














Galaxy Quenching at the High Redshift Frontier: A Fundamental Test of Cosmological Models in the Early Universe with JWST-CEERS

Asa F. L. Bluck¹ , Christopher J. Conselice² , Katherine Ormerod² , Joanna M. Piotrowska³ , Nathan Adams² ,
Duncan Austin² , Joseph Caruana^{4,5} , K. J. Duncan⁶ , Leonardo Ferreira⁷ , Paul Goubert¹, Thomas Harvey² ,
James Trussler², and Roberto Maiolino^{8,9} 

¹ Stocker AstroScience Center, Dept. of Physics, Florida International University, 11200 S.W. 8th Street, Miami, FL 33199, USA; abluck@fiu.edu

² Jodrell Bank Centre for Astrophysics, Dept. of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL UK

³ Department of Astronomy, California Institute of Technology, 1200 E. California Boulevard, Pasadena, CA 91125, USA

⁴ Department of Physics, University of Malta, Msida MSD 2080, Malta

⁵ Institute of Space Sciences & Astronomy, University of Malta, Msida MSD 2080, Malta

⁶ Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh, EH9 3HJ, UK

⁷ Centre for Astronomy and Particle Theory, University of Nottingham, Nottingham, UK

⁸ Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

⁹ Cavendish Laboratory Astrophysics Group, University of Cambridge, 19 JJ Thomson Avenue, Cambridge, CB3 0HE, UK

Received 2023 September 4; revised 2023 November 2; accepted 2023 November 4; published 2024 January 23

Abstract

We present an analysis of the quenching of star formation in massive galaxies ($M_* > 10^{9.5} M_\odot$) within the first 0.5–3 Gyr of the Universe’s history utilizing JWST-CEERS data. We utilize a combination of advanced statistical methods to accurately constrain the intrinsic dependence of quenching in a multidimensional and intercorrelated parameter space. Specifically, we apply random forest classification, area statistics, and a partial correlation analysis to the JWST-CEERS data. First, we identify the key testable predictions from two state-of-the-art cosmological simulations (IllustrisTNG and EAGLE). Both simulations predict that quenching should be regulated by supermassive black hole mass in the early Universe. Furthermore, both simulations identify the stellar potential (ϕ_*) as the optimal proxy for black hole mass in photometric data. In photometric observations, where we have no direct constraints on black hole masses, we find that the stellar potential is the most predictive parameter of massive galaxy quenching at all epochs from $z = 0$ –8, exactly as predicted by simulations for this sample. The stellar potential outperforms stellar mass, galaxy size, galaxy density, and Sérsic index as a predictor of quiescence at all epochs probed in JWST-CEERS. Collectively, these results strongly imply a stable quenching mechanism operating throughout cosmic history, which is closely connected to the central gravitational potential in galaxies. This connection is explained in cosmological models via massive black holes forming and growing in deep potential wells, and subsequently quenching galaxies through a mix of ejective and preventative active galactic nucleus feedback.

Unified Astronomy Thesaurus concepts: [Galaxy evolution \(594\)](#); [Galaxy formation \(595\)](#)

1. Introduction

The existence of quenched (quiescent, non-star-forming) galaxies remains a significant puzzle in the field of extragalactic astrophysics. The reason for this is that galaxies cease forming stars long before the baryon content of their dark matter haloes is exhausted. Indeed, only $\sim 5\%$ – 10% of baryons collapse into stars by the present epoch (Fukugita & Peebles 2004; Shull et al. 2012). Yet, simple models of galaxy formation (utilizing only cosmological expansion, gravitation, and cooling processes) unanimously predict that the vast majority of baryons should have been collated into stars and stellar remnants by $z = 0$ (e.g., Cole et al. 2000; Bower et al. 2006; Croton et al. 2006; Bower et al. 2008; Henriques et al. 2013, 2015; Somerville & Dave 2015). Closely related to this problem is the question as to why the majority of baryons in high-mass galaxies, groups, and clusters reside in a hot gas halo ($T_{\text{halo}} \sim 10^{7-8} \text{K}$), which despite being embedded in a $\sim 3(1+z) \text{K}$ ambient medium (the cosmic microwave background), is

stable from cooling and collapse for *billions of years* (e.g., Fabian 1994, 1999, 2012; McNamara et al. 2007; Voit et al. 2002; McNamara & Nulsen 2007; Hlavacek-Larrondo et al. 2012, 2015, 2018).

Typically, theorists think about these problems in terms of cosmological star formation efficiency, i.e., the fraction of baryons converted into stars in a given dark matter halo (e.g., Henriques et al. 2015; Somerville & Dave 2015; Henriques et al. 2019). The current best strategy to reduce the cosmological efficiency of star formation (and keep hot gas haloes from collapse) is to introduce baryonic feedback in two principal forms: (i) stellar and supernova feedback (Cole et al. 2000; Bower et al. 2008; Schaye et al. 2015); and (ii) active galactic nucleus (AGN) feedback (Croton et al. 2006; Sijacki et al. 2007; Bluck et al. 2011; Maiolino et al. 2012; Sijacki et al. 2015; Vogelsberger et al. 2014a, 2014b; Schaye et al. 2015; Zinger et al. 2020). The former is highly effective at slowing the rate of star formation in low-mass systems, but fails to reduce star formation in the deep potential wells of massive galaxies (e.g., Fabian 2012; Henriques et al. 2019). On the other hand, AGN feedback becomes highly effective in high-mass systems, which host more massive central black holes (e.g., Sijacki et al. 2007, 2015; Weinberger et al. 2017, 2018). Hence, contemporary theoretical models are in near consensus



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

in their use, and requirement, of AGN feedback to explain the existence of massive quenched galaxies (e.g., Henriques et al. 2013, 2015, Vogelsberger et al. 2014a, 2014b; Schaye et al. 2015; Weinberger et al. 2017; Nelson et al. 2018; Pillepich et al. 2018; Weinberger et al. 2018; Dave et al. 2019; Zinger et al. 2020; Piotrowska et al. 2022).

Observationally, much skepticism surrounds the role of AGN in galaxy evolution, and especially their regulation of galaxy quenching. The primary reason for this is that a clear link between instantaneous AGN luminosity and reduction in star formation has not been forthcoming (see Bluck et al. 2023 for a discussion). Despite several early claims to support this link (e.g., Nandra et al. 2007; Bundy et al. 2008), the majority of contemporary observational studies of luminous AGN place these objects in actively star-forming galaxies (e.g., Hickox et al. 2009; Aird et al. 2012; Rosario et al. 2013; Heckman & Best 2014; Trump et al. 2015), apparently in contradiction to the notion of AGN feedback quenching galaxies. However, this objection (though logical) is the result of a fundamental misconception about the nature of AGN feedback in contemporary simulations, i.e., how AGN are actually theorized to quench galaxies.

As an example of this, in Piotrowska et al. (2022) we find that no strong link between specific star formation rate (sSFR) and current AGN accretion rate, or luminosity, is predicted by three state-of-the-art contemporary models (EAGLE, Illustris, and IllustrisTNG). This is in spite of the fact that all three of these simulations quench massive galaxies exclusively via AGN feedback. This result was confirmed in Ward et al. (2022), who found that the most powerful AGN reside in gas-rich star-forming galaxies, the antithesis of quenching galaxies. Hence, it becomes of paramount importance to resolve this apparent discrepancy.

The answer lies in the distinction between power (i.e., AGN luminosity) and work done to the galactic system (i.e., some fraction of the energy emitted by AGN across cosmic time). Of course, energy is simply the time integral of power, so the two are fundamentally related, but they are by no means identical. Contemporary cosmological simulations unanimously predict that quenching should scale fundamentally with *energy* input from AGN over the lifetime of galaxies, *not* instantaneous power, luminosity, or identification of AGN (see Terrazas et al. 2020; Zinger et al. 2020; Piotrowska et al. 2022; Bluck et al. 2023). The goal for an effective observational test of the AGN feedback paradigm of massive galaxy quenching is thus to find an observable that is strongly correlated with AGN energy, not AGN luminosity. Fortunately, in essentially all physical AGN models, the energy released over the lifetime of a black hole is directly proportional to its mass (see Soltan 1982; Silk & Rees 1998; Bluck et al. 2020a).

In Piotrowska et al. (2022), we find that all three of the simulations mentioned above predict that black hole mass should be the optimal observable predictor of massive galaxy quenching. Indeed, black hole mass is predicted to be far superior to dark matter halo mass, stellar mass, and black hole accretion rate (and hence AGN luminosity) in separating star-forming and quiescent galaxies. In Bluck et al. (2023), we extend this analysis to moderately high redshifts. Across 10 Gyr of cosmic history ($z = 0-2$), these simulations predict that the fundamental signature of AGN feedback should be a close dependence of quiescence on black hole mass. This is clearly identified via machine-learning classification with random forests (RFs; Piotrowska et al. 2022; Bluck et al. 2023). The

distinct advantage of the RF approach is that it can be used to carefully disambiguate intercorrelated *nuisance* parameters from the fundamental *causal* relationships (see Bluck et al. 2022 for a thorough demonstration, where we show that this approach can reverse-engineer the “LGalaxies” semi-analytic model, Henriques et al. 2013, 2015).

In the local Universe, the key prediction of AGN feedback-driven quenching has been tested in two complementary ways: (i) using a small sample of dynamically measured black hole masses (Terrazas et al. 2016, 2017; Piotrowska et al. 2022), and (ii) via the use of the black hole mass (M_{BH})–central velocity dispersion (σ) relationship (e.g., Ferrarese & Merritt 2000; Saglia et al. 2016), and other calibrations, to estimate black hole masses for samples of ~ 500 k galaxies (Bluck et al. 2016, 2020a, 2020b, 2022; Piotrowska et al. 2022). The results from both tests are in close agreement with each other, and moreover, are precisely as predicted by AGN feedback models.

Additionally, a fundamental link between quenching and the disordered (pressure-supported) kinematics in local galaxies is found in Brownson et al. (2022), with essentially no connection between quiescence and galactic rotation. This strongly suggests that quenching is critically linked to catastrophic redistribution of angular momentum within galaxies (from mergers or violent disk instabilities), which leads to both bulge and supermassive black hole growth (e.g., Hopkins et al. 2006, 2008). Hence, this kinematic result may be seen as yet further support for the AGN feedback paradigm in the local Universe.

At higher redshifts there is a dearth of dynamically measured black hole masses, and even measurements of central velocity dispersions are extremely rare. This is due to the fact that the vast majority of high- z galaxy surveys are predominantly photometric, with limited (or no) spectroscopic counterparts. In the coming years, this profound lack in our observational knowledge of early galaxies will be significantly remedied by the VLT-MOONRISE survey (Cirasuolo et al. 2020; Maiolino et al. 2020). However, in the meantime, we have found effective ways to test the AGN feedback paradigm in photometric data, through the application of the virial theorem (see Bluck et al. 2022, 2023).

In Bluck et al. (2023), we find that cosmological models predict that, in lieu of black hole mass, the stellar gravitational potential (given approximately as the ratio of stellar mass to galaxy half-light radius, $\phi_* \sim M_*/R_h$), ought to become the best predictor of massive galaxy quiescence. Remarkably, we find that this prediction is precisely met across 10 Gyr of cosmic history, through analyses of the Hubble Space Telescope-Cosmic Assembly Near-IR Deep Extragalactic Legacy Survey (HST-CANDELS) and the Sloan Digital Sky Survey (SDSS). Additionally, these comparisons with models provide a natural explanation to many prior observational studies which have found a close dependence of quenching on galaxy density, light concentration, and bulge mass (see, e.g., Wuyts et al. 2011; Cheung et al. 2012; Fang et al. 2013; Bluck et al. 2014, 2016, 2022; Lang et al. 2014).

With the launch of the James Webb Space Telescope¹⁰ (JWST; Gardner et al. 2006), and the subsequent first data release in 2022 (Pontoppidan et al. 2022), we are now in a position to extend the test of the AGN feedback paradigm

¹⁰ JWST Project: <https://webb.nasa.gov/>

essentially to the very first quiescent systems in the Universe. The reason JWST is needed for this task is twofold. First, the significantly larger JWST primary mirror, compared to HST, enhances the resolution of galaxies observed by (at least) a factor of 3, where there is wavelength overlap. This is essential for accurate measurements of galaxy sizes at high redshifts. Second, and arguably even more importantly, the longer wavelength detection capabilities of JWST NIRCам (compared to HST ACS and NICMOS) provide a rest-frame optical (and near-IR) measurement of galaxy structure at the earliest cosmic epochs for the first time. This enables accurate constraints of the stellar mass distributions of early galaxies since longer wavelengths are largely uncontaminated by bright young stellar populations (which typically contribute only a tiny fraction of the underlying mass budget of galaxies, yet dominate shorter wavelength measurements).

In this paper we apply techniques developed in Bluck et al. (2022, 2023), Piotrowska et al. (2022), and Brownson et al. (2022) to the largest public JWST galaxy survey (as of the time of writing): JWST-CEERS (Finkelstein et al. 2023b; Kartaltepe et al. 2023; Kocevski et al. 2023). We also provide updated quenching predictions from two state-of-the-art simulations at the very early epochs we can now probe with JWST (EAGLE: Schaye et al. 2015; IllustrisTNG: Nelson et al., IllustrisTNG: Nelson et al. 2018; Pillepich et al. 2018), as well as consistent comparisons to lower-redshift galaxy surveys (i.e., HST-CANDELS: Grogin et al. 2011; Koekemoer et al. 2011 and SDSS: Abazajian et al. 2009). This paper represents the first rigorous statistical test on the AGN feedback paradigm at cosmic dawn (the epoch in which the first galaxies form).

The paper is structured as follows. In Section 2, we discuss the various observational and simulated data sets used throughout this work. In Section 3, we describe our methodology for categorizing star-forming and quenched galaxies throughout cosmic time. Additionally, we provide a brief summary of our machine-learning classification technique utilizing RFs (which is extensively discussed in prior publications). In Section 4, we present our results, split between theoretical predictions and observational tests, including the primary JWST-CEERS analysis. In Section 5, we discuss the results of this paper in the context of the theory of galaxy formation and evolution. We summarize our contributions in Section 6.

2. Observations and Data

In this paper we test the high- z quenching predictions from hydrodynamical simulations using data from JWST-CEERS¹¹ (Finkelstein et al. 2023a, 2023b; Bagley et al. 2023; Kartaltepe et al. 2023; Kocevski et al. 2023). We additionally make comparisons of observations of later cosmic times in HST-CANDELS¹² (Faber 2011; Grogin et al. 2011; Koekemoer et al. 2011), and SDSS¹³ (Abazajian et al. 2009). Full details on the survey design, parameter measurements, and data access are provided in the above references. Here we give a review of the most important aspects of these data for this work.

2.1. HST+JWST Data

We utilize novel data from the public JWST-CEERS photometric survey (Finkelstein et al. 2023b; Kartaltepe et al. 2023;

Kocevski et al. 2023) to probe galaxy quenching at intermediate ($z = 0.5-2$), high ($z = 2-4$), and ultrahigh ($z = 4-8$) redshifts. Our analysis of this data represents the first test of quenching mechanisms in the very early Universe, where we are just now discovering that AGN are likely very common (e.g., Juodžbalis et al. 2023; Marshall et al. 2023; Ubler et al. 2023). Furthermore, the JWST data, in combination with lower- z surveys (see below), enable a unified analysis of star formation quenching in galaxies across essentially all of cosmic time.

To analyze the structural properties of our galaxies, we utilize NIRCам JWST observations processed through a customized reduction process from the Cosmic Evolution Early Release Science Survey (CEERS; ID = 1345). Specifically, we work with the subset of JWST data that overlaps with CANDELS; Grogin et al. 2011; Koekemoer et al. 2011) in the Extended Groth Strip (EGS) field. The data are independently reduced using a custom setup of the JWST pipeline version 1.6.2 (Bushouse et al. 2022), with calibration files from CDRS 0942. For a detailed account of this process and the resulting data quality, refer to Ferreira et al. (2023).

Our parent sample comprises the same 3965 sources at redshifts $z > 1.5$ from the CANDELS catalogs, as described in Ferreira et al. (2023). At present, this represents one of the most extensive catalogs of galaxy classifications at $z > 1.5$ using JWST. To conduct our analysis, we rely on well-calibrated and reliable photometric redshifts, star formation rates (SFRs), and stellar masses derived from the HST-CANDELS data (Duncan et al. 2014, 2019; Whitney et al. 2021). The selection of sources does not involve morphological information, nor do we use magnitude cuts, ensuring the inclusion of sources that might appear faint in HST but are clearly visible in JWST observations. Additionally, we incorporate information on galaxies at lower redshifts within this field, processed in the same manner, and compare them with the comprehensive multi-field HST-CANDELS data at redshifts where galaxies are well enough resolved without JWST imaging (i.e., at $z < 2$, see Dimauro et al. 2018).

As explained in Ferreira et al. (2023), we adopt two distinct approaches to galaxy morphology and structure. First, we conduct visual classifications for all sources (see also Conselice et al. 2024). Second, we employ quantitative structural measures using GALFIT (Peng et al. 2002), conducting parametric light profile fitting to measure sizes and Sérsic indices for each galaxy in each wave band available in the JWST observations (see Ormerod et al. 2024). A comparison of visual morphologies with the quantitative structural measurements can be found in Ferreira et al. (2023). For this work, our focus primarily centers on the quantitative structures, which enable robust comparisons with cosmological models. However, the visual morphologies have been used in the JWST observations to test and validate the more sophisticated quantitative structural measurements.

