of the fifth nearest neighbor density, δ_5 for the EAGLE simulation.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.

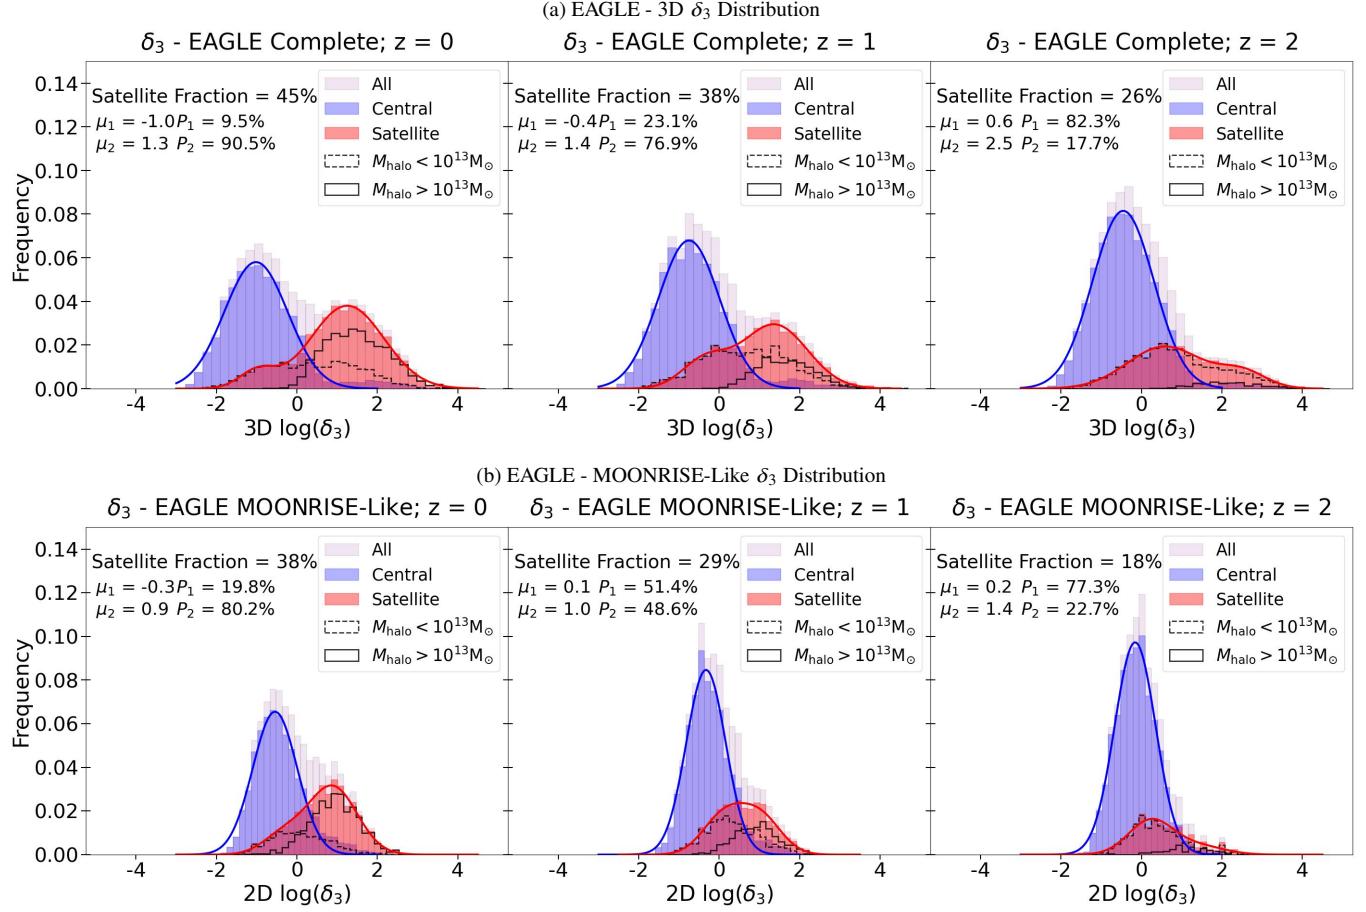


Figure A4. Identical in structure and method to Fig. 4, but here presenting the evolution of the distribution of galaxy over-density for the third nearest neighbor density, δ_3 , for the EAGLE simulation.