2.1.1. Photometric Redshifts, Stellar Masses, and SFRs

In this paper, we utilize photometric redshifts in the EGS field derived by Duncan et al. (2019). The software EAZY (Brammer et al. 2008) is employed for template fitting, where three distinct template sets are used and fitted to the photometric bands. These templates incorporate zero-point offsets to modify input fluxes and address wavelength-dependent errors. In addition to these redshifts, a Gaussian process code (GPz; Almosallam et al. 2016) is employed to obtain further empirical estimates using a

¹¹ JWST-CEERS data access: <https://ceers.github.io/overview.html>

¹² CANDELS data access: <https://archive.stsci.edu/hlsp/candels>

¹³ SDSS data access: <https://classic.sdss.org/dr7/access>

subset of the photometric bands. Individual redshift posteriors are calculated, and all measurements are combined in a statistical framework through a hierarchical Bayesian approach to arrive at a final redshift estimate. For a detailed explanation of the process, more details are available in Section 2.4 of Duncan et al. (2019).

The data used to determine these redshifts are obtained from the original CANDELS+GOODS WFC3/ACS imaging and data, Spitzer/IRAC S-CANDELS (Ashby et al. 2015), and ground-based observations with CFHT (Stefanon et al. 2017). The methodology we use for this is described in detail in Duncan et al. (2019). The photometric redshifts are highly accurate for the redshift range studied in this paper, as validated when compared with available spectroscopic redshifts within the EGS field (see Duncan et al. 2019 for various statistical tests).

To determine galaxy stellar masses, a modified version of the spectral energy distribution (SED) code (described in Duncan et al. 2014) is employed. This code allows for stellar mass measurement at all redshifts within the photo- z fitting range. The stellar masses are also accompanied by a *template error function*, described in Brammer et al. (2008), which accounts for uncertainties arising from the template set and any wavelength effects. Typical uncertainties on these measurements are estimated to be 0.2–0.3 dex.

More specifically, the stellar mass measurement technique employs Bruzual & Charlot (2003) model spectral templates, encompassing a wide range of stellar population parameters, and we adopt and use a Chabrier 2003 initial mass function (IMF). The applied star formation histories follow exponential τ -models for both positive and negative values of τ , enabling the representation of both rising and declining star formation histories (which is crucial at high redshifts). The timescales utilized include $|\tau| = 0.25, 0.5, 1, 2.5, 5,$ and 10 , along with a short burst model ($\tau = 0.05$), and continuous star formation models.

To ensure the reliability of our measured stellar masses, we compare these with other values independently calculated by other teams within the CANDELS collaboration (see Santini et al. 2015). This comparison aims to rule out systematic biases affecting the stellar mass estimates. While some scatter is observed between the two mass estimates ($\sigma \sim 0.3$ dex), the masses used in this study remain unaffected by any significant biases compared to others. Further details regarding the method and models can be found in Section 2.5 of Duncan et al. (2019).

The SFRs used in this work are derived from star formation histories extracted from the SED fitting of the galaxies (i.e., from the weighted sum of simple stellar populations of the best composite fit). From the best SED fit, the UV light from these galaxies is examined, and corrections are made for dust using the UV slope of the spectral energy distribution (β). This correction provides a measure of the dust attenuation in the galaxy, enabling the determination of the total SFR corrected for extinction. Our SFRs are found to agree extremely well with SFRs derived directly from the SED fitting (see Duncan et al. 2014).

2.1.2. Galaxy Sizes from Surface Brightness Fitting

Galaxy sizes are a crucial aspect of this paper. As such, we describe how these are measured here (although for full details see Conselice et al. 2024; Ormerod et al. 2024).

We determine half-light radii and Sérsic indices for JWST-CEERS galaxies using the GALFIT software. GALFIT operates by fitting 2D Sérsic light profiles to each wave band image of each galaxy (Peng et al. 2002). Ultimately, GALFIT is a least squares fitting algorithm designed to find the optimum global solution to a 2D galaxy light profile, operating at its core is a Levenberg–Marquardt algorithm. GALFIT determines the goodness of fit via the χ^2 statistic for each free parameter configuration, and then systematically adjusts parameters until the χ^2 is minimized. More precisely, GALFIT performs this optimization using the reduced chi-squared loss function, χ_ν^2 , determined by

$$\chi_\nu^2 = \frac{1}{N_{\text{DOF}}} \sum_{x=1}^{n_x} \sum_{y=1}^{n_y} \frac{(f_{\text{data}}(x, y) - f_{\text{model}}(x, y))^2}{\sigma(x, y)^2}, \quad (1)$$

which is summed over n_x and n_y pixels, and where N_{DOF} is the number of degrees of freedom in the model. GALFIT uses a data image, $f_{\text{data}}(x, y)$, and a weight map, $\sigma(x, y)$, which are utilized together to calculate the best-fit model image, $f_{\text{model}}(x, y)$ (see Peng et al. 2002). A full description of this method (along with numerous quality tests) is presented in Ormerod et al. (2024).

We utilize the obtained sizes and Sérsic indices for JWST-CEERS galaxies to explore the statistical connection between galaxy structures and the processes of star formation quenching at high redshifts. The JWST data brings a major advantage here, as we can now leverage these radius measurements to construct an accurate estimate of the stellar potential ($\phi_* \sim M_*/R_h$), and quantify the 2D and 3D stellar mass densities of galaxies within the CEERS field in the rest-frame optical. This is the first time that rest-frame optical structural measurements have been available at redshifts $z > 2$, thanks to the unique capabilities of JWST.

2.2. Ancillary Observational Data I: HST-CANDELS

We incorporate the full multi-field HST-CANDELS survey into our analysis to probe intermediate redshifts ($z = 0.5$ – 2). The limiting factor determining the maximum redshift is set by galaxy resolution and the need for rest-frame optical measurements (see, e.g., Dimauro et al. 2018). As such, HST-CANDELS acts as a test on the intermediate redshift analysis from JWST-CEERS. The advantage is that the full HST-CANDELS survey is much wider, and hence probes a much larger volume at these epochs (with correspondingly higher galaxy counts).

Specifically, we utilize data from value-added catalogs¹⁴ for the HST-CANDELS survey provided by Dimauro et al. (2018). From this data release, we utilize photometric redshifts, stellar masses, SFRs, and galaxy half-light radii, as well as rest-frame optical magnitudes and colors (for various data quality checks). All of the above are constructed via a comprehensive morphological SED pipeline, which also produces bulge and disk masses, as well as a host of other structural information. The SED fitting is performed using FAST (Kriek et al. 2009), employing Bruzual & Charlot (2003) stellar population synthesis models, and a Calzetti et al. (2000) extinction law. Full details on these procedures are provided in Dimauro et al. (2018).

¹⁴ CANDELS VAC: <https://mhurtascompany.weebly.com/data-releases-and-codes.html>

2.3. Ancillary Observational Data II: SDSS

We also incorporate SDSS-DR7 legacy survey (Abazajian et al. 2009) into our analysis to compare our high- z results to the low- z baseline. From SDSS, we utilize a variety of public value-added catalogs. We take spectroscopic redshifts from the parent data release (Abazajian et al. 2009). We take SFRs from Brinchmann et al. (2004). These are computed via two distinct modes: (i) for emission line, non-AGN galaxies, SFRs are estimated from dust corrected emission lines (particularly dependent on $H\alpha$); (ii) for non-emission line or AGN contaminated galaxies, SFRs are computed from an empirical relationship between sSFR and the strength of the 4000 Å break (D4000). We take stellar masses from the SED catalogs of Mendel et al. (2014), and galaxy half-light radii from the morphological catalogs of Simard et al. (2011). A concise summary of the techniques used in all of these measurements, as well as numerous data quantity checks, are provided in Bluck et al. (2014, 2016).

2.4. Cosmological Simulations

In this paper, we extract the testable predictions for high-mass galaxy quenching within the first few gigayears of cosmic history from two state-of-the-art cosmological hydrodynamical simulations: (i) EAGLE¹⁵ (Crain et al. 2015; Schaye et al. 2015; McAlpine et al. 2016); and (ii) IllustrisTNG¹⁶ (hereafter TNG; Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018; Springel et al. 2018; Nelson et al. 2019). In previous work, we have also extracted quenching predictions from Illustris (Vogelsberger et al. 2014a, 2014b). However, this latter simulation fails to predict quiescent objects in significant numbers at $z > 2$, so cannot be used at the epochs targeted in this work. Moreover, TNG replaces Illustris as a more realistic and successful version of the latter.

Full details on the simulations are given in the references above, as well as a concise overview and comparison between them in Piotrowska et al. (2022). Additionally, a thorough description of the parameters used in this work is presented in Bluck et al. (2023). Here we extract the exact same parameters and apply the exact same quality checks as in our previous work, the only difference being that we now view the $z = 2-4$ snapshots, concatenated together, for our quenching analyses. A docker¹⁷ describing how to extract these data from the simulation websites is provided in Piotrowska et al. (2022). The only modification needed is the snapshot redshifts (given above).

In this work, for simulations, we select only high-mass galaxies ($M_* > 10^{9.5} M_\odot$) in relatively high-mass haloes ($M_{\text{halo}} > 10^{11} M_\odot$). This ensures that all galaxies studied are well resolved in the simulations. Furthermore, it enables robust comparison with high- z JWST observations, which are essentially mass complete at these masses (e.g., Ferreira et al. 2023). We apply identical stellar mass cuts in observations, but cannot apply halo mass cuts since these are not available in the JWST and HST galaxy surveys. Nonetheless, we have checked that the halo mass criteria make no significant difference to the simulation predictions. It is included primarily to exclude the possibility of studying galaxies without seed black holes (which are extremely rare at these stellar masses).

Given that these simulations are very well discussed in the literature, and that we have previously outlined our data extraction and quality assurance methodology, we give only a brief summary of the most relevant details for this work (i.e., with respect to galaxy quenching).

2.4.1. EAGLE

We utilize EAGLE-RefL0100N1504 (Schaye et al. 2015), which has a box size of $\sim 100 \text{ cMpc}^3$, and is run with the GADGET-3 smoothed particle hydrodynamics (SPH) code (Springel 2005). The input cosmology is given by Planck Collaboration I (2014), assuming a spatially flat, dark energy Λ CDM universe. The sole quenching mechanism for massive galaxies in EAGLE is AGN feedback. Supermassive black holes are seeded in low-mass dark matter haloes and grow via Eddington-limited Bondi–Hoyle accretion (Hoyle & Lyttleton 1939; Bondi & Hoyle 1944).

Feedback from AGN is modeled as a single mode in EAGLE. A fraction of the accreted rest energy of matter is converted to energy, with a fraction of that energy coupled to the interstellar medium (ISM). Energy exchange between the black hole accretion disk and the galaxy is modeled via a pure thermal injection, applied in a stochastic *bursty* manner (see Crain et al. 2015). This mechanism functions both to trigger quenching, via ISM outflows, and to keep massive galaxies quenched, by heating the surrounding circumgalactic medium (CGM). However, this model is known to be less effective at maintaining quiescence in high-mass galaxies compared to observations (see Piotrowska et al. 2022).

2.4.2. IllustrisTNG

We utilize TNG-100-1 (Nelson et al. 2018; Pillepich et al. 2018), which has a box size of $\sim 100 \text{ cMpc}^3$, and is run with the AREPO unstructured moving mesh code (Springel 2010), updated to implement magnetic fields in addition to gravity and hydrodynamics. The input cosmology is given by Planck Collaboration (2016). As in EAGLE, supermassive black holes are seeded in low-mass dark matter haloes and evolve in accordance with Eddington-limited Bondi–Hoyle accretion (Hoyle & Lyttleton 1939; Bondi & Hoyle 1944).

Unlike in EAGLE, TNG operates three distinct forms of AGN feedback (see Sijacki et al. 2007; Nelson et al. 2018): (i) *quasar mode* (high Eddington ratio); (ii) *kinetic mode* (low Eddington ratio); and (iii) radiative ionization (all Eddington ratios). However, only the second mechanism (kinetic mode) significantly impacts quenching in this simulation (see Weinberger et al. 2017; Zinger et al. 2020; Piotrowska et al. 2022). In this mode (which is only operational at $M_{\text{BH}} \gtrsim 10^8 M_\odot$), a fraction of the accreted rest energy of matter is converted to the kinetic energy of gas cells in a stochastic manner (see Weinberger et al. 2017). This leads to turbulence and outflows in the ISM, as well as long-term (preventative) heating of the CGM through gas percolation and shocks. The combination of these effects leads to quenching of massive galaxies in this simulation (see Terrazas et al. 2020; Zinger et al. 2020; and Piotrowska et al. 2022 for illuminating discussions). Note that the kinetic mode replaces the overly zealous *radio mode* in Illustris (see Sijacki et al. 2007), which had the undesired consequence of effectively destroying the CGM around massive galaxies.

¹⁵ EAGLE data access: <http://icc.dur.ac.uk/Eagle/>

¹⁶ IllustrisTNG data access: www.tng-project.org/.

¹⁷ Simulations docker: <https://hub.docker.com/u/jpiotrowska>

Table 1
Data Summary

Sample	N_{gal}	N_{SF}	N_{Q}
Observations			
JWST ($0.5 < z < 2$)	384	243	141
JWST ($2 < z < 4$)	389	300	89
JWST ($4 < z < 8$)*	57	52	5
CANDELS ($0.5 < z < 2$)	9,417	6,034	3,383
SDSS ($0.02 < z < 0.2$)	498,206	247,364	250,842
Simulations			
EAGLE ($2 < z < 4$)	5,999	5,289	710
IllustrisTNG ($2 < z < 4$)	10,890	10,011	879

Note. All data are selected at $M_* > 10^{9.5} M_{\odot}$. Simulations are additionally required to have $M_{\text{halo}} > 10^{11} M_{\odot}$. The quenching threshold is set at $\text{sSFR} = 1/H(z)$ (see Figure 1), except for SDSS (which is set at $\text{sSFR} = 10^{-11} \text{ yr}^{-1}$). * Not used in machine-learning analysis. Additionally note that the numbers of galaxies decrease rapidly with redshift, as a result of high- z surveys having much smaller areas than low- z surveys (see the source papers).

2.5. Sample Selection

For all machine-learning analyses in this paper, we select massive galaxies with $M_* > 10^{9.5} M_{\odot}$. For simulations, we additionally require that $M_{\text{halo}} > 10^{11} M_{\odot}$. In simulations, we select only central galaxies (defined as being the most massive member in each dark matter halo) for extracting quenching predictions. This enables us to focus on intrinsic quenching mechanisms in simulations. For observations, we suspect that we are most sensitive to intrinsic quenching as well as these masses, but this is explicitly tested through comparison with the simulations (see also Bluck et al. 2023 for a discussion on this point).

We separate our sample into various redshift bins (see Table 1). The redshift cuts are motivated by studying galaxies at redshifts below and above cosmic noon ($z \sim 2$; see Madau & Dickinson 2014). Additionally, we separate galaxies into star-forming and quiescent systems based on their sSFR (see Section 3.1). The breakdown of total galaxy counts per sample, and the numbers of star-forming and quiescent objects are provided in Table 1.

For the JWST observations we impose two quality cuts on the data: (i) *chi-square-robust* = *True*, which ensures high-quality SED fits (see Duncan et al. 2014, 2019); and (ii) *Good-Fit* = *True*, which ensures high-quality structural fits with GALFIT (see Ormerod et al. 2024). For the SDSS and CANDELS data, we apply the exact same quality cuts as in Bluck et al. (2022, 2023).

To obtain the redshift range in simulations, we concatenate the $z = 2-4$ snapshots from both EAGLE and TNG. We have also checked that our analysis is stable for viewing each snapshot individually.

Toward the end of Section 4, we extend our analysis to explore galaxies at lower masses in the JWST data ($M_* > 10^8 M_{\odot}$), where we are sensitive to environmental quenching routes in addition to intrinsic quenching.

3. Methods

3.1. Identifying Quiescent Galaxies across Cosmic Time

In Figure 1, we present the evolution in sSFR ($\text{sSFR} \equiv \text{SFR}/M_*$) as a function of redshift for galaxies observed in JWST-CEERS. Gray contours show the location of all galaxies. The geometric

mean sSFR for galaxies with $\text{sSFR} > 10^{-10} \text{ yr}^{-1}$ (the quenched threshold at $z = 0.5$, see Bluck et al. 2023) is represented by blue stars, with error bars indicating the full 1σ range. It is clear that the typical sSFR of star-forming galaxies rises significantly with redshift, all the way out to $z = 8$, i.e., to ~ 650 Myr after the Big Bang. This justifies the cut in sSFR above, which is used as a coarse method to remove obviously quenched systems before assessing evolution in star formation, and ultimately constructing a more accurate quenching threshold (discussed below).

From the density contours in Figure 1, it is clear that there is a significant population of galaxies with sSFR much lower (and higher) than the median relation. To identify quiescent systems we utilize the common approach at intermediate-to-high redshifts of considering the inverse Hubble time (e.g., Tacchella et al. 2019; Bluck et al. 2022, 2023). This is displayed in Figure 1 as a dashed orange line, which clearly does a good job of identifying low sSFR systems (relative to the norm of the star-forming population). Additionally, we include a more stringent cut at a value of sSFR a factor of 3 lower (shown as a red dashed line). Using these thresholds, we see in Figure 1 that quiescent galaxies exist up to $z = 6-7$ (as confirmed in, e.g., Looser et al. 2023), yet the vast majority of quenched galaxies in JWST-CEERS are at $z < 4$.