Table B1. The best-fit Schechter function parameters in TNG. Each row corresponds to a sample subset for the TNG simulation at each redshift, while the columns list the parameters along with their associated uncertainties.

Sample Subset	$\log(M^*)/\text{M}_\odot$	$\phi [10^{-3}\text{Mpc}^{-3}]$	α	$\phi_2 [10^{-3}\text{Mpc}^{-3}]$	α_2
TNG - All Galaxies; $z = 0$	$11.22^{+0.02}_{-0.01}$	$1.911^{+0.035}_{-0.049}$	$-1.359^{+0.005}_{-0.006}$
TNG - All Star Forming; $z = 0$	$10.55^{+0.01}_{-0.02}$	$2.428^{+0.111}_{-0.094}$	$-1.221^{+0.014}_{-0.009}$
TNG - All Quenched; $z = 0$	$10.48^{+0.03}_{-0.03}$	$1.931^{+0.103}_{-0.103}$	$0.834^{+0.119}_{-0.109}$	$0.301^{+0.039}_{-0.030}$	$-1.597^{+0.028}_{-0.026}$
TNG - Quenched Satellites; $z = 0$	$10.45^{+0.05}_{-0.05}$	$0.905^{+0.056}_{-0.080}$	$0.669^{+0.222}_{-0.203}$	$0.333^{+0.061}_{-0.052}$	$-1.574^{+0.037}_{-0.041}$
TNG - Quenched Centrals; $z = 0$	$10.51^{+0.02}_{-0.03}$	$1.067^{+0.041}_{-0.064}$	$0.942^{+0.125}_{-0.072}$
TNG - All Galaxies; $z = 1$	$11.12^{+0.02}_{-0.02}$	$1.129^{+0.058}_{-0.041}$	$-1.354^{+0.008}_{-0.007}$
TNG - All Star Forming; $z = 1$	$10.76^{+0.02}_{-0.02}$	$1.679^{+0.068}_{-0.080}$	$-1.297^{+0.013}_{-0.014}$
TNG - All Quenched; $z = 1$	$10.46^{+0.03}_{-0.03}$	$0.624^{+0.058}_{-0.060}$	$1.316^{+0.117}_{-0.131}$	$0.129^{+0.0185}_{-0.0137}$	$-1.517^{+0.034}_{-0.034}$
TNG - Quenched Satellites; $z = 1$	$10.27^{+0.05}_{-0.03}$	$0.153^{+0.044}_{-0.023}$	$1.825^{+0.175}_{-0.311}$	$0.209^{+0.020}_{-0.034}$	$-1.424^{+0.036}_{-0.050}$
TNG - Quenched Centrals; $z = 1$	$10.58^{+0.07}_{-0.07}$	$0.497^{+0.079}_{-0.133}$	$0.945^{+0.402}_{-0.238}$
TNG - All Galaxies; $z = 2$	$11.14^{+0.03}_{-0.03}$	$0.546^{+0.032}_{-0.034}$	$-1.449^{+0.009}_{-0.008}$
TNG - All Star Forming; $z = 2$	$10.85^{+0.02}_{-0.02}$	$0.813^{+0.048}_{-0.050}$	$-1.414^{+0.010}_{-0.012}$
TNG - All Quenched; $z = 2$	$10.43^{+0.04}_{-0.04}$	$0.268^{+0.044}_{-0.043}$	$1.54^{+0.203}_{-0.212}$	$0.023^{+0.007}_{-0.005}$	$-1.706^{+0.077}_{-0.0725}$
TNG - Quenched Satellites; $z = 2$	$10.31^{+0.07}_{-0.08}$	$0.077^{+0.017}_{-0.031}$	$1.337^{+0.637}_{-0.337}$	$0.032^{+0.015}_{-0.008}$	$-1.653^{+0.083}_{-0.095}$
TNG - Quenched Centrals; $z = 2$	$10.51^{+0.04}_{-0.06}$	$0.235^{+0.038}_{-0.038}$	$1.367^{+0.256}_{-0.202}$

Table B2. The best-fit Schechter function parameters in EAGLE. Each row corresponds to a sample subset for the EAGLE simulation at each redshift, while the columns list the parameters along with their associated uncertainties.

Sample Subset	$\log(M^*)/\text{M}_\odot$	$\phi[10^{-3}\text{Mpc}^{-3}]$	α	$\phi_2[10^{-3}\text{Mpc}^{-3}]$	α_2
EAGLE - All Galaxies; $z = 0$	$11.03^{+0.05}_{-0.04}$	$1.246^{+0.121}_{-0.123}$	$-1.337^{+0.018}_{-0.014}$
EAGLE - All Star Forming; $z = 0$	$10.74^{+0.04}_{-0.04}$	$1.589^{+0.143}_{-0.135}$	$-1.236^{+0.020}_{-0.022}$
EAGLE - All Quenched; $z = 0$	$10.98^{+0.16}_{-0.15}$	$0.592^{+0.200}_{-0.175}$	$-0.904^{+0.338}_{-0.238}$	$0.011^{+0.041}_{-0.010}$	$-2.159^{+0.319}_{-0.573}$
EAGLE - Quenched Satellites; $z = 0$	$10.60^{+0.15}_{-0.53}$	$0.505^{+0.68}_{-0.108}$	$-0.709^{+1.38}_{-0.291}$	$0.0349^{+0.197}_{-0.024}$	$-1.911^{+0.384}_{-0.490}$
EAGLE - Quenched Centrals; $z = 0$	$11.24^{+0.06}_{-0.05}$	$0.255^{+0.020}_{-0.025}$	$-0.794^{+0.031}_{-0.037}$
EAGLE - All Galaxies; $z = 1$	$11.12^{+0.03}_{-0.04}$	$0.761^{+0.085}_{-0.065}$	$-1.452^{+0.015}_{-0.012}$
EAGLE - All Star Forming; $z = 1$	$10.84^{+0.02}_{-0.03}$	$1.029^{+0.082}_{-0.067}$	$-1.429^{+0.015}_{-0.012}$
EAGLE - All Quenched; $z = 1$	$10.44^{+0.04}_{-0.04}$	$0.861^{+0.023}_{-0.020}$	$0.121^{+0.101}_{-0.092}$	$0.003^{+0.002}_{-0.001}$	$-2.705^{+0.152}_{-0.129}$
EAGLE - Quenched Satellites; $z = 1$	$10.24^{+0.06}_{-0.11}$	$0.364^{+0.026}_{-0.054}$	$0.375^{+0.322}_{-0.260}$	$0.023^{+0.021}_{-0.010}$	$-2.263^{+0.175}_{-0.179}$
EAGLE - Quenched Centrals; $z = 1$	$10.62^{+0.03}_{-0.04}$	$0.484^{+0.017}_{-0.018}$	$-0.015^{+0.092}_{-0.071}$
EAGLE - All Galaxies; $z = 2$	$11.09^{+0.07}_{-0.07}$	$0.318^{+0.049}_{-0.056}$	$-1.585^{+0.016}_{-0.021}$
EAGLE - All Star Forming; $z = 2$	$10.80^{+0.10}_{-0.08}$	$0.445^{+0.106}_{-0.095}$	$-1.592^{+0.033}_{-0.030}$
EAGLE - All Quenched; $z = 2$	$10.12^{+0.07}_{-0.08}$	$0.510^{+0.041}_{-0.052}$	$0.646^{+0.264}_{-0.204}$	$0.0006^{+0.0013}_{-0.0004}$	$-3.454^{+0.484}_{-0.533}$
EAGLE - Quenched Satellites; $z = 2$	$9.82^{+0.12}_{-0.07}$	$0.111^{+0.040}_{-0.025}$	$1.182^{+0.285}_{-0.558}$	$0.020^{+0.010}_{-0.012}$	$-2.232^{+0.193}_{-0.320}$
EAGLE - Quenched Centrals; $z = 2$	$10.25^{+0.07}_{-0.06}$	$0.390^{+0.019}_{-0.015}$	$0.429^{+0.188}_{-0.172}$