Throughout this paper, we define quenched galaxies to be any system that has an sSFR below the inverse Hubble time (at its epoch in cosmic history), with star-forming systems being any galaxy with an sSFR above this threshold. Additionally, we define fully quenched systems to be galaxies with sSFR values below a third of the inverse Hubble time (as in Tacchella et al. 2019). All results remain unchanged for either quenching threshold, indicating that our results are highly stable to the exact threshold adopted for separating star-forming and quiescent objects. We apply the exact same quenching thresholds to all observational and simulated data sets used in this work to ensure a fair comparison (with the sole exception of the $z = 0$, SDSS data, which is discussed below).

Utilizing data from HST-CANDELS, we found that the sSFR quenching threshold approach leads to quiescent populations that are extremely similar to *UVJ* color diagnostics, as well as to a simple -1 dex offset from the peak sSFR at each epoch (see Bluck et al. 2022, 2023). This indicates that the choice of method used to identify quiescent systems is not a significant source of bias or uncertainty at $z < 2$. However, at very high redshifts galaxy color is more problematic as a quenching indicator because essentially all stellar populations are relatively young (due simply to the age of the Universe). As such, an sSFR approach calibrated on finding evidence of very massive (short-lived) stars is a superior approach (see, e.g., Carnall et al. 2023).

At very low redshifts (i.e., in the local Universe), the above-defined quenching thresholds break down. In fact, the star-forming main-sequence relationship perfectly intersects with the inverse Hubble time at $z = 0$, making it clearly invalid as a quenching definition. The reason for this is likely due to dark energy leading to accelerated expansion, and hence, exponential decline in gas accretion (see Henriques et al. 2013, 2015). As such, for the SDSS data (which we use as a low- z comparison to our high- z results), we adopt the standard $z = 0$ quenching threshold of $\text{sSFR} < 10^{-11} \text{ yr}^{-1}$ (see, e.g., Bluck et al. 2022; Piotrowska et al. 2022; Bluck et al. 2023). This yields a similar offset from the mean sSFR of star-forming

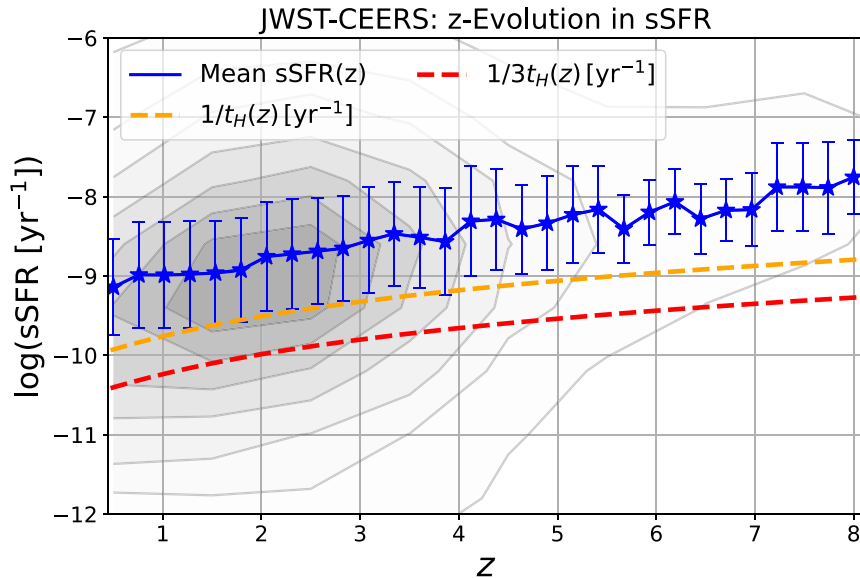


Figure 1. Evolution in sSFR as seen in JWST-CEERS. The mean $\log(\text{sSFR})$ at each epoch for star-forming systems is represented by blue stars, with error bars displaying the full 1σ range. The location of all galaxies in the sSFR - z plane is displayed by gray contours (with linearly spaced intervals). Two quenching thresholds are displayed on the figure, representing (i) $\text{sSFR} < 1/t_{\text{Hubble}}(z)$ (displayed in orange); and (ii) $\text{sSFR} < 1/3t_{\text{Hubble}}(z)$ (displayed in red). These thresholds clearly do a good job at all epochs studied in JWST-CEERS of separating star formers from quiescent systems. By design, the red threshold is more robust for identifying quiescence than the orange threshold, at the price of reduced numbers of quiescent systems. We use both definitions of quiescence throughout this paper for various purposes. Clearly, quiescent galaxies exist in the JWST-CEERS field up to $z = 6-7$. However, the vast majority of quenched galaxies reside at $z < 4$.

systems as the inverse Hubble time definition at higher redshifts (see Bluck et al. 2022).

One potential issue with the sSFR approach is that the star-forming main-sequence (SFR- M_*) relation is actually slightly sublinear (with a gradient of 0.7–0.8; see, e.g., Bluck et al. 2016, 2020b), whereas a threshold in sSFR implicitly assumes a gradient of unity. As such, a fixed threshold in sSFR slightly overidentifies quiescent galaxies at very high masses. However, this effect is very small and we have checked that alternative Δ MS methods (i.e., offsets from a main-sequence fit) yield equivalent results to the more straightforward sSFR criterion.

Ultimately, the main problem with measuring offsets from the star-forming main sequence is that this relation is quantitatively different in simulations and observations, in addition to its intrinsic redshift evolution. Conversely, the sSFR method is an objective definition of quenching, which can be used in simulations and observations without the need for fitting bespoke star-forming main-sequence relations (which may obscure important differences between models and observations). As such, we choose the sSFR method for its objectivity and reproducibility (see Bluck et al. 2023 for further discussion on this point). Nonetheless, we note that this yields identical conclusions to a Δ MS approach in any case (see also Bluck et al. 2022 for a demonstration of this).

3.2. RF Classification

In Bluck et al. (2022), we present a detailed overview of RF classification, including a series of tests that demonstrate its capacity to accurately distinguish causal from nuisance parameters in simulated data. Ultimately, this is the key advantage of this machine-learning technique over other approaches. The RF approach has been used in numerous previous papers to hone in on the fundamental drivers of galaxy quiescence (see Bluck et al. 2020a, 2020b, 2022; Brownson et al. 2022; Piotrowska et al. 2022; Bluck et al. 2023). As such,

we give only a brief outline of the most salient aspects of this method here.

RF classification is a type of supervised machine learning that provides an effective route to categorize data into types, based on given input features (or parameters). It is capable of discerning highly nonlinear boundaries in a complex multi-dimensional parameter space (see Bluck et al. 2022 for several demonstrations). However, unlike deep learning with, e.g., artificial neural networks (ANNs), the RF approach is emphatically not a *black box*. Every step in the iteration may be output, and given the simple criteria for developing each tree (discussed below), clearly understood by humans.¹⁸ Consequently, this approach yields an optimal compromise between sophistication and interpretability, which is ideal for the application to astronomical data. The power of this approach has been thoroughly demonstrated by its capacity to accurately reverse-engineer both semi-analytic and hydrodynamical simulations, i.e., revealing the input physics in these models from their output galaxy catalogs (see Bluck et al. 2022; Piotrowska et al. 2022; Bluck et al. 2023).

Briefly, an RF is a set of decision trees with differences enforced from bootstrapped random sampling. Within a given tree, the classifier chooses the most effective feature and threshold to minimize the impurity (quantified by the Gini parameter) of a sample. This continues throughout each level of the decision tree until all data are perfectly segregated, or else a predefined limit is reached. This limit can be used to mitigate the potential for overfitting (see below). Finally, the *votes* on the class of each object in the sample are averaged over each tree in the forest to yield a final class prediction and a probability distribution, which encodes the uncertainty of that prediction.

¹⁸ However, note that strategies exist to achieve this with deep learning as well via the use of interpreters. Yet, such methods are highly complex and infrequently incorporated in standard machine learning.

In this work, we construct an RF classifier using the `scikit-learn` Python package (Pedregosa et al. 2011). Following our exploration in Bluck et al. (2022), we make every feature available to the classifier at every step in the RF, which enables disambiguation of intercorrelated nuisance parameters up to correlations of $\rho = 0.99$, or higher (see Piotrowska et al. 2022). As such, our bespoke RF approach offers one route to move beyond the canonical problem of correlation not implying causation. By using RF classification one can definitively prove that a parameter is only useful in a classification problem due to its intercorrelation with another parameter (i.e., that it is noncausal). An interesting case in point is that of stellar mass, which while well known to correlate strongly with quiescence (see Baldry et al. 2006; Peng et al. 2010, 2012), is clearly established to be noncausal via RF classification in the local Universe (see Bluck et al. 2020a, 2022; Brownson et al. 2022; Piotrowska et al. 2022; Bluck et al. 2023).

As in all supervised machine learning, it is always possible to perfectly solve the problem on test data (seen by the classifier). However, the real value of the RF classification is that it can be effectively applied to novel data (unseen by the classifier). To ensure this, we split the sample into training (typically 50%–70% of the data) and testing (typically 30%–50% of the data), and require that the performance (as measured by the area under the true positive–false positive receiver operator curve, AUC, see Teimoorinia et al. 2016) is extremely similar in both samples. Specifically, we require $|\Delta\text{AUC}| < 0.02$ (see Bluck et al. 2023). This ensures very similar performance on unseen samples to the training data, which prevents the classifier from learning pathological features in the data. Consequently, the RF classifier learns to classify novel data almost as accurately as its training sample.

Due to the transparency of the RF architecture, the importance of a given feature for solving the classification problem can be extracted. It is given simply as the weighted average across the entire RF of the improvements in Gini coefficients engendered by each feature in turn, normalized by the total reduction in Gini engendered by all parameters (see Bluck et al. 2022, 2023). As such, this gives a clear intuitive sense of how valuable (or important) a given parameter is for solving the classification problem, within the context of the chosen set of parameters. Consequently, it is of *relative importance*: an importance relative to the available list of parameters. As such, one can always rule out causality in a given parameter for a given application (e.g., quenching), but can only find causality in the limit where every conceivable parameter has been assessed. Hence, one approaches causality from RF classification via a process of elimination.

In this work, we use RF classification to separate star-forming and quiescent galaxies in observations and simulations (based on the training thresholds discussed in Section 3.1). We train the RF classifier with a variety of parameters, which are known or suspected to be connected to various quenching mechanisms. The RF classifier then sorts the parameters by how valuable they are for predicting quiescence in each data set. We then use these results to test theoretical models, and ultimately identify the underlying causes of transition from star formation to quiescence in galaxies across cosmic time. We also provide other simpler statistical tests on the main RF classification results to establish confidence in these relatively new methods. Readers interested in the full mathematical

details of RF classification are encouraged to read Appendix B in Bluck et al. (2022). Additionally, in Appendix A of this work, we provide full details on the RF hyperparameters used for each classification analysis in this paper. This is to ensure the reproducibility of our results.

4. Results

4.1. High- z Quenching Predictions from Hydrodynamical Simulations

In Bluck et al. (2023), we expand on the low- z analysis of Piotrowska et al. (2022) to reveal the key quenching predictions of hydrodynamical simulations at $z = 0$ – 2 . We find that in EAGLE, Illustris, and TNG, black hole mass (M_{BH}) is clearly the best parameter for predicting quiescence. This is explained as a consequence of the total energy released by AGN feedback over the lifetime of a galaxy being directly proportional to the mass of the central black hole (see Bluck et al. 2020a). In lieu of a dynamical measurement of black hole mass, we find that the stellar potential (ϕ_*) is predicted by these simulations to act as an excellent proxy, outperforming many other parameters in predicting quiescence in models (Bluck et al. 2023).

In this subsection, we extend the AGN feedback predictions from EAGLE and TNG (Illustris has essentially no quiescent galaxies at $z > 2$) to the highest redshifts at which quiescent galaxies exist in these simulations.

In Figure 2, we show results from a set of RF classification runs used to separate central galaxies (at $M_* > 10^{9.5} M_\odot$) into star-forming and quenched categories, based on the threshold: $\text{sSFR} = 1/t_H(z)$ (orange dashed line in Figure 1). We have also tested the impact of using the more severe cut in quenching ($\text{sSFR} = 1/3t_H(z)$, red dashed line in Figure 1) and find that the following conclusions remain completely stable. We use central galaxies in the simulations to extract the quenching predictions for AGN feedback in isolation (since only satellites experience strong environmental quenching in these models).

In the top panel of Figure 2, we show results from using a large variety of parameters to train the RF classifier: black hole mass (M_{BH}); black hole accretion rate (\dot{M}_{BH}); stellar gravitational potential ($\phi_* \equiv M_*/R_H$); dark matter halo mass (M_H); galaxy half-mass radius (R_h); stellar mass (M_*); mass times radius (M_*R_h); and finally a random number (Rdm), for comparison purposes. These parameters have been carefully chosen to represent various potential quenching channels (e.g., supernova feedback–stellar mass; virial shock–halo mass; AGN feedback–black hole mass or black hole accretion rate, depending on the mode of operation, etc.).

In the top panel of Figure 2, it can be seen that black hole mass is overwhelmingly the most valuable parameter for identifying quenched systems, vastly outperforming all other variables in both the EAGLE and TNG simulations at very high redshifts ($2 < z < 4$). Note that this is the first time we have determined the key observable predictors of quiescence in these simulations at these very early cosmic times (but see Piotrowska et al. 2022; Bluck et al. 2023 for the predictions at later epochs, $z < 2$). This is essential for accurate comparison to the JWST observations (see Section 4.2). At even higher redshifts, the number of quiescent galaxies falls to zero in TNG and near zero in EAGLE, so further statistical analysis at even earlier cosmic times is not possible in these simulations.




Figure 2. RF quenching classification analyses to predict when galaxies will be star-forming or quiescent in cosmological simulations at high redshifts. Results are shown separately for the TNG and EAGLE cosmological simulations (see legends). The top panel displays results for training with all parameters (M_{BH} , \dot{M}_{BH} , ϕ_* , M_H , R_h , M_* , M_*R_h , Rdm). The bottom panel displays results for a reduced set of parameters, which are possible to measure accurately in extant observational data at high- z from JWST. The relative importance of quenching is displayed on the y-axis for each parameter. Uncertainties on the quenching importance are inferred from the variance over multiple RF classification runs. For the full parameter set (top panel), it is evident that supermassive black hole mass is clearly the critical parameter for setting quenching in the models. Interestingly, there is essentially no importance given to black hole accretion rate, dark matter halo mass, stellar mass, or galaxy stellar potential, once black hole mass is controlled for. However, in the restricted set of parameters (bottom panel), we see that ϕ_* is predicted to be found as the best observable parameter for quenching in lieu of black hole mass. This is a key testable *prediction* of AGN feedback quenching in the early universe for photometric surveys. See Bluck et al. (2023), and Figures 2 and 3 for equivalent results at lower redshifts. See Table 2 in Appendix A for full details on the RF hyperparameter setup.

Note that the lack of importance between stellar mass, halo mass, stellar potential, and quiescence in Figure 2 (upper panel) does not indicate a lack of correlation. In the RF, all other parameters are controlled when assessing the importance of any given training parameter. Hence, the lack of importance indicates a lack of residual information relevant to quenching once black hole mass is made available to the classifier. This clearly establishes that the well-known correlations between quenching and these parameters are ultimately spurious, i.e., of noncausal origin, originating solely from intercorrelation with the causal black hole mass.

It is especially important to highlight the lack of importance of accretion rate in models for predicting galaxy quiescence in

the very early Universe. Hence, efforts to test AGN feedback in the early Universe by using tracers of accretion rate (i.e., AGN luminosity or detection fractions) are destined to fail. The correct approach to test the modern AGN feedback paradigm is to look for parameters closely connected with black hole mass, and hence, the integrated energy released from AGN over the lifetimes of their host galaxies (see Section 1, Piotrowska et al. 2022; Bluck et al. 2023 for further discussion).

Given that extremely few black hole masses have been dynamically measured in the very early Universe (and none in our sample), we must now look for a predicted proxy in order to test the hydrodynamical predictions in observations. Following the logic of Bluck et al. (2023), there is strong

evidence to suspect that supermassive black hole mass should scale strongly with the gravitational potential of the galaxy (a consequence of the virial theorem in combination with the $M_{\text{BH}}-\sigma$ relation). Here we estimate the gravitational potential with the stellar gravitational potential, which we take to be approximately given by the ratio of stellar mass to galaxy half-mass radius (see Bluck et al. 2023 for many caveats). See Appendix B of this paper for an explicit demonstration of the strong relationship between M_{BH} and ϕ_* predicted by both TNG and EAGLE at early cosmic times.

In the bottom panel of Figure 2, we restrict the number of variables used to train the RF classifier to those that can be accurately constrained in JWST data. In this reduced parameter set, both EAGLE and TNG predict that quenching should be best constrained by ϕ_* , which significantly outperforms M_* , R_h , and M_*R_h . The latter variables all perform (at best) only marginally better than a random number, once ϕ_* is available to the classifier. This result is expected from the simple arguments sketched above (and in Bluck et al. 2023). Hence, the simulations predict that black hole quenching will look like stellar potential quenching in photometric data.