Table C1. Random Forest hyper-parameters for the classifications tests performed for centrals of the complete data sets, presented in Figs. 5 & 6

Sample Set - Centrals	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	100	42	0.98	0.008
TNG; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	35	42	0.98	0.014
TNG; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	18	42	0.99	0.012
EAGLE; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	95	42	0.84	0.022
EAGLE; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	25	42	0.97	0.028
EAGLE; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	35	42	0.96	0.027

Table C2. Random Forest hyper-parameters for the classifications tests performed for satellites ('HM' & 'LM' are used for high and low-mass) of the complete data sets, presented in Figs. 5 & 6

Sample Set - Satellites	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG HM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	51	42	0.84	0.032
TNG HM; $z = 1$	60/40	Yes	Yes	RFclass	100	200	500	64	42	0.88	0.030
TNG HM; $z = 2$	70/30	Yes	Yes	RFclass	100	200	500	31	42	0.90	0.041
TNG LM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	300	42	0.87	0.007
TNG LM; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	100	42	0.85	0.017
TNG LM; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	47	42	0.89	0.026
EAGLE HM; $z = 0$	70/30	Yes	Yes	RFclass	100	200	500	63	42	0.78	0.047
EAGLE HM; $z = 1$	60/40	Yes	Yes	RFclass	100	200	500	52	42	0.85	0.038
EAGLE HM; $z = 2$	50/50	No	Yes	RFclass	100	200	500	50	42	0.90	0.028
EAGLE LM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	200	42	0.88	0.009
EAGLE LM; $z = 1$	70/30	Yes	Yes	RFclass	100	200	500	85	42	0.88	0.028
EAGLE LM; $z = 2$	70/30	No	Yes	RFclass	100	200	500	81	42	0.83	0.057

Table C3. Random Forest hyper-parameters for the classifications tests performed for centrals of the 2D-complete data sets, presented in Figs. 5 & 6

2D Sample Set - Centrals	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	100	42	0.97	0.006
TNG; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	50	42	0.98	0.015
TNG; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	15	42	0.99	0.011
EAGLE; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	100	42	0.84	0.024
EAGLE; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	50	42	0.96	0.027
EAGLE; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	32	42	0.96	0.031

Table C4. Random Forest hyper-parameters for the classifications tests performed for satellites ('HM' & 'LM' are used for high and low-mass) of the 2D-complete data sets, presented in Figs. 5 & 6

2D Sample Set - Satellites	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG HM; $z = 0$	60/40	Yes	Yes	RFclass	100	200	500	54	42	0.85	0.024
TNG HM; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	42	42	0.89	0.031
TNG HM; $z = 2$	60/40	No	Yes	RFclass	100	200	500	40	42	0.89	0.031
TNG LM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	220	42	0.87	0.007
TNG LM; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	170	42	0.84	0.010
TNG LM; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	52	42	0.88	0.027
EAGLE HM; $z = 0$	70/30	Yes	Yes	RFclass	100	200	500	64	42	0.79	0.05
EAGLE HM; $z = 1$	60/40	Yes	Yes	RFclass	100	200	500	52	42	0.86	0.036
EAGLE HM; $z = 2$	70/30	No	Yes	RFclass	100	200	500	50	42	0.89	0.039
EAGLE LM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	180	42	0.88	0.011
EAGLE LM; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	65	42	0.87	0.021
EAGLE LM; $z = 2$	70/30	No	Yes	RFclass	100	200	500	78	42	0.84	0.05

Table C5. Random Forest hyper-parameters for the classifications tests performed for centrals of the MOONRISE-Like (ML) data sets, presented in Figs. 5 & 6