Taken as a whole, Figure 2 shows that (in both EAGLE and TNG) M_{BH} is the fundamental parameter that regulates central galaxy quenching at early cosmic times. Moreover, in observational data, AGN feedback-driven quenching should look like a strong dependence of quiescence on ϕ_* . This represents a clear testable prediction of the AGN feedback paradigm in the very early Universe.

It is important to note the stability of this prediction to the specifics of the AGN quenching model (which are very different in EAGLE and TNG; compare Schaye et al. 2015 with Weinberger et al. 2017, 2018). The reason for the similarity in their predictions is a consequence of both models ultimately extracting energy from around supermassive black holes to quench massive galaxies. The exact nature of the extraction, its coupling to the galaxy and halo, and the triggering mechanism (s) are all secondary to the crucial fact of the energy source. Consequently, there is a clear way to test the entire paradigm of AGN feedback, not just one specific instantiation of it (see Bluck et al. 2020a, 2023 for more details on this important point). In the next subsection, we perform this test in JWST-CEERS.

4.2. HST+JWST Results

4.2.1. RF Classification

In Figure 3, we present results from an RF classification to predict quiescence in galaxies, applied to JWST-CEERS (at $z=0.5-2$ and $2-4$). For comparison, we show equivalent results from HST-CANDELS (at $z=0.5-2$), and SDSS (at $z=0.02-0.2$). The parameters used in training the classifier are listed along the x -axis, and are identical to those used in the bottom panel of Figure 2 for simulations. The quenching threshold is taken as the orange dashed line in Figure 1. We have also checked and confirmed that a more constraining threshold (red dashed line in Figure 1) yields essentially identical results.

For the structural measurements, we choose JWST wave bands $\lambda = 3.56 \mu\text{m}$ (at $z = 2-4$) and $\lambda = 2.00 \mu\text{m}$ (at $z = 0.5-2$); HST-CANDELS H band: $\lambda = 1.65 \mu\text{m}$ (at $z = 0.5-2$); and SDSS: r band $\lambda = 0.65 \mu\text{m}$ (at $z = 0.02-0.2$). This ensures that at all epochs probed, spanning 13 Gyr of cosmic evolution, we

probe approximately the red optical to near-infrared in the rest frame of galaxies. The exact wave bands used are, of course, limited by availability and the quality of structural measurements achieved. For example, in HST-CANDELS we use the reddest available band, to best approximate a mass distribution; but in SDSS the r -band measurements are chosen for their superior quality. We find only a modest dependence on wave band for most ranges in JWST (see later in this section). Moreover, these small variations in rest-frame wavelength coverage could at most lead to more significant *differences* between the varying surveys (and our principal result is a lack of variation, see below).

From Figure 3 it is clear that the stellar gravitational potential, ϕ_* , is by far the most predictive variable of galaxy quenching at all epochs. The stellar potential greatly outperforms stellar mass, galaxy size, and mass times radius (all of which perform similarly to a random variable). This is in remarkable agreement with the predictions from the EAGLE and TNG hydrodynamical simulations (see Figure 2, and Bluck et al. 2023 for lower- z simulations results). In the models, this is a result of AGN feedback driving quenching, and central black holes growing in dense potential wells.

We emphasize the remarkable stability in the importance of parameters for predicting quenching in galaxies across 13 Gyr of cosmic history seen in Figure 3. This strongly suggests that quenching must be a stable mechanism operating across cosmic time. That is, at the base, there is no fundamental observable difference between quenching at high or low redshifts. Again, this result is precisely predicted by hydrodynamical simulations (compare the results at $z < 2$ in Bluck et al. 2023 to Figure 3). This is a truly impressive achievement for the AGN feedback paradigm, which was designed and calibrated using only low redshift data.

Unlike in simulations, in observational data, there are uncertainties in the measurements of parameters used to train the classifier. As such, it is important to carefully consider whether these effects could lead to an erroneous conclusion. The parameters used in Figure 3 have been carefully chosen to mitigate this potential problem.

The uncertainties on stellar mass are estimated to be $\sim 0.2-0.3$ dex, accounting for uncertainties in both SED fitting and the SED template libraries used (see Duncan et al. 2019). The uncertainties on galaxy sizes are $\sim 0.1-0.2$ dex (see Ormerod et al. 2024). Since both ϕ_* and M_*R_h are formed from linear combinations of mass and radius, their errors are given (by adding in quadrature) as $\sim 0.25-0.35$ dex. In Bluck et al. (2022), we demonstrate that RF classification results are extremely stable for parameters with the same, or similar, relative errors, but that differential measurement uncertainty can lead to less well-measured parameters performing worse than they should.

Crucially, in the example of Figure 3, the worst measured parameters are ϕ_* and M_*R_h . This is a necessary consequence of these parameters taking the other two parameters as inputs.¹⁹ Yet, the former clearly outperforms stellar mass and radius. Hence, this result cannot be explained by differential uncertainty. Furthermore, the relative uncertainties on ϕ_* and

¹⁹ For a relationship $C=A/B$, the relative error on C is given by: $(eC/C) = \sqrt{(eA/A)^2 + (eB/B)^2} > (eA/B) | (eB/B)$, where, e.g., eA denotes the uncertainty on A (measured in units of variable A). Note also that it is the relative error that matters for the RF analysis because we median subtract (and normalize by the interquartile range) all parameters to ensure a fair comparison (see Appendix A).

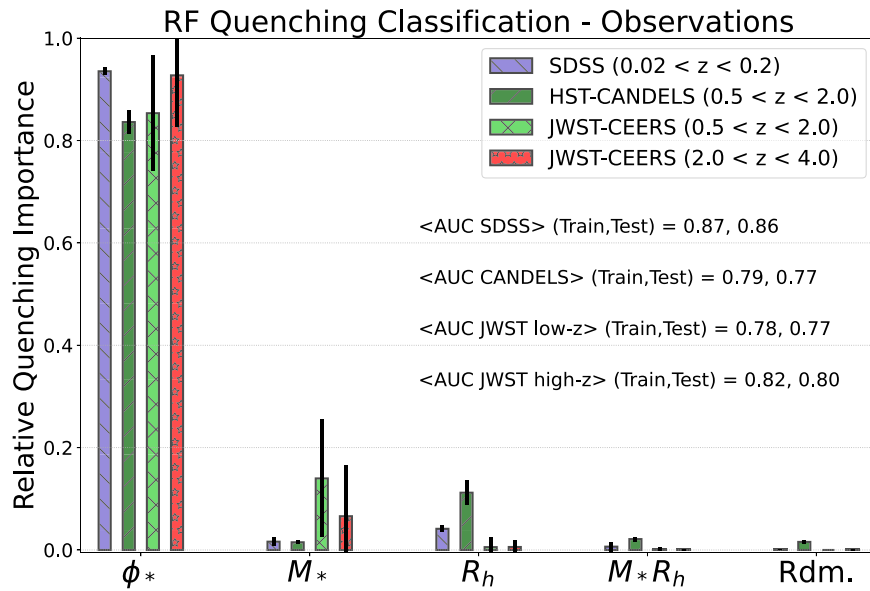


Figure 3. RF quenching classification comparing observational data sets targeting galaxies across ~ 13 Gyr of cosmic history ($z = 0-4$). The x-axis displays the parameters used to train the RF classifier, ordered from most to least predictive of quiescence (averaged across all epochs). The y-axis displays the importance of each variable for solving the classification problem at each epoch (see the legend). Uncertainties on the quenching importance are given from the variance of multiple independent RF classification runs, randomly sampling the data in each galaxy survey. At all epochs, the stellar potential (ϕ_*) is found to be overwhelmingly the most important parameter for predicting quiescence in galaxies. Stellar mass (M_*), galaxy half-light radius (R_h), and mass times radius (M_*R_h) are all either equivalent in importance to a random variable (Rdm), or else only marginally superior to it. This result clearly implies a stable quenching mechanism operating across cosmic time, which is closely linked to the galaxy’s gravitational potential. Moreover, this result is precisely as predicted by contemporary cosmological models utilizing AGN feedback to quench galaxies (compare to Figure 2 lower panel, and Bluck et al. 2023 for lower- z results). See Table 3 in Appendix A for full details on the RF hyperparameter setup.

M_*R_h are identical (each being the sum of the relative errors on M_* and R_h in quadrature). Yet, ϕ_* vastly outperforms M_*R_h . Therefore, these results cannot be explained, at any epoch, by measurement uncertainty.

Furthermore, the inclusion of M_*R_h in the list of parameters clearly demonstrates that an arbitrary (nonphysically motivated) combination of mass and radius performs worse than mass and radius individually, and no better than a random variable. Yet, a carefully chosen theoretically motivated combination (i.e., $\phi_* \sim M_{\text{BH}}$, according to simulations) performs as clearly the best predictor of quiescence in all galaxy surveys, and at all epochs. This motivates the, perhaps unusual, choice to include M_*R_h .

It is worth emphasizing this point a little further. An arbitrary combination of variables obscures the information contained within each component. For example, a phone number plus a random number is of no use for making contact. On the other hand, a well-chosen combination can reveal great insight. For example, the square root of a phone number multiplied by itself does, of course, enable contact. In terms of the problem at hand, ϕ_* ($=M_*/R_h$) is of great predictive value for identifying quiescence, yet M_*R_h is of no predictive power at all. Hence, one cannot claim that the superiority of ϕ_* in predicting quenching is a trivial consequence of it being based on M_* and R_h . Rather, the superiority of ϕ_* is a consequence of it being a *valuable* combination of M_* and R_h (as predicted by simulations, see Section 4.1). Nonetheless, this does not necessarily imply that the combination is optimal. To test this further we next explore a greater number of variables.

In Figure 4 we show an additional set of RF quenching classifications applied to the JWST-CEERS data. Here we expand further the list of parameters used to train the classifier, by including galaxy surface stellar mass density ($\Sigma_* = M_*/R_h^2$),

galaxy volumetric mass density ($\rho_* = M_*/R_h^3$), and Sérsic index (n_S). Additionally, we compare structural results in four wave bands, spanning the entire JWST-CEERS wavelength range. Note that in this figure we analyze the entire redshift range in JWST-CEERS together since no significant differences as a function of redshift are seen in Figure 3.

As in Figure 4, ϕ_* is clearly found to be the best parameter for predicting quiescence in JWST-CEERS in the larger parameter space, and at all wavelength ranges. The only exception to this is in the bluest band, where ϕ_* and M_* become indistinguishable. At the median redshift in JWST-CEERS, this band corresponds to rest-frame UV. As such it is no longer a good measurement of the underlying mass distribution in galaxies. Instead, it traces the location of bright, blue, massive young stars (which make up only a tiny fraction of the mass budget in essentially all massive galaxies).

As a result of the above, the reduction in the importance of ϕ_* in the bluest JWST wave band is actually the exception, which proves the rule. In the AGN feedback paradigm, the value of ϕ_* is that it measures the galaxy potential, which correlates strongly with black hole mass. However, once the half-light radius is no longer a good approximation of the half-mass radius (as in the bluest bands which are contaminated by young stellar populations), this is no longer true, and the link to black hole mass is lost. Indeed, if the results of Figure 4 were inverted, this would be exceptionally worrying, since it would suggest that the value in combining galaxy radius with stellar mass stems not from the importance of the gravitational potential but, rather, from a trivial tracing of young stars. Ultimately, the fact that this is emphatically not the case is very reassuring for the robust interpretation of these quenching results from JWST-CEERS.

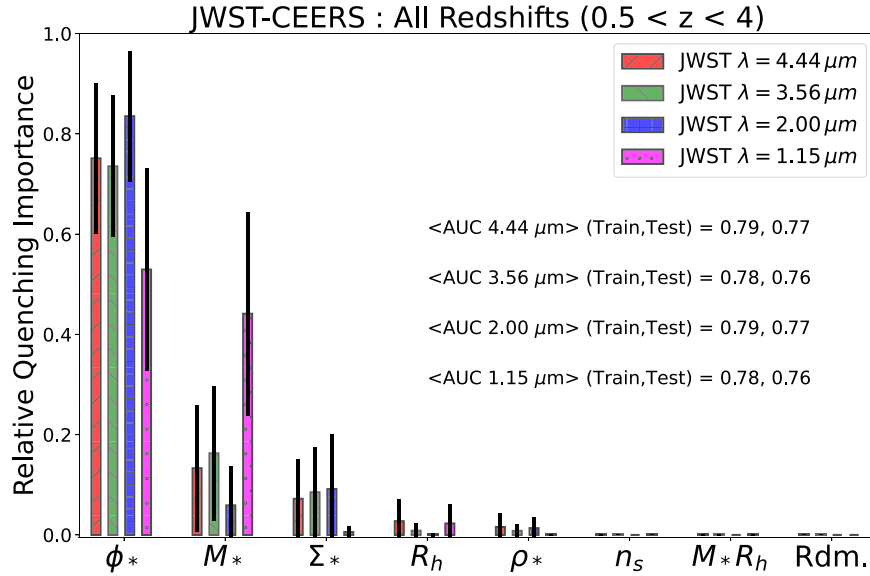


Figure 4. Extended RF classification of JWST-CEERS data. In this figure, we increase the number of observable parameters used to train the classifier in Figure 3, including the Sérsic index (n_s), surface density (Σ_*), and volume density (ρ_*). Additionally, we explore the impact of wavelength on the importance of the structural parameters to quenching (see the legend). Here we analyze the entire redshift range in JWST-CEERS since the importance of variables to quenching does not appear to change with cosmic time in Figure 3. The stellar potential remains the clear best parameter for predicting quiescence in the larger parameter set, when measured in red bands, but becomes equivalent in importance to stellar mass in the bluest band available. This strongly implies that the value of ϕ_* is connected to the mass distribution (traced at longer wavelengths). When this measurement is heavily biased by the distribution of bright star-forming regions (i.e., in the bluest band) it is less valuable for predicting quiescence. Note also that the uncertainties (measured from the variance of independent classification runs) increase with respect to Figure 3. This is an expected consequence of adding more closely intercorrelated variables, which provides more pathways to the segregation of galaxies into star-forming and quiescent classes. See Table 3 in Appendix A for full details on the RF hyperparameter setup.

In terms of interpretation, the results shown for simulations clearly show that a dominant dependence of quenching on ϕ_* is expected for AGN quenching, which ultimately depends on M_{BH} . There are no dynamical black hole mass estimates for our high- z samples. Hence, the best that can currently be done to test the AGN feedback paradigm at these epochs is to see how quenching depends on measurable photometric parameters. This provides clear indirect support for the AGN feedback paradigm (the phenomenology of Figures 2 and 3 are identical). Additionally, note again that a dependence of quenching on accretion rate (and hence AGN luminosity and detection) is not predicted to be evident in models. Therefore, the current approach is essentially the optimal methodology for testing the AGN feedback paradigm at high redshifts.

4.2.2. Area Statistics Check

The results from the RF classification analyses of JWST-CEERS galaxies (shown above) are extremely clear. The stellar potential, ϕ_* , is identified as the most important observable variable for predicting quiescence in galaxies at all epochs studied, spanning ~ 13 Gyr of cosmic history. These results are precisely as predicted in the AGN feedback-driven quenching paradigm for these photometric parameters (see Section 4.1, and Piotrowska et al. 2022; Bluck et al. 2023). However, given both the novelty and importance of the results of the previous subsection, it is sensible to check the validity of the RF classification conclusions via a simpler, and arguably more transparent, technique.

In Figure 5 we present the quenched fraction relationships with ϕ_* (top panels), M_* (lower-left panel), and galaxy half-light radius (lower-right panel). Results are shown for the highest JWST-CEERS redshift range used in the RF classification above ($2 < z < 4$). Each quenched fraction relationship is subdivided into separate relationships for percentile ranges of a

third variable. For example, in the top-left panel, the f_Q - ϕ_* relationship is displayed in percentile ranges of M_* (i.e., upper 25%, upper 50%, lower 50%, lower 25%). The lower-left panel inverts this, displaying the f_Q - M_* relationship in ranges of ϕ_* . Clearly, varying ϕ_* at fixed M_* engenders a far greater impact on the quenched fraction than the other way around, implying that the stellar potential must be more fundamentally linked to quiescence than stellar mass.

We quantify the tightness of each relationship using the area statistics approach (see Bluck et al. 2016, 2020a, 2022). We integrate the quenched fraction from the upper to the lower percentile range (for the upper and lower 50th and 25th percentiles in each secondary variable, labeled by the legend in each panel). The integration area is then normalized by the total area displayed in each panel, for a fair comparison. This method yields a relative tightness, dependent upon a third variable (since an absolute tightness is not obtainable for a simple fraction). This quantifies, e.g., the impact of varying M_* at fixed ϕ_* (top-left panel), and ϕ_* at fixed M_* (lower-left panel). In the right panels, the same process is applied to R_h and ϕ_* . As such, Figure 5 should be read column-wise to compare the relative impact of each pair of variables to quenching.

We estimate statistical uncertainties in the area statistics through bootstrapped random sampling of the data. Typical uncertainties are found to be ~ 0.02 . However, the value of the area statistics varies systematically with bin size, which we hold constant in all plots to remove this issue. Additionally, we have checked that the general trends (i.e., which parameter is the tightest) are invariant to reasonable bin size choices. Consequently, we conclude that the improvement in parameterization of quenching with ϕ_* over both M_* and R_h is highly significant, despite the increase in relative error on the former with respect to the latter (see above), and is analysis methodology independent.

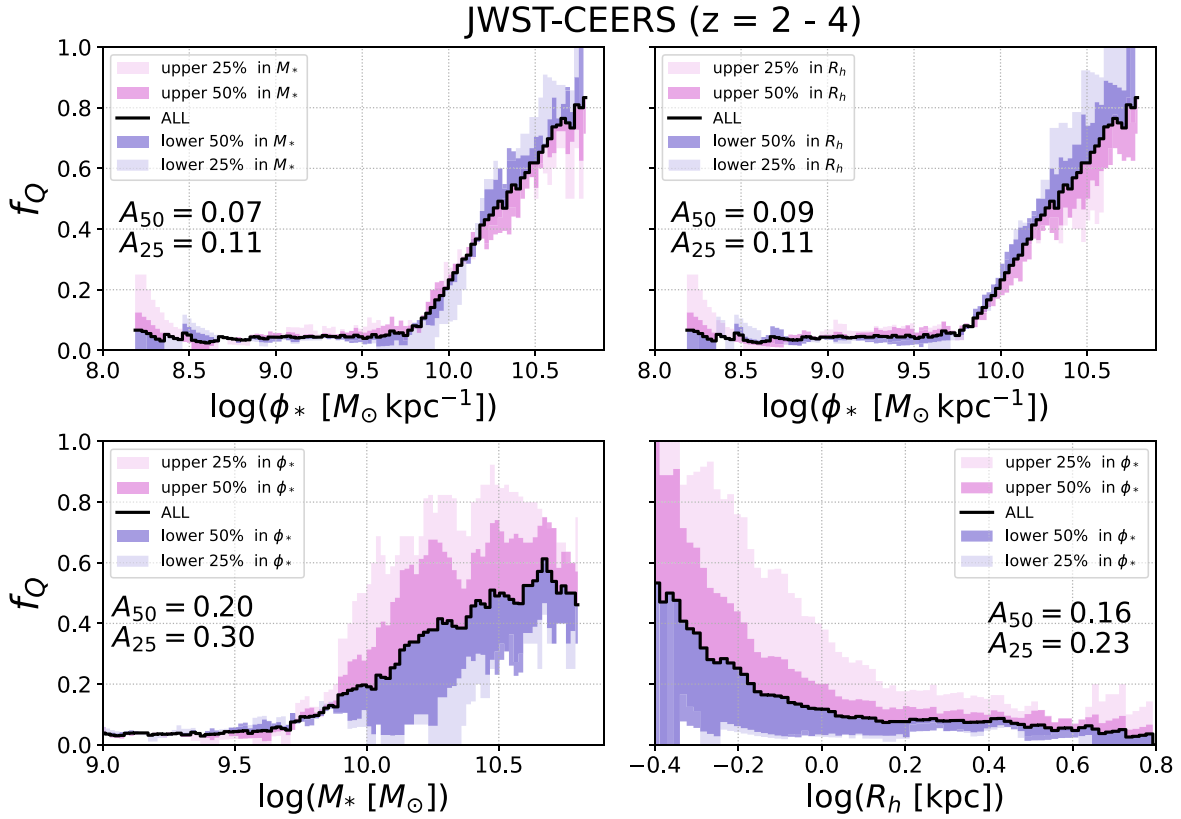


Figure 5. Area statistics analysis to check the validity of the RF results in JWST-CEERS at $z = 2-4$. The quenched fraction relationships are displayed as a function of stellar potential (top panels), stellar mass (lower-left panel), and galaxy radius (lower-right panel). Each quenched fraction relationship is subdivided into percentile ranges of a secondary parameter (as labeled by the legends). For example, the top-left panel displays the quenched fraction relationship with ϕ_* split into percentile ranges of M_* at each step. The bottom-left panel inverts this, showing the quenched fraction relationship with M_* split into percentile ranges of ϕ_* at each step. It is clear that varying ϕ_* at fixed M_* or R_h engenders far more variation in the quenched fraction than the other way around. This strongly indicates that ϕ_* is a more fundamental parameter for quenching (exactly as concluded in the RF classification analyses, see Figures 3 and 4). We quantify this effect with the area statistics (A_{50} and A_{25} , which show the area subtended between the upper and lower 50% and 25% of the data at each bin, respectively, and are displayed on each panel). Higher values of these statistics indicate less tight relationships.

From Figure 5, it is clear that the f_Q - ϕ_* relationship is much less impacted by either M_* or R_h , than the other way around. This clearly indicates that ϕ_* is the most fundamental (i.e., least impacted by other variables) parameter for driving quiescence in massive galaxies in JWST-CEERS data at very high redshifts, out of the list considered. Hence, the conclusions from the area statistics approach exactly match the conclusions found in the RF classification analyses above. Additionally, the area statistics approach yields consistent results with the RF classifications at lower redshifts as well (see Bluck et al. 2023).

At this point, one might wonder what the need for the RF classification is. The answer to this is twofold. First, the RF classifier can simultaneously control for an arbitrary number of parameters, whereas the area statistics approach can consider only two at a time. Second, the classification analysis explicitly solves the problem of identifying quiescent objects, whereas the area statistics approach only suggests an approach for doing so. In summary, the area statistics offer a simple, but useful, check on the results of the more sophisticated RF classification, but are not a substitute for it.

4.2.3. Quenching in the ϕ_* - M_* Plane

In Figure 6, we present the location of star-forming and quiescent galaxies from JWST-CEERS in the ϕ_* - M_* plane, shown separately for intermediate redshifts ($0.5 < z < 2$), high

redshifts ($2 < z < 4$), and ultrahigh redshifts ($4 < z < 8$). Note that this is the first time we introduce an analysis of ultrahigh redshift systems. The reason for this is that the number of quiescent galaxies is far too low in JWST-CEERS at these epochs for the previous statistical analyses (i.e., RF classification in Section 4.2.1; and area statistics in Section 4.2.2). As such, our conclusions for the ultrahigh redshift galaxies will be preliminary.

In all panels of Figure 6, the location of star-forming galaxies (identified as objects above the dashed orange line in Figure 1) are shown by blue density contours. The median relationship for star formers is shown by a solid blue line. Clearly, at all epochs there is a very strong relationship between stellar mass and stellar potential, such that more massive galaxies have deeper potential wells (with correlations, $\rho = 0.77$, 0.78 , and 0.91 for low, high, and ultrahigh redshifts, respectively). However, crucially for this analysis, there is physically meaningful scatter in the relationship owing to variation in galaxy size (and hence mass distribution) at a fixed stellar mass.

We also show the location of quenched (below the orange line in Figure 1) and fully quenched (below the red line in Figure 1) high- and low-mass galaxies on the ϕ_* - M_* plane in Figure 6 (color coded as described in the legends). Note that we include galaxies in the high-mass sample if $>30\%$ of their probability density function extends into the high-mass regime.

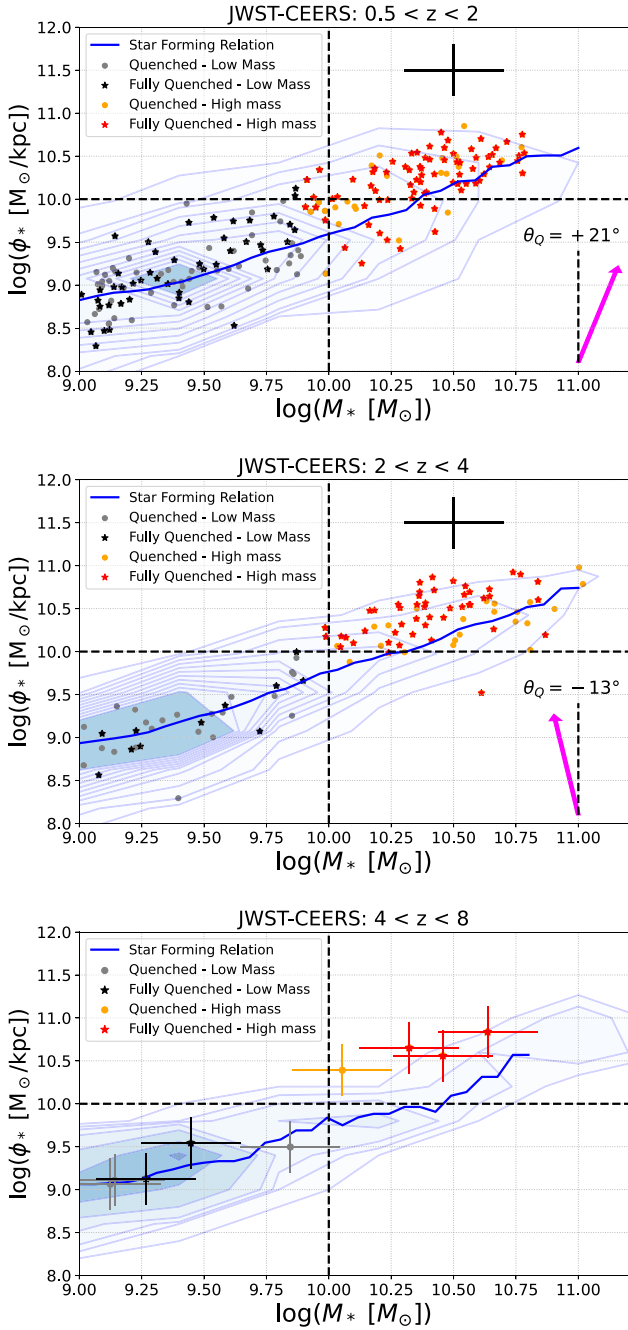


Figure 6. The stellar potential–stellar mass (ϕ_* – M_*) relationship in JWST-CEERS. Results are shown separately for intermediate redshifts (top panel), high redshifts (middle panel), and ultrahigh redshifts (bottom panel). In each panel, the locations of star-forming galaxies (above the orange dashed line in Figure 1) are shown by blue density contours. The median ϕ_* – M_* relationship for star formers is indicated by a solid blue line. There is clearly a very strong relationship at all epochs, whereby more massive galaxies have deeper gravitational potentials, with physically meaningful scatter engendered by varying galaxy size at a fixed mass. The location of high-mass fully quenched (red), high-mass quenched (orange), low-mass fully quenched (black), and low-mass quenched (gray) galaxies are shown by data points. The typical uncertainty on these measurements is indicated by the black error bars on the top two panels, and by individual error bars in the lower panel. At all epochs, high-mass quiescent systems are offset to higher ϕ_* values at fixed M_* , whereas low-mass quiescent galaxies are not significantly offset in this diagnostic relative to star formers. For intermediate and high redshifts, we quantify the effect for high-mass systems with the quenching angle (see text), which shows the optimal direction to move in a 2D plot in order to quench. Since $\|\theta_Q\| < 45^\circ$, we conclude that ϕ_* is more important to quenching than M_* (as in the RF and area statistics analyses above).

To guide the eye, we overlay dashed black lines to indicate separation into high and low stellar mass systems (located at $M_* = 10^{10} M_\odot$) and show a separation into deep and shallow potentials (located at $\phi_* = 10^{10} M_\odot \text{ kpc}^{-1}$). This separates the parameter space into four quadrants of interest for further discussion.

At low masses, quiescent objects trace the median star-forming relationship at all epochs. This suggests that there is no impact of varying ϕ_* at a fixed M_* for these systems. Utilizing the interpretation of simulations implies that low-mass quenching is unrelated to AGN feedback (as expected). Consequently, these quiescent objects are most likely the result of environmental quenching in the early Universe.

For the high-mass quenched galaxies in Figure 6, we see an offset in ϕ_* at fixed M_* , such that quiescent galaxies reside in deeper potential wells for their stellar masses than star-forming systems. This offset can be naturally explained in the AGN feedback paradigm by these systems hosting more massive central black holes, which will have injected more energy into their circumgalactic media. Interestingly, we see evidence of this offset all the way out to the most extreme redshift range ($z = 4$ –8).

As yet another test on the relative importance of stellar mass and stellar potential for quenching high-mass galaxies, in Figure 6 we plot the quenching angle (θ_Q). This is defined as the arctangent of the ratio of partial correlations, i.e., treating each partial correlation as a basis vector. Explicitly, we compute

$$\theta_Q \equiv -\tan^{-1} \left\{ \frac{\rho(\text{sSFR}/(1/t_H(z)) - \phi_* | M_*)}{\rho(\text{sSFR}/(1/t_H(z)) - M_* | \phi_*)} \right\}, \quad (2)$$

where the numerator indicates the rank-ordered partial correlation between normalized sSFR and stellar potential, at fixed stellar mass; and the denominator indicates the rank-ordered partial correlation between normalized sSFR and stellar mass, at a fixed stellar potential. The normalization of the sSFR is used to account for redshift evolution in the quenching threshold (see Figure 1), and the negative sign ensures that the angle points in the direction of increasing quiescence (rather than star formation). The partial correlations are given as standard by

$$\rho_{A-B|C} = \frac{\rho_{AB} - \rho_{AC} \cdot \rho_{BC}}{\sqrt{1 - \rho_{AC}^2} \sqrt{1 - \rho_{BC}^2}}, \quad (3)$$

for three arbitrary parameters, A , B , and C . For example, ρ_{AB} indicates the Spearman rank correlation between variables A and B , and $\rho_{A-B|C}$ indicates the partial correlation between A and B , controlling for C . See Bluck et al. (2020a) for a full explanation of the quenching angle, and applications at low- z . Ultimately, this statistic quantifies the optimal route through a 2D parameter space to a lower sSFR.

A quenching angle of 0° indicates pure ϕ_* dependence, an angle of 90° indicates full M_* dependence, with 45° indicating an equal contribution of both. Since $\|\theta_Q\| < 45^\circ$ in both panels where we have the statistics to run this test, we conclude that ϕ_* is more important to quenching than M_* in JWST-CEERS (as concluded in the RF and area statistics analyses above).

4.2.4. Distributions in Offsets from the ϕ_* - M_* Relation

In Figure 7, we present the distributions in the offsets of quiescent galaxies from the ϕ_* - M_* median relation (defined for star-forming objects, shown as blue solid lines in Figure 6). The left panels show results for high-mass quiescent galaxies, with the right panels showing results for low-mass quiescent galaxies. It is clear that high-mass quenched objects are offset to higher ϕ_* at fixed M_* , relative to star-forming galaxies. Conversely, there is no apparent offset in the low-mass quenched systems, relative to star-forming galaxies.

We quantify the significance and magnitude of the offsets (or lack thereof) via two statistics displayed on each panel of Figure 7: (i) the probability of being drawn from the same distribution via the Kolmogorov–Smirnov (K-S) test ($P_{\text{K-S}}$); and (ii) the difference in mean offset between star-forming and quenched populations, expressed in units of the combined error on the means of the two distributions ($\Delta\mu [\sigma_{\text{err}}]$).

High-mass galaxies have significant offsets from the star-forming ϕ_* - M_* relationship with confidence (probability of being spurious) ranging from 1/5000 (ultrahigh redshifts) to 1/10¹³ (high redshifts), with intermediate redshift galaxies being offset to a confidence of one in a million. Interestingly, the largest magnitude in offset is seen at the highest redshifts and systematically declines to lower redshifts. This hints at a possible trend whereby in order to transition to quiescence in the very early Universe, galaxies must have preferentially deeper potential wells for their stellar masses, compared to later epochs. Within the AGN feedback paradigm, this would be interpreted as these systems forming exceptionally massive central black holes for their stellar masses, leading to enhanced AGN feedback integrated over their histories (see Section 5.1). However, better statistics at the highest redshifts are needed to confirm these interesting preliminary results.

On the other hand, for low-mass quiescent galaxies, no significant offsets from the star-forming ϕ_* - M_* relationship are seen at any epoch. This strongly suggests that this low-mass quenching is unrelated to both stellar mass and gravitational potential, and hence, also likely unrelated to black hole mass and AGN feedback. These quiescent systems are almost certainly a result of environmental, rather than intrinsic, quenching mechanisms (see, e.g., Bluck et al. 2016, 2020b; P. Goubert et al. 2024, in preparation, for low- z analogs). Hence, we find evidence for both high-mass intrinsic quenching and low-mass environmental quenching in JWST out to the most extreme redshifts ever probed. However, note that our results are preliminary in the sense that we have a photometric-only sample with limited parameter availability, and have a modest sample size at the highest redshifts. Both of these limitations will be improved with upcoming JWST studies in the coming few years.

4.2.5. Evolution in the f_Q - ϕ_* Relation

In Section 4.2.1 we find that the best parameter for predicting quenching is stable throughout cosmic time (see back to Figure 3). However, this does not mean that there is no explicit dependence of quenching on redshift, only that any such evolution must leave ϕ_* as the most predictive parameter at any given epoch. To explore the possibility of redshift dependence on quenching directly, in Figure 8 we perform an RF classification analysis for the full JWST-CEERS range, comparing the dominant quenching parameter at each epoch

(ϕ_*) to redshift and a random number. It is clear that while ϕ_* remains the dominant parameter, there is a significant secondary dependence on redshift in the multi-epoch sample (which is much superior to random). This implies the need for evolution in the f_Q - ϕ_* relationship with cosmic time.

We test this expectation explicitly in Figure 9 (left panel), where we show the f_Q - ϕ_* relationship for JWST-CEERS galaxies in bins of redshift (as labeled by the legends). As predicted by the RF analysis, there is a significant evolution in the f_Q - ϕ_* relationship, such that higher ϕ_* values are required to reach the same level of quenching at earlier cosmic times. For comparison, we also show the redshift evolution in the (less fundamental) f_Q - M_* relation. Again, this shows marked evolution, such that higher masses are needed to quench galaxies at earlier epochs.