ML Sample Set - Centrals	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	70	42	0.96	0.013
TNG; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	40	42	0.97	0.02
TNG; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	15	42	0.99	0.015
EAGLE; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	65	42	0.85	0.043
EAGLE; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	25	42	0.96	0.031
EAGLE; $z = 2$	50/50	No	Yes	RFclass	100	200	500	70	42	0.93	0.012

Table C6. Random Forest hyper-parameters for the classifications tests performed for satellites of the MOONRISE-Like (ML) data sets, presented in Figs. 5 & 6

ML Sample Set - Satellites	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG HM; $z = 0$	70/30	Yes	Yes	RFclass	100	200	500	44	42	0.83	0.024
TNG HM; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	34	42	0.87	0.040
TNG HM; $z = 2$	60/40	No	Yes	RFclass	100	200	500	20	42	0.90	0.052
TNG LM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	100	42	0.86	0.022
TNG LM; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	48	42	0.83	0.032
TNG LM; $z = 2$	70/30	No	Yes	RFclass	100	200	500	40	42	0.91	0.034
EAGLE HM; $z = 0$	70/30	Yes	Yes	RFclass	100	200	500	38	42	0.82	0.055
EAGLE HM; $z = 1$	70/30	Yes	Yes	RFclass	100	200	500	40	42	0.86	0.053
EAGLE HM; $z = 2$	60/40	No	Yes	RFclass	100	200	500	34	42	0.86	0.054
EAGLE LM; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	90	42	0.86	0.021
EAGLE ML - LM Sat; $z = 1$
EAGLE ML - LM Sat; $z = 2$

Table C7. Random Forest hyper-parameters for the classifications tests performed for centrals of the observation MOONRISE-Like (ML_{obs}) data sets, presented in Figs. 7 & 8

ML _{obs}	Sample Set - Centrals	Train:Test	Balanced	Normalized	Method	N_{int}	N_{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
	TNG; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	120	42	0.91	0.010
	TNG; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	40	42	0.92	0.022
	TNG; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	15	42	0.98	0.025
	EAGLE; $z = 0$	50/50	Yes	Yes	RFclass	100	200	500	70	42	0.76	0.031
	EAGLE; $z = 1$	50/50	Yes	Yes	RFclass	100	200	500	60	42	0.84	0.032
	EAGLE; $z = 2$	50/50	Yes	Yes	RFclass	100	200	500	40	42	0.84	0.041

Table C8. Random Forest hyper-parameters for the classifications tests performed for satellites of the observation MOONRISE-Like (ML_{obs}) data sets, presented in Figs. 7 & 8

ML _{obs}	Sample Set - Satellites	Train:Test	Balanced	Normalized	Method	N _{int}	N _{trees}	Max Depth	MSL	Random State	AUC	Δ AUC
TNG HM; z = 0	60/40	No	Yes	RFclass	100	200	500	66	42	0.74	0.054	
TNG HM; z = 1	70/30	Yes	Yes	RFclass	100	200	500	58	42	0.75	0.054	
TNG HM; z = 2	60/40	No	Yes	RFclass	100	200	500	36	42	0.75	0.058	
TNG LM; z = 0	50/50	Yes	Yes	RFclass	100	200	500	95	42	0.81	0.021	
TNG LM; z = 1	70/30	Yes	Yes	RFclass	100	200	500	49	42	0.80	0.048	
TNG LM; z = 2	70/30	No	Yes	RFclass	100	200	500	46	42	0.81	0.048	
EAGLE HM; z = 0	70/30	Yes	Yes	RFclass	100	200	500	40	42	0.72	0.062	
EAGLE HM; z = 1	60/40	No	Yes	RFclass	100	200	500	49	42	0.74	0.07	
EAGLE HM; z = 2	70/30	No	Yes	RFclass	100	200	500	46	42	0.71	0.073	
EAGLE LM; z = 0	50/50	Yes	Yes	RFclass	100	200	500	70	42	0.81	0.027	
EAGLE LM; z = 1	
EAGLE LM; z = 2	