These results can also be explained elegantly within the AGN feedback paradigm. In simple models of structure formation, the accretion rate into dark matter haloes scales with the physical density of baryons in the intergalactic medium, i.e., proportional to $(1+z)^3$. Hence, at earlier cosmic times there is greater accretion into dark matter haloes, at a fixed total mass. Thus, in order to shut off this accretion (and prevent the hot gas atmosphere from collapse), greater energetic feedback from AGN will be required, leading to higher black hole mass quenching thresholds (and hence higher ϕ_* values) at earlier epochs. In turn, the ϕ_* - M_* relationship (Figure 6) requires that there will be redshift dependence on the f_Q - M_* relation as well, which is a noncausal reflection of the underlying evolution in the black hole quenching threshold. See Bluck et al. (2022) for an analytic derivation of this redshift dependence in a semi-analytic model.

5. Discussion

5.1. AGN Quenching in the Early Universe

A comparison of the results from Section 4.2 (observations) with Section 4.1 (simulations) clearly demonstrates that AGN feedback *can* explain the statistical properties of galaxy quenching across cosmic time. Indeed, the direct predictions from EAGLE and TNG are met precisely in observations across 13 Gyr of cosmic history, incorporating data from three major galaxy surveys: (i) JWST-CEERS; (ii) HST-CANDELS; and (iii) SDSS.

In both EAGLE and TNG, black hole mass is predicted to act as the most important variable for quenching at very early cosmic times, and remain stable throughout cosmic history to the present epoch (see Figure 2; and Piotrowska et al. 2022; Bluck et al. 2023). Crucially, black hole accretion rate (and hence AGN luminosity and detection) are not predicted to be strongly connected to quiescence at any epoch. This is in spite of the fact that AGN feedback is explicitly the cause of high-mass quenching in both simulations (see Crain et al. 2015; Schaye et al. 2015; Weinberger et al. 2017, 2018; Zinger et al. 2020). This implies that searching for the location of AGN on the star-forming main sequence is not a productive route to testing the AGN feedback paradigm, despite its ubiquity in the observational literature (see Piotrowska et al. 2022; Ward et al. 2022; Bluck et al. 2023 for further discussion).

The reason for this is that in order for a galaxy to remain quiescent for long periods of time, cold gas inflow into the system must be prevented. As such, it is not enough to strip a galaxy of its ISM. Hence, powerful energetic feedback (often




Figure 7. Distributions in offsets from the star-forming ϕ_*-M_* relationship ($\Delta(\phi_*-M_*)$). Offsets are measured relative to the star-forming median relationship at each epoch (blue solid lines in Figure 6). The left panels show a comparison of quenched galaxies to star-forming galaxies at high mass ($M_* > 10^{10} M_\odot$), with the right panels showing the same comparison for low-mass galaxies ($M_* < 10^{10} M_\odot$). The top row displays results for intermediate redshifts, the middle row for high redshifts, and the bottom row for ultrahigh redshifts. It is clear that quiescent galaxies are offset to higher ϕ_* at high masses, relative to star formers at a fixed stellar mass. Additionally, it is also clear that there is no significant offset for low-mass systems. We quantify these effects with two statistics (displayed on the individual plots). First, we quantify the significance of the difference between each pair of distributions by the Kolmogorov–Smirnov p -value (P_{KS}). Second, we display the difference in mean values of each pair of distributions, in units of the combined error on the mean ($\Delta\mu$ [σ_{err}]). We note a trend whereby high-mass quiescent galaxies become more offset to higher ϕ_* at higher redshifts (relative to star formers at a fixed stellar mass).

referred to as the *quasar mode*; see Hopkins et al. 2006, 2008) is not sufficient to ensure a fully quenched system. Additionally, even without ISM removal, a galaxy can still quench if gas inflows are permanently switched off. In this scenario, galaxies will continue to transform their ISM into stars, and quench only once it is fully depleted. This latter (slow) quenching leads to the prediction that the stellar metallicities ought to rise during the transition to quiescence (which is confirmed in Peng et al. 2015; Bluck et al. 2020b; Trussler et al. 2020).

Therefore, high Eddington ratio, energetic AGN (which are by far the easiest type of AGN to identify) are neither necessary nor sufficient to quench massive galaxies. The essential condition of massive galaxy quenching in modern simulations is the formation of low Eddington ratio AGN, which acts over cosmological times, preventing cooling of the CGM, and hence, gas accretion into the system. These conditions occur when black hole masses become very large, and the predominant accretion mechanism switches from the cold to

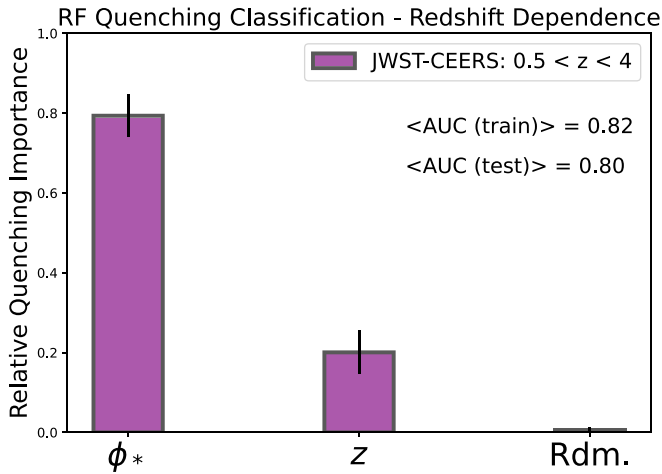


Figure 8. RF classification analysis to predict the existence of quenched galaxies in the full JWST-CEERS redshift range, utilizing only ϕ_* , z , and a random variable. It is clear in this analysis that high-mass galaxy quenching has a significant secondary dependence on redshift, even after controlling for the dominant parameter (ϕ_*). This suggests an evolution in the f_Q - ϕ_* relationship (see Figure 9).

hot mode (e.g., Sijacki et al. 2007; Fabian 2012; Weinberger et al. 2018; Terrazas et al. 2020; Zinger et al. 2020; Piotrowska et al. 2022). Observationally, this mode of AGN is extremely hard to detect. However, their long-term effects on the hot gas halo are clearly seen in galaxy clusters (see Hlavacek-Larrondo et al. 2015, 2018; Brownson et al. 2019; Jones et al. 2023). Additionally, low luminosity core radio detection is identified preferentially in high-mass quenched galaxies in the local Universe (see, e.g., Hickox et al. 2009; Heckman & Best 2014), which is obviously consistent with this paradigm.

The total energy released into galactic haloes through long-term, low luminosity AGN heating in the *preventative mode* is directly proportional to the mass of supermassive black holes (see Bluck et al. 2020a for a modern derivation; and Soltan 1982; Silk & Rees 1998 for the original arguments). Moreover, it is completely uncorrelated with the current AGN accretion rate, AGN bolometric luminosity, and hence, AGN detection in any given wave band. Ultimately, this reasoning explains the phenomenological results from simulations in Section 4.1, and in Piotrowska et al. (2022), Bluck et al. (2023), and Ward et al. (2022). Contemporary AGN feedback models explicitly predict that quiescence ought to intrinsically depend on black hole mass, not accretion rate or luminosity. As such, the correct method to test the AGN feedback paradigm is *not* to look for the location of AGN on the star-forming main sequence (as many prior observational works have done), but rather to look for a connection between quiescence and the fossil record of past accretion, i.e., black hole mass (or its best available proxies).

In Bluck et al. (2023), we find that the stellar gravitational potential is predicted by several simulations to act as an excellent proxy for central black hole mass. In Section 4.1, we find that this remains true up to the earliest cosmic times (when simulations cease to predict the existence of quiescent systems). This provides a route to test the paradigm of AGN feedback quenching at early epochs, using photometric imaging data from JWST-CEERS, in addition to comparisons at lower redshifts with HST-CANDELS and SDSS. Interestingly, recent papers have found that dynamically measured black hole mass does not correlate strongly with stellar mass in

the early Universe, but is still closely connected to central velocity dispersion, and hence, the galaxy potential (see Maiolino et al. 2023; Pacucci et al. 2023).

In Section 4.2, we find that the stellar potential, ϕ_* , is unambiguously the best predictor of high-mass galaxy quenching at very early cosmic times, exactly as predicted by simulations. Furthermore, we find that the predictive power of ϕ_* is completely stable over 13 Gyr of cosmic history (see Figure 3). This is in remarkable accord with predictions from the AGN feedback paradigm. Moreover, independent of the likely physical mechanism(s), this result demands a stable quenching process operating throughout cosmic time (see Bluck et al. 2022, 2023 for a similar conclusion at $z < 2$).

In Section 4.2, we also find that ϕ_* outperforms a host of other parameters in JWST-CEERS data, including stellar mass, galaxy size, galaxy 2D and 3D density, and Sérsic index (among others). These results are naturally explained by the AGN feedback paradigm. Supermassive black holes form and grow in dense potential wells (e.g., Sijacki et al. 2007, 2015; Weinberger et al. 2017). Once they reach a critical mass (typically $M_{\text{BH}} \sim 10^{7-8} M_{\odot}$; e.g., Weinberger et al. 2018), they transition to hot-mode accretion, form a stable central engine, and keep hot gas haloes stable from cooling and collapse for many gigayears. This process simultaneously resolves the cooling problem in galaxy groups and clusters (e.g., Fabian 2012); the cosmological low efficiency of star formation (e.g., Fukugita & Peebles 2004), and the existence and demographics, of quiescent galaxies (e.g., Peng et al. 2010, 2012; Bluck et al. 2014).

5.1.1. Alternative Explanations to AGN

We suspect that the AGN feedback paradigm is indeed the best explanation for the multi-epoch quenching results of the previous section. Certainly, all results we have found from JWST-CEERS, HST-CANDELS, and SDSS are consistent with its predictions. Nevertheless, it is still worth considering if any other paradigm for massive galaxy quenching could also be consistent with our results.

An important early hypothesis for high-mass galaxy quenching was to evoke galaxy–galaxy mergers. In such a scenario the ISM of the merging systems will be compressed leading to enhanced star formation, followed by a depression in star formation as a result of depleted gas reservoirs (e.g., Hopkins et al. 2006, 2008). Additionally, energetic feedback from quasar-mode AGN could aid in the depletion of the ISM by driving powerful outflows, ultimately triggering quenching. However, without some other form of long-term prevention of gas inflows (either from cold flows in low-mass haloes, or from cooling of the CGM in high-mass haloes) into galaxies would inevitably rejuvenate, and no long-term quiescence would be apparent. This issue is brought into sharper focus by considering the *cooling catastrophe*, whereby the CGM around massive galaxies, groups, and clusters ought to be thermodynamically unstable, leading to runaway cooling, yet is observed to be stable across cosmic times (see, e.g., Fabian et al. 2006; Fabian 2012). As such, some heating source is required to maintain long-term quiescence.

Given the logic of the cooling catastrophe (see Fabian 2012 for a review), a heating source is required to keep hot gaseous haloes buoyant, shielding galaxies from cold inflows (e.g., Dekel & Birnboim 2006), and ultimately enabling quiescence through starvation of gas needed as fuel for future star

JWST-CEERS: Evolution in Quenched Fraction




Figure 9. Evolution in the quenched fraction relationship with stellar potential (left panel) and stellar mass (right panel). Solid lines show the quenched fraction as a function of ϕ_* and M_* , with shaded regions showing the 1σ uncertainty in the quenched fraction (colored by epoch as displayed in the legends). Uncertainties are inferred from Poisson counting statistics. It is clear that there is a strong positive trend of quenched fraction with both stellar mass and stellar potential, which evolves with redshift such that higher values of both ϕ_* and M_* are required at earlier epochs in order to quench.

formation (e.g., Peng et al. 2015; Bluck et al. 2020b). Astrophysically, there are a finite number of possibilities. The main three are (i) AGN feedback (as discussed at length above, see Somerville & Dave 2015 for a review); (ii) stellar and supernova feedback (e.g., Cole et al. 2000; Henriques et al. 2015); and (iii) virial shocks (e.g., Dekel & Birnboim 2006; Woo et al. 2013). Additionally, more exotic alternatives could be speculated to be important, such as dark matter annihilation, magnetic field heating, and cosmic rays. However, the latter three are not well calibrated, have weak-to-no observational evidence, or are not predicted by cosmological models to impact the star formation efficiency problem. As such, we restrict the following discussion to the three numbered sources above.

In Bluck et al. (2020a), we demonstrate that supernova (or more generally stellar) feedback must scale fundamentally with stellar mass, as the time integral of SFR. In Section 4.2 (see also Bluck et al. 2014, 2016, 2020a, 2020b; Brownson et al. 2022; Piotrowska et al. 2022; Bluck et al. 2023), we find that stellar mass is of no intrinsic importance for galaxy quenching at any epoch. Indeed, the importance of stellar mass tends to that of a random variable, once more fundamental parameters are included (see Figures 3 and 4). This emphatically rules out supernovae as the fundamental cause of high-mass galaxy quiescence. Indeed, this result is entirely as expected in modern cosmological models, which introduced AGN feedback explicitly to deal with the failure of supernovae to quench galaxies (see Bower et al. 2006; Croton et al. 2006; Sijacki et al. 2007; Bower et al. 2008).

Also in Bluck et al. (2020a), we show that the heating induced by virial shocks scales as a power law with halo mass (consistent with the presentation in Dekel & Birnboim 2006). Hence, halo mass is the key observable of virial shock quenching, which is why it has been dubbed *halo mass quenching* (see Woo et al. 2013, 2015). Although we do not

have access to reliable estimates of halo masses in the early Universe yet, at lower redshifts we have found that halo mass is of no greater predictive power over quenching than a random variable, once an estimate of black hole mass is included (see Bluck et al. 2016, 2020a, 2022; Piotrowska et al. 2022). This rules out halo mass quenching in the local Universe. Moreover, the extreme stability of the photometric parameters’ quenching importance (see Figure 3) strongly indicates that there is no variation in quenching mechanism across cosmic time. Hence, we *predict* that halo mass will not be important at high redshifts, exactly as in the local Universe. This prediction may be tested relatively soon, using abundance matching in the wide-field spectroscopic VLT-MOONRISE survey (Maiolino et al. 2020).

One final possibility is that the strong dependence of quenching on the galactic potential (traced by ϕ_*) in observations could be a result of dynamical stabilization (e.g., Martig et al. 2009; Gensior et al. 2020). In such a scenario, dynamical torques from a central mass distribution could provide stabilization against the gravitational collapse of the ISM in disk structures. Consequently, the gas content of galaxies is predicted to *rise* during quenching because in quiescent systems gas conversion into stars halts, yet gas accretion into the system (either from cold gas accretion or CGM cooling) continues. However, in the local Universe, the opposite is seen, whereby quiescent systems have significantly lower gas fractions than their star-forming counterparts (e.g., Saintonge et al. 2017, 2018; Brownson et al. 2020; Ellison et al. 2020, 2021a, 2021b; Piotrowska et al. 2020). Hence, this cannot be the dominant quenching mechanism at low redshifts.

At higher redshifts, it remains to be seen whether the gas content of quiescent objects is lower than star-forming systems. However, long-term quiescence cannot be maintained without shutting off gas inflow into the system (as demonstrated by simulations, see, e.g., Vogelsberger et al. 2014a, 2014b; Schaye et al. 2015;

Nelson et al. 2018). No amount of dynamical stabilization could prevent the vast quantities of gas accreted into galaxies from forming stars, if cooling from the CGM is not prevented (see Fabian 2012 for a discussion). Hence, we reject the dynamical stabilization possibility as a probable quenching mechanism for high-mass galaxies across cosmic time.

We remain agnostic about the possibility of future proposed heating sources driving galaxy quiescence. However, any such mechanism must simultaneously explain (i) why ϕ_* is the best photometric predictor of quiescence; (ii) why this feature is completely stable across cosmic time; and (iii) why the observed long-term heating from AGN (e.g., Fabian 2012; Hlavacek-Larondo et al. 2015, 2018), which is consistent with the two criteria above, does not stabilize the CGM around massive galaxies and hence requires another quenching source. Applying Occam’s razor, it seems that the AGN feedback paradigm is the most probable quenching mechanism for massive galaxies across the entire history of the Universe. Furthermore, this conclusion is in perfect accord with the current (near-consensus) theoretical view.

5.2. Environmental Quenching in the Early Universe

In this paper, we have focused primarily on massive galaxy quenching ($M_* > 10^{9.5} M_\odot$ systems). As such, we have sought to test the AGN feedback paradigm across cosmic time, with the main novelty being the high and ultrahigh redshift regimes accessible for the first time with JWST-CEERS. However, in Section 4.2.3 we expand our view of galaxy quiescence to lower stellar masses, and hence, are sensitive to other quenching mechanisms.

In Figures 6 and 7 we see two distinct populations of quiescent galaxies, both existing at all epochs studied (from $z = 0.5$ –8). For high-mass galaxies ($M_* > 10^{10} M_\odot$), we see clear evidence of quenched systems having higher ϕ_* , at fixed M_* , than their star-forming counterparts. This is completely consistent with the AGN feedback paradigm (see Section 5.1). Additionally, we find a population of quenched low-mass galaxies ($M_* < 10^{10} M_\odot$) that are *not* offset in ϕ_* , at fixed M_* . Moreover, since these systems are defined to have low stellar masses, this population of quiescent galaxies has their quenching uncorrelated with both stellar mass and central potential. Consequently, this type of quenching is likely uncorrelated with black hole mass, and hence, AGN feedback.

In the local Universe, a similar population of low-mass (and shallow gravitational potential) quiescent objects are identified, which are found to have their quiescence critically connected to environment (e.g., Peng et al. 2010, 2012; Bluck et al. 2014, 2016; P. Goubert et al. 2024, in preparation). An analogous population is also identified out to $z \sim 2$ by, e.g., Lang et al. (2014). As such, it is logical to speculate that the low-mass quenched galaxies observed out to $z \sim 8$ in JWST-CEERS are similarly environmentally quenched systems. We also note that the simulations used in this work (EAGLE and TNG) both predict the existence of low-mass quenched systems, uncorrelated with black hole mass, which are predominantly satellites (see P. Goubert et al. 2024, in preparation, for an analysis of satellite quenching at early cosmic times).

At present, it is not at all straightforward to quantify the environments of JWST-CEERS galaxies. The primary limitation is the lack of spectroscopic redshifts for a large contiguous galaxy sample at high- z , which is required for accurate group

and cluster identification. In the coming ~ 5 yr the lack of wide-field spectroscopic surveys targeting the early Universe will be addressed with VLT-MOONRISE (Maiolino et al. 2020), as well as various planned JWST spectroscopic surveys. As such, we take the opportunity to make a prediction: the observed low-mass quenched galaxies in JWST-CEERS will turn out to be environmentally quenched objects, as evidenced by a strong correlation between their star formation and galaxy environmental metrics (such as local density, central satellite categorization, and halo properties).

Using the rationale above, we suspect that we have found (indirect) evidence for environmental quenching at the earliest cosmic times ever probed. Testing this hypothesis will likely become a high priority of the next generation of high- z spectroscopic galaxy surveys.

6. Summary

In this paper, we start by extracting the testable predictions for AGN feedback-driven quenching at very high redshifts from the EAGLE and TNG cosmological simulations. We then focus on testing these predictions using data from JWST-CEERS data in the very early Universe, complementing this analysis with observational tests at later cosmic epochs utilizing HST-CANDELS and the SDSS.

Throughout the paper, we use a consistent definition of quiescence in galaxies, defined as sSFR values lower than the inverse of the Hubble time at each redshift (see Figure 1). We also test the stability of our results against a more stringent quenching threshold (a factor of 3 lower than this). The only exception to this general approach is at $z \sim 0$ (probed by SDSS), where the inverse Hubble time ceases to be a viable threshold (see Bluck et al. 2023). In this case, we adopt the commonly used threshold at $\text{sSFR} < 10^{-11} \text{ yr}^{-1}$ (as in Piotrowska et al. 2022; Bluck et al. 2023).

To analyze both the observational and simulated galaxy surveys, we utilize a sophisticated statistical approach with RF classifications, segregating star-forming and quiescent galaxies. We have demonstrated that our RF classification technique is extremely effective at identifying causal relationships in complex intercorrelated data (see Bluck et al. 2022; Brownson et al. 2022; Piotrowska et al. 2022).

Our primary results are as follows:

- (1) Both TNG and EAGLE predict that supermassive black hole mass should be the key predictor of quiescence in the early Universe at $z = 2$ –4 (see Figure 2, top panel). This is identical to the predictions from these simulations at lower redshifts, which is indicative of a single quenching mechanism operating throughout the history of the Universe in these models (see Piotrowska et al. 2022; Bluck et al. 2023 for results at $z < 2$).
- (2) Both TNG and EAGLE also predict that, in lieu of supermassive black hole mass, the stellar potential ($\phi_* = M_*/R_h$) will outperform stellar mass, and several other parameters, in predicting quenching (see Figure 2, bottom panel).
- (3) In JWST-CEERS, at both intermediate ($0.5 < z < 2$) and high ($2 < z < 4$) redshifts, we find that ϕ_* is unambiguously identified as the best parameter for predicting quiescence in galaxies, from an RF classification. Once ϕ_* is accounted for, there is essentially no importance of quenching contained in the stellar mass, galaxy size,

stellar surface (or volume) density, or several other galaxy structural parameters (see Figures 3–5).

- (4) Through comparison with HST-CANDELS and SDSS, we find that the key parameter for predicting high-mass galaxy quiescence is stable across 13 Gyr of cosmic history. This indicates a very high probability of a single high-mass galaxy quenching mechanism operating throughout all of cosmic time (see Figure 3).
- (5) We identify two distinct populations of quiescent galaxies at high redshifts in JWST-CEERS: (i) high-mass quenched galaxies, which have significant offsets to higher ϕ_* at a fixed M_* (as expected in the AGN feedback paradigm); and (ii) low-mass quenched galaxies, which have no offsets in ϕ_* (see Figures 6 and 7). Through analogy with low- z observations, we speculate that the latter population is likely a result of environmental quenching in the early Universe (see Section 5.2).
- (6) Finally, we find that, even though the importance of ϕ_* for predicting quenching remains stable across cosmic time, the f_Q - ϕ_* relation does evolve, such that higher values of ϕ_* are required in order for galaxies to quench at earlier cosmic times (see Figures 8 and 9).

Taken as a whole, the observational results from this paper represent a critical test of the AGN feedback paradigm, spanning over 13 Gyr of cosmic history. The direct predictions from this paradigm (as given by cosmological simulations) are in remarkable agreement with the observations. This strongly suggests that AGN feedback can explain high-mass galaxy quenching at all epochs in the history of the Universe. We emphasize that a more *direct* test utilizing a putative correlation between sSFR and AGN luminosity is *not* predicted to exist by contemporary AGN feedback models (see Section 5.1 for a discussion on this crucial point). Hence, the methodology of this work is essentially the optimal strategy to test the AGN feedback paradigm across cosmic time (at least in photometric data).

We also identify the need for low-mass galaxy quenching in JWST-CEERS, which is incompatible with AGN feedback (since it occurs in galaxies that, almost certainly, host low-mass central black holes). In the local Universe, low-mass quenching is nearly exclusively associated with environment. Hence, our high- z results suggest that environment may already suppress star formation for some low-mass systems as early as 1–2 Gyr after the Big Bang. Further exploration of this population will require wide-field spectroscopic galaxy surveys targeting the early Universe, which will be realized in the next ~ 5 yr with VLT-MOONRISE, and various upcoming JWST surveys.

Data Access

All of the data used in this paper is publicly available, or else currently in preparation for public release.

HST+JWST-CEERS data: for access to photometric redshift and stellar mass catalogs see Duncan et al. (2019). For access to structural catalogs see Ormerod et al. (2024). Or, for immediate requirements, message the corresponding author with reasonable requests. Full details on the JWST observations and source data are provided at <https://ceers.github.io/>

HST-CANDELS source data: <https://archive.stsci.edu/hls/candels>. VAC: <https://mhurtascompany.weebly.com/data-releases-and-codes.html>

SDSS source data: <https://classic.sdss.org/dr7/access/>. SFR catalog: see Brinchmann et al. (2004). Morphological catalog: see Simard et al. (2011). Stellar mass catalog: see Mendel et al. (2014).

Simulations: EAGLE data access <http://icc.dur.ac.uk/Eagle/>. IllustrisTNG data access www.tng-project.org/. Docker access <https://hub.docker.com/u/jpiotrowska>

Analysis: all analyses in this work were performed using PYTHON-3 (<https://www.python.org/>), including NumPy, Astropy, SciPy, pandas, Seaborn, and matplotlib. All machine-learning analyses were performed using scikit-learn (Pedregosa et al. 2011; see <https://scikit-learn.org>).

Acknowledgments

We thank the anonymous referee for a constructive report, which helped us improve the presentation of our work. We thank all contributors to JWST-CEERS, HST-CANDELS, SDSS, the EAGLE Collaboration, and the Illustris and IllustrisTNG collaborations. This paper would not have been possible without the work done by hundreds of scientists, engineers, and support staff over the past decade. We also thank the teams behind the analysis tools we have used from PYTHON, especially those involved in `scikit-learn`.

A.F.L.B. gratefully acknowledges a faculty start-up grant at the Florida International University. C.J.C. acknowledges support from the ERC Advanced Investigator grant EPOCHS (788113), as well as studentships from STFC. N.J.A. acknowledges support from the ERC Advanced Investigator grant EPOCHS (788113). D.A. and T. H. acknowledge the support of studentships from the STFC. K.J.D. acknowledges support from the STFC through an Ernest Rutherford Fellowship (grant No. ST/W003120/1). R.M. acknowledges a Royal Society Research Professorship, as well as support from the STFC and ERC Advanced grant (695671) “QUENCH.”

Appendix A

Machine-learning Reproducibility

In order to enable accurate reproduction of our machine-learning results by other researchers, we provide in this appendix a series of tables listing the hyperparameters used in the RF classification analyses (for Figures 2–4 see Tables 2–4, respectively).

The table rows indicate each classification run, which all categorize star-forming and quiescent galaxies via a threshold set at the inverse Hubble time (or $\text{sSFR} < 10^{-11} \text{yr}^{-1}$ for SDSS at $z=0.1$). All classifications are run using `scikit-learn` (Pedregosa et al. 2011; see <https://scikit-learn.org>).

In every classification run, we set `max_features` = “None.” This ensures that every feature is available at every node in each decision tree within the RF. In Bluck et al. (2022), we have demonstrated that this mode represents the most effective RF classification architecture for disentangling intercorrelated *nuisance* parameters from the underlying causal relations in simple models, simulations, and observations. Note that this is *not* the default method of RF classification in `scikit-learn`. However, it is the most effective for revealing physical insight.

The table columns are as follows:

Table 2
RF Hyperparameters for Figure 2

	Method	N_{trees}	Max Depth	$N_{\text{iterations}}$	Balanced	Normalized	Train:Test	MSL	AUC	ΔAUC
EAGLE (ALL)	RFclass	250	250	10	Yes	Yes	70:30	60	0.95	0.02
EAGLE (Obs)	RFclass	250	250	10	Yes	Yes	70:30	60	0.86	0.01
IllustrisTNG (ALL)	RFclass	250	250	10	Yes	Yes	70:30	25	0.98	0.02
IllustrisTNG (Obs)	RFclass	250	250	10	Yes	Yes	70:30	25	0.95	0.02

Table 3
RF Hyperparameters for Figure 3

	Method	N_{trees}	Max Depth	$N_{\text{iterations}}$	Balanced	Normalized	Train:Test	MSL	AUC	ΔAUC
JWST ($2 < z < 4$)	RFclass	250	250	100	Yes	Yes	70:30	30	0.87	0.02
JWST ($0.5 < z < 2$)	RFclass	250	250	100	Yes	Yes	70:30	50	0.79	0.02
HST-CANDELS	RFclass	250	250	25	Yes	Yes	50:50	110	0.78	0.02
SDSS	RFclass	250	250	10	Yes	Yes	50:50	1000	0.82	0.02

Table 4
RF Hyperparameters for Figure 4

	Method	N_{trees}	Max Depth	$N_{\text{iterations}}$	Balanced	Normalized	Train:Test	MSL	$\langle\text{AUC}\rangle$	$\langle\Delta\text{AUC}\rangle$
JWST (λ)	RFclass	250	250	100	Yes	Yes	70:30	65	0.79	0.02

Method. Type of machine learning used. Here always RandomForestClassifier from `scikit-learn` (abbreviated to RFclass in the tables).

N_{trees} . Number of decision trees used in the RF. Here set universally to 250. We have tested that increasing the number of trees does not further improve the results of any of these classifications.

Max Depth. Number of decision forks permitted within each decision tree. This is set universally to 250, which is deeper than ever required (given the MSL limit, see below).

$N_{\text{iterations}}$. The number of times we run each entire RF classification, using different randomly selected subsets of the data for training and testing. These iterations are used to infer statistical uncertainties on the feature importance. Note that this number varies from 10 to 100, to account for very different data sizes (and hence run times and sensitivity to random sampling).

Balanced. Whether the training and validation data has an equal number of star-forming and quenched galaxies within it. We follow standard best practices by evenly sampling each class in every RF classification run (see Teimoorinia et al. 2016 for a discussion).

Normalized. Whether the training data are normalized prior to being given to the classifier. In this work we always median subtract, and normalize by the interquartile range, every feature in the classifier. This ensures a fair comparison between data with different numerical values and distribution shapes (see Bluck et al. 2022 for further discussion).

Train:Test. The ratio of data sizes used in training and testing. In general, a 50:50 split is ideal. However, for smaller data sets it is recommended to weight training more heavily. As such, we adopt 70:30 for some of the runs. We have tested that this leads to more stable results (lower errors) and higher overall performance in those cases.

Min-Samples-Leaf (MSL). This hyperparameter sets the minimum number of data in a given node of a given decision tree needed to attempt further splits in the data. Setting this value to its logical minimum of two enables perfect classification of the training data. However, this leads to overfitting. We

adjust this parameter to simultaneously maximize the performance in training (AUC, penultimate column) while maintaining a minimal difference in performance with unseen testing data (ΔAUC , final column). Explicitly, in this work we require $|\Delta\text{AUC}| < 0.02$, yielding extremely similar performance in unseen data to the training sample, thus solving the overfitting problem.

AUC. The area under the true positive–false positive receiver operator curve, which quantifies the overall performance of the RF classification (see Teimoorinia et al. 2016). A value of $\text{AUC} = 0.5$ indicates a pure random performance, whereas a value of $\text{AUC} = 1$ indicates a perfect segregation of star-forming and quiescent systems. Note that in the table the AUC is given for training data.

ΔAUC . The difference in AUC between training and testing data. Following Bluck et al. (2022), we adopt a maximum threshold of 0.02 to count as a successful fit, i.e., free of significant overfitting.

We emphasize that the parameters shown in the tables indicate the optimal settings for each classification run. In addition to these runs, we have explored a wide hyperparameter space, involving significant variation in essentially all of the variables discussed above. All results in this paper are extremely stable to these alternative modes of classification, within reasonable bounds. As such, the results presented in the main body of the paper are easily obtainable without carefully matching our prescriptions. Nonetheless, to fully reproduce our results the above settings should be used.

Appendix B Predicted $M_{\text{BH}}-\phi_*$ Relation in Simulations at Early Cosmic Times

In Section 4.1, we show that the fundamental parameter regulating central galaxy quenching in cosmological simulations is supermassive black hole mass (see Figure 2, upper panel). However, we do not have access to black hole masses for galaxies in the early Universe for the photometric surveys

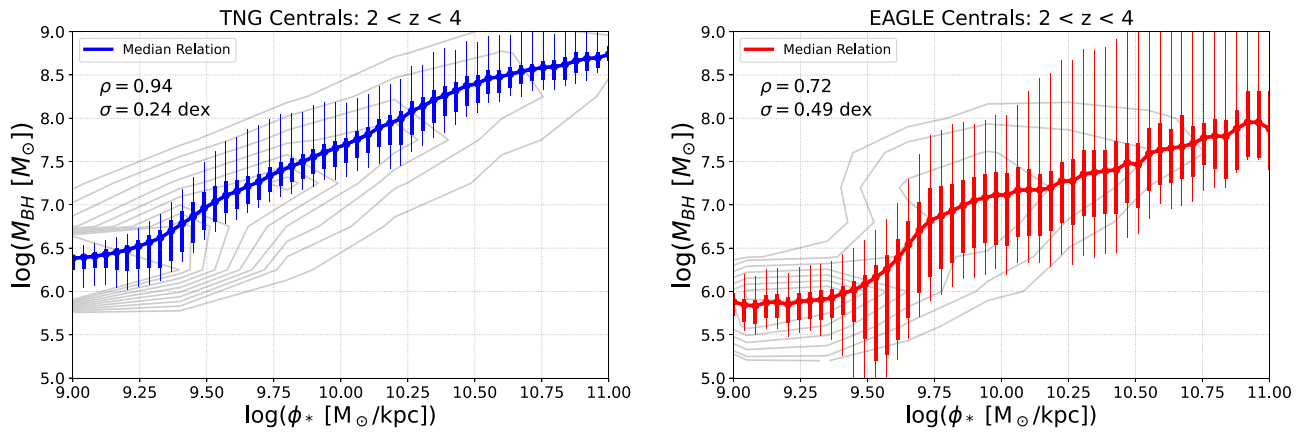


Figure 10. The relationship between supermassive black hole mass and stellar gravitational potential in simulations of the early Universe, shown separately for TNG (left) and EAGLE (right). The solid lines indicate the median relationship, with solid error bars indicating the 1σ dispersion, and light error bars indicating the 2σ dispersion. Linearly spaced density contours are overlaid in gray. Additionally, on each panel we present the Spearman rank correlation coefficient (ρ), and the average galaxy-weighted dispersion from the median relation (σ). It is clear that both simulations predict a strong relationship between M_{BH} and ϕ_* , which is leveraged in this work to test the integrated effect of AGN feedback at early cosmic times in photometric data. However, note that this relationship exhibits a stronger correlation and tighter dispersion in TNG compared to EAGLE.

used in this work. Furthermore, accurate dynamical measurements for black hole masses will likely only ever become available for less than 10% of galaxies at these epochs (i.e., for systems that present as powerful AGN with the appropriate orientation of molecular torus to measure broad line regions). Hence, dynamical measurements of black hole masses for early galaxies will never yield a fully representative sample, which is needed for statistical analysis of galaxy demographics (e.g., star-forming versus quiescent populations, as studied in this paper).

To combat this issue, we have looked for a proxy for black hole mass to use in extant photometric data, in order to test the AGN feedback paradigm from simulations at early cosmic times. In the lower panel of Figure 2, we find that in lieu of measurement of a black hole mass, the stellar potential (ϕ_*) is predicted to become the most effective predictor of quenching in simulations. This is a clear *prediction* of the simulations in the acausal photometric parameter space. The primary goal of the rest of the paper has been to test this prediction across cosmic time in JWST-CEERS, HST-CANDELS, and SDSS. This expands on previous work in Piotrowska et al. (2022) and Bluck et al. (2023).

The reason ϕ_* works as a good proxy for M_{BH} is that the two parameters are strongly correlated. In Figure 10 we show the $M_{\text{BH}}-\phi_*$ relationship for TNG (left) and EAGLE (right) at early cosmic times ($2 < z < 4$). See Bluck et al. (2023) for a presentation of these relationships at later cosmic epochs. In both simulations, there is a strong correlation between the two parameters, with a relatively tight scatter (see statistics presented on each panel). In fact, the scatter on this relationship is comparable, or tighter, to that found between black hole mass and central velocity dispersion in the local Universe (see Saglia et al. 2016; Terrazas et al. 2016, 2017; Piotrowska et al. 2022). We have tested that this relationship is the most effective to use in photometric data to constrain supermassive black hole masses via a variety of techniques. Most importantly, in the lower panel of Figure 2 we find that, in lieu of a measurement of black hole mass, the stellar potential becomes the most predictive observable parameter of quiescence. In essence, Figure 10 shows why this works.

In the simulations, the origin of these strong relationships is due to supermassive black holes forming and growing preferentially in

deep potential wells. See Bluck et al. (2023) for a discussion. Note that the correlations between these parameters are not seen explicitly in the RF analysis because when supermassive black hole mass is controlled for (as in the top panel of Figure 2), there is no residual dependence of quenching on the stellar potential. This exposes the stellar potential’s correlation with quenching as being spurious. Nonetheless, the stellar potential is still a highly useful parameter because we can leverage it as a proxy for black hole mass in photometric data, much as the central velocity dispersion is often used as a proxy for black hole mass in spectroscopic data (see, e.g., Bluck et al. 2022; Brownson et al. 2022; Piotrowska et al. 2022).

ORCID iDs

Asa F. L. Bluck <https://orcid.org/0000-0001-6395-4504>
 Christopher J. Conselice <https://orcid.org/0000-0003-1949-7638>
 Katherine Ormerod <https://orcid.org/0000-0003-2000-3420>
 Joanna M. Piotrowska <https://orcid.org/0000-0003-1661-2338>
 Nathan Adams <https://orcid.org/0000-0003-4875-6272>
 Duncan Austin <https://orcid.org/0000-0003-0519-9445>
 Joseph Caruana <https://orcid.org/0000-0002-6089-0768>
 K. J. Duncan <https://orcid.org/0000-0001-6889-8388>
 Leonardo Ferreira <https://orcid.org/0000-0003-1445-0590>
 Thomas Harvey <https://orcid.org/0000-0002-8649-2936>
 Roberto Maiolino <https://orcid.org/0000-0002-4985-3819>

References

- Abazajian, K. N., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2009, *ApJS*, 182, 543
 Aird, J., Coil, A. L., Moustakas, J., et al. 2012, *ApJ*, 746, 90
 Almosallam, I. A., Jarvis, M. J., & Roberts, S. J. 2016, *MNRAS*, 462, 726
 Ashby, M. L. N., Willner, S. P., Fazio, G., et al. 2015, *ApJS*, 218, 33
 Bagley, M. B., Finkelstein, S. L., Koekemoer, A. M., et al. 2023, *ApJ*, 946, L12
 Baldry, I. K., Balogh, M. L., & Bower, R. G. 2006, *MNRAS*, 373, 469
 Bluck, A. F. L., Conselice, C. J., & Almaini, O. 2011, *MNRAS*, 410, 1174
 Bluck, A. F. L., Maiolino, R., Brownson, S., et al. 2022, *A&A*, 659, 160
 Bluck, A. F. L., Maiolino, R., Piotrowska, J. M., et al. 2020b, *MNRAS*, 499, 230

- Bluck, A. F. L., Maiolino, R., Sánchez, S. F., et al. 2020a, *MNRAS*, **492**, 96
- Bluck, A. F. L., Mendel, J. T., & Ellison, S. L. 2014, *MNRAS*, **441**, 599
- Bluck, A. F. L., Mendel, J. T., Ellison, S. L., et al. 2016, *MNRAS*, **462**, 2559
- Bluck, A. F. L., Piotrowska, J. M., & maiolino, R. 2023, *ApJ*, **944**, 108
- Bondi, H., & Hoyle, F. 1944, *MNRAS*, **104**, 273
- Bower, R. G., Benson, A. J., Malbon, R., et al. 2006, *MNRAS*, **370**, 645
- Bower, R. G., McCarthy, I. G., & Benson, A. J. 2008, *MNRAS*, **390**, 1399
- Brammer, G. B., van Dokkum, P. G., & Coppi, P. 2008, *ApJ*, **686**, 1503
- Brinchmann, J., Charlot, S., & White, S. D. M. 2004, *MNRAS*, **351**, 1151
- Brownson, S., Belfiore, F., Maiolino, R., Lin, L., & Carniani, S. 2020, *MNRAS*, **498**, L66
- Brownson, S., Bluck, A. F. L., Maiolino, R., & Jones, G. C. 2022, *MNRAS*, **511**, 1913
- Brownson, S., Maiolino, R., Tazzari, M., Carniani, S., & Henden, N. 2019, *MNRAS*, **490**, 5134
- Bruzual, G., & Charlot, S. 2003, *MNRAS*, **344**, 1000
- Bundy, K., Georgakakis, A., Nandra, K., et al. 2008, *ApJ*, **681**, 931
- Bushouse, H., Eisenhammer, J., Dencheva, N., et al. 2022, JWST Calibration Pipeline, v1.6.2, Zenodo, doi:10.5281/zenodo.7041998
- Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, *ApJ*, **533**, 682
- Carnall, A. C., Begley, R., McLeod, D. J., et al. 2023, *MNRAS*, **518L**, 45
- Chabrier, G. 2003, *PASP*, **115**, 763
- Cheung, E., Faber, S. M., Koo, D. C., et al. 2012, *ApJ*, **760**, 131
- Cirasuolo, M., Fairley, A., Rees, P., et al. 2020, *ESO Messenger*, **180**, 10
- Cole, S., Lacey, C. G., Baugh, C. M., & Frenk, C. S. 2000, *MNRAS*, **319**, 168
- Conselice, C. J., Basham, J. T. F., & Bettaney, D. 2024, *MNRAS*, submitted
- Crain, R., Schaye, J., Bower, R. G., et al. 2015, *MNRAS*, **450**, 1937
- Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *MNRAS*, **365**, 11
- Dekel, A., & Bimboim, Y. 2006, *MNRAS*, **368**, 2
- Dimauro, P., Huertas-Company, M., Daddi, E., et al. 2018, *MNRAS*, **478**, 5410
- Duncan, K., Conselice, C. J., Mortlock, A., et al. 2014, *MNRAS*, **444**, 2960
- Duncan, K., Conselice, C. J., Mundy, C., et al. 2019, *ApJ*, **876**, 110
- Ellison, S. L., Lin, L., Thorp, M. D., et al. 2021a, *MNRAS*, **501**, 4777
- Ellison, S. L., Lin, L., Thorp, M. D., et al. 2021b, *MNRAS*, **502**, L6
- Ellison, S. L., Thorp, M. D., Lin, L., et al. 2020, *MNRAS*, **493**, L39
- Faber, S. 2011, The Cosmic Assembly Near-IR Deep Extragalactic Legacy Survey (“CANDELS”), STScI/MAST, doi:10.17909/T94S3X
- Fabian, A. C. 1994, *ARA&A*, **32**, 277
- Fabian, A. C. 1999, *MNRAS*, **308**, L39
- Fabian, A. C. 2012, *ARA&A*, **50**, 455
- Fabian, A. C., Sanders, J. S., Taylor, G. B., et al. 2006, *MNRAS*, **366**, 417
- Fang, J. J., Faber, S. M., Koo, D. C., & Dekel, A. 2013, *ApJ*, **776**, 63
- Ferrarese, L., & Merritt, D. 2000, *ApJL*, **539**, L9
- Ferreira, L., Conselice, C. J., Sazonova, E., et al. 2023, *ApJ*, **955**, 94
- Finkelstein, S., Bagley, M., Guang, Y., et al. 2023a, Data from The Cosmic Evolution Early Release Science Survey (CEERS), STScI/MAST, doi:10.17909/Z7P0-8481
- Finkelstein, S. L., Bagley, M. B., Ferguson, H. C., et al. 2023b, *ApJL*, **946**, 13
- Fukugita, M., & Peebles, P. J. E. 2004, *ApJ*, **616**, 643
- Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, *SSRv.*, **123**, 485
- Gensior, J., Kruijssen, J. M. D., & Keller, B. W. 2020, *MNRAS*, **495**, 199
- Grogan, N. A., Kocevski, D. D., Faber, S. M., et al. 2011, *ApJS*, **197**, 35
- Heckman, T. M., & Best, P. N. 2014, *ARA&A*, **52**, 589
- Henriques, B. M. B., White, S. D. M., & Lilly, S. J. 2019, *MNRAS*, **485**, 3446
- Henriques, B. M. B., White, S. D. M., Thomas, P. A., et al. 2013, *MNRAS*, **431**, 3373
- Henriques, B. M. B., White, S. D. M., Thomas, P. A., et al. 2015, *MNRAS*, **451**, 2663
- Hickox, R. C., Jones, C., Forman, W. R., et al. 2009, *ApJ*, **696**, 891
- Hlavacek-Larrondo, J., Fabian, A. C., & Edge, A. C. 2012, *MNRAS*, **421**, 1360
- Hlavacek-Larrondo, J., Gendron-Marsolais, M.-L., Fecteau-Beaucage, D., et al. 2018, *MNRAS*, **475**, 2743
- Hlavacek-Larrondo, J., McDonald, M., Benson, B. A., et al. 2015, *ApJ*, **805**, 35
- Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2006, *ApJS*, **163**, 1
- Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2008, *ApJS*, **175**, 356
- Hoyle, F., & Lyttleton, R. A. 1939, *PCPS*, **35**, 405
- Jones, G. C., Maiolino, R., & Carniani, S. 2023, *MNRAS*, **522**, 275
- Juodžbalis, I., Conselice, C. J., Singh, M., et al. 2023, *MNRAS*, **525**, 1353
- Kartalpe, J. S., Rose, C., Vanderhoof, B. N., et al. 2023, *ApJL*, **946**, 15
- Kocevski, D. D., Barro, G., McGrath, E. J., et al. 2023, *ApJL*, **946**, 14
- Koekemoer, A. M., Faber, S. M., Ferguson, H. C., et al. 2011, *ApJS*, **197**, 36
- Kriek, M., van Dokkum, P. G., Labbe, I., et al. 2009, *ApJ*, **700**, 221
- Lang, P., Wuyts, S., Somerville, R. S., et al. 2014, *ApJ*, **788**, 11
- Looser, T. J., D'Eugenio, F., Maiolino, R., et al. 2023, *Natur*, submitted (arXiv:2302.14155)
- Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415
- Maiolino, R., Cirasuolo, M., Afonso, J., et al. 2020, *Msng*, **180**, 24
- Maiolino, R., Gallerani, S., Neri, R., et al. 2012, *MNRAS*, **425L**, 66
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2023, *A&A*, submitted (arXiv:2308.01230)
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, **480**, 5113
- Marshall, M. A., Perna, M., Willott, C. J., et al. 2023, *A&A*, **678**, A191
- Martig, M., Bournaud, F., Teyssier, R., & Dekel, A. 2009, *ApJ*, **707**, 250
- McAlpine, S., Helly, J. C., Schaller, M., et al. 2016, *A&C*, **15**, 72
- McNamara, B. R., & Nulsen, P. E. J. 2007, *ARA&A*, **45**, 117
- Mendel, J. T., Simard, L., Palmer, M., Ellison, S. L., & Patton, D. R. 2014, *ApJS*, **210**, 3
- Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **477**, 1206
- Nandra, K., Georgakakis, A., Willmer, C. N. A., et al. 2007, *ApJ*, **660**, L11
- Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **475**, 624
- Nelson, D., Springel, V., Pillepich, A., et al. 2019, *ComAC*, **6**, 2
- Ormerod, K., Conselice, C. J., Adams, N. J., et al. 2024, *MNRAS*, **527**, 6110
- Pacucci, F., Nguyen, B., Carniani, S., Maiolino, R., & Fan, X. 2023, *ApJL*, **957**, L3
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
- Peng, C. Y., Ho, L. C., Impey, C. D., & Rix, H.-W. 2002, *AJ*, **124**, 266
- Peng, Y., Lilly, S., Kováč, K., et al. 2010, *ApJ*, **721**, 193
- Peng, Y., Lilly, S. J., Renzini, A., & Carollo, M. 2012, *ApJ*, **757**, 4
- Peng, Y., Maiolino, R., & Cochrane, R. 2015, *Natur*, **521**, 192
- Pillepich, A., Nelson, D., Hernquist, L., et al. 2018, *MNRAS*, **475**, 648
- Piotrowska, J. M., Bluck, A. F. L., Maiolino, R., Concas, A., & Peng, Y. 2020, *MNRAS*, **492**, L6
- Piotrowska, J. M., Bluck, A. F. L., Maiolino, R., & Peng, Y. 2022, *MNRAS*, **512**, 1052
- Planck Collaboration 2016, *A&A*, **594**, A13
- Pontoppidan, K. M., Barrientes, J., & Blome, C. 2022, *ApJL*, **936**, L14
- Rosario, D. J., Santini, P., Lutz, D., et al. 2013, *ApJ*, **771**, 63
- Saglia, R. P., Opitsch, M., Erwin, P., et al. 2016, *ApJ*, **818**, 47
- Saintonge, A., Catinella, B., Tacconi, L., et al. 2017, *ApJS*, **233**, 22
- Saintonge, A., Wilson, C. D., Xiao, T., et al. 2018, *MNRAS*, **481**, 3497
- Santini, P., Ferguson, H. C., & Fontana, A. 2015, *ApJ*, **801**, 97
- Shaye, J., Crain, R., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521
- Skill, J. M., Smith, B. D., & Danforth, C. W. 2012, *ApJ*, **759**, 23
- Sijacki, D., Springel, V., Di Matteo, T., & Hernquist, L. 2007, *MNRAS*, **380**, 877
- Sijacki, D., Vogelsberger, M., & Genel, S. 2015, *MNRAS*, **452**, 575
- Silk, J., & Rees, M. J. 1998, *A&A*, **331**, L1
- Simard, L., Mendel, J. T., Patton, D. R., Ellison, S. L., & McConnachie, A. W. 2011, *ApJS*, **196**, 11
- Soltan, A. 1982, *MNRAS*, **200**, 115
- Somerville, R. S., & Dave, R. 2015, *ARA&A*, **53**, 51
- Springel, V. 2005, *MNRAS*, **364**, 1105
- Springel, V. 2010, *ARA&A*, **48**, 391
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, **475**, 676
- Stefanon, M., Yan, H., Mobasher, B., et al. 2017, *ApJS*, **229**, 32
- Tacchella, S., Diemer, B., Hernquist, L., et al. 2019, *MNRAS*, **487**, 5416
- Teimoorinia, H., Bluck, A. F. L., & Ellison, S. L. 2016, *MNRAS*, **457**, 2086
- Terrazas, B. A., Bell, E. F., & Henriques, B. M. B. 2016, *ApJL*, **830**, L12
- Terrazas, B. A., Bell, E. F., Pillepich, A., et al. 2020, *MNRAS*, **493**, 1888
- Terrazas, B. A., Bell, E. F., Woo, J., & Henriques, B. M. B. 2017, *ApJ*, **844**, 170
- Trump, J. R., Sun, M., Zeimann, G. R., et al. 2015, *ApJ*, **811**, 26
- Trussler, J., Maiolino, R., Maraston, C., et al. 2020, *MNRAS*, **491**, 5406
- Ubler, H., Maiolino, R., Curtis-Lake, E., et al. 2023, *A&A*, **677**, A145
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014a, *Natur*, **509**, 177
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014b, *MNRAS*, **444**, 1518
- Voit, G. M., Bryan, G. L., Balogh, M. L., & Bower, R. G. 2002, *ApJ*, **576**, 601
- Ward, S. R., Harrison, C. M., Costa, T., & Mainieri, V. 2022, *MNRAS*, **513**, 2936
- Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 329
- Weinberger, R., Springel, V., Pakmor, R., et al. 2018, *MNRAS*, **479**, 4056
- Whitney, A., Ferreira, L., Conselice, C. J., & Duncan, K. 2021, *ApJ*, **919**, 139
- Woo, J., Dekel, A., Faber, S. M., et al. 2013, *MNRAS*, **428**, 3306
- Woo, J., Dekel, A., Faber, S. M., & Koo, D. C. 2015, *MNRAS*, **448**, 237
- Wuyts, S., Förster Schreiber, N. M., van der Wel, A., et al. 2011, *ApJ*, **742**, 96
- Zinger, E., Pillepich, A., Nelson, D., et al. 2020, *MNRAS*, **499**, 768