






















The Cliff: A Metal-Poor Little Red Dot Hosting an Overmassive Black Hole at $z = 3.55$

L. R. Ivey, ^{1,2,*} F. D'Eugenio, ^{1,2} R. Maiolino, ^{1,2,3} Y. Isobe, ^{1,2,4}, I. Juodžbalis, ^{1,2} S. Koudmani, ^{1,5}
M. Perna, ⁶ S. Zhang, ^{7,8} V. Bromm, ^{7,9} A. J. Bunker, ¹⁰ S. Carniani, ¹¹ A. C. Fabian, ¹² K.
Inayoshi, ¹³ X. Ji, ^{1,2} G. C. Jones, ^{1,2} B. Liu, ¹⁴ R. Pascalau, ^{1,2} P. Rinaldi, ^{15,16} B. Robertson, ¹⁷
J. Scholtz, ^{1,2} S. Tacchella ^{1,2}

¹Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

²Cavendish Laboratory, University of Cambridge, 19 JJ Thomson Avenue, Cambridge CB3 0HE, UK

³Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

⁴Waseda Research Institute for Science and Engineering, Faculty of Science and Engineering, Waseda University, 3-4-1, Okubo, Shinjuku, Tokyo 169-8555, Japan

⁵Centre for Astrophysics Research, Department of Physics, Astronomy and Mathematics, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, UK

⁶Centro de Astrobiología (CAB), CSIC-INTA, Cra. de Ajalvir Km. 4, 28850 – Torrejón de Ardoz, Madrid, Spain

⁷Weinberg Institute for Theoretical Physics, Texas Center for Cosmology and Astroparticle Physics, University of Texas at Austin, Austin, TX 78712, USA

⁸Department of Physics, University of Texas at Austin, Austin, TX 78712, USA

⁹Department of Astronomy, University of Texas at Austin, Austin, TX 78712, USA

¹⁰Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

¹¹Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy

¹²Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

¹³Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China

¹⁴Institut für Theoretische Astrophysik, Zentrum für Astronomie, Universität Heidelberg, Albert Ueberle Str. 2, D-69120 Heidelberg, Germany

¹⁵Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218, USA

¹⁶Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA

¹⁷Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

JWST has revealed a large population of massive black holes (BHs) in the early Universe with unusual properties which mark them as distinct from low-redshift active galactic nuclei. Such findings have prompted the development of new models of BH formation and growth, and of their co-evolution with host galaxies. Linking the gas-phase metallicity of BH environments to seed masses is key to understanding which evolutionary pathways could explain the population of *JWST*-discovered BHs. We present new high-resolution *JWST* NIRSpect/IFU observations covering the rest-frame optical emission lines of a Little Red Dot (LRD) at $z = 3.55$, known as *The Cliff*, from the ‘Red Unknowns: Bright Infrared Extragalactic Survey’ (RUBIES). We find evidence for low metallicity ($Z = 0.017 \pm 0.004 Z_{\odot}$) based on the low narrow-line [O III] λ 5007/H β ratio, supported by the non-detection of low-ionisation emission lines such as [O II] λ 3727,3729 and [N II] λ 6548,6583. We find that the observed properties of *The Cliff*, including its overmassive BH, can be reproduced by some simulations of black hole growth and evolution down to $z \sim 3.5$. However, these simulation runs require high seed masses ($10^4 - 10^5 M_{\odot}$) and appear as rarely in the simulation volume as in the RUBIES survey volume over redshifts $3 < z < 4$, highlighting the unusual nature of *The Cliff*. Future simulations and numerical models will help to uncover how such a metal poor system managed to develop a massive black hole and persist to such low redshift.

Key words: galaxies: high-redshift - galaxies: active - galaxies: nuclei - galaxies: abundances - galaxies: evolution

1 INTRODUCTION

The James Webb Space Telescope (*JWST*) has opened a new observational frontier for studying the formation and evolution of early black holes. The sensitivity and wavelength coverage of *JWST* have enabled the detection of active galactic nuclei (AGN) at high redshift

($z > 4$) with luminosities $L_{\text{bol}} \sim 10^{42} - 10^{46} \text{ erg s}^{-1}$, rendering it possible to study low-luminosity AGN at higher redshifts than ever before (e.g., Kocovski et al. 2023; Übler et al. 2023; Harikane et al. 2023; Matthee et al. 2024; Maiolino et al. 2024b,a; Kokorev et al. 2023; Greene et al. 2024; Taylor et al. 2025a,b; Juodžbalis et al. 2024b, 2026; Mazzolari et al. 2025; Chisholm et al. 2024; Scholtz et al. 2025; Adamo et al. 2025). The AGN most readily detected by *JWST* are broad line AGN (BLAGN; i.e. Type 1 AGN), identified

* E-mail: li247@cam.ac.uk

through the broad component of their permitted lines (usually H α or H β). The absence of a corresponding broad component in the forbidden lines (e.g. [O III] λ 5007) confirms an AGN Broad Line Region (BLR) as the origin for the broad emission, distinguishing it from potential outflows. The low luminosities of *JWST*-discovered AGN are likely due in part to lower black hole (BH) masses and/or lower accretion rates compared to the luminous high-redshift quasars previously identified by ground-based surveys (Maiolino et al. 2024b; Juodžbalis et al. 2024b, 2026).

However, the *JWST*-discovered low-luminosity AGN population has been found to exhibit a number of peculiarities. Crucially, these AGN are not simply the ‘high-redshift version’ of local ($z < 1$) AGN, nor are they ‘scaled-down versions’ of high-redshift quasars, meaning *JWST* AGN are likely a distinct population. Firstly, they show characteristic X-ray weakness (Ananna et al. 2024; Yue et al. 2024b; Maiolino et al. 2025b) and radio weakness (Mazzolari et al. 2024, 2025), while both weak optical variability (e.g. Ji et al. 2025; Furtak et al. 2025; Naidu et al. 2025; Zhang et al. 2025a; Lin et al. 2026a) and non-variability (e.g. Kokubo & Harikane 2025; Zhang et al. 2025c) have been observed. The chemical enrichment associated with *JWST* AGN is also poorly understood. Many *JWST* AGN exhibit a lack of strong high-ionisation lines (e.g. Lambrides et al. 2024; Zucchi et al. 2026) or classical broad-line iron emission bumps (e.g. Trefoloni et al. 2025). However, many *JWST* AGN also exhibit significant nitrogen enhancement (Ji et al. 2024; Isobe et al. 2025), suggesting unexpectedly rapid chemical enrichment, potentially driven by extreme star formation or exotic stellar populations.

Many *JWST* AGN also appear overmassive compared to the local black hole to stellar mass relation ($M_{\text{BH}}-M_*$; e.g., Maiolino et al. 2024b; Furtak et al. 2024; Juodžbalis et al. 2024b, 2026; Chen et al. 2025; Jones et al. 2025; Übler et al. 2023). The presence of relatively massive BHs in early low-mass galaxies is remarkable (Mezcua et al. 2024), and raises fundamental questions about the co-evolution of black holes and their host galaxies (e.g. Smith & Bromm 2019; Inayoshi et al. 2020). To a certain extent, the overmassive nature of *JWST* AGN may be ascribed to selection effects (Juodžbalis et al. 2026; Li et al. 2025; Geris et al. 2026; Habouzit 2025) or could point to $M_{\text{BH}} - M_{\text{dyn}}$ as a more fundamental relation between BHs and their host galaxies (McClymont et al. 2026; Danhaive et al. 2026b; Maiolino et al. 2024b). While selection effects alone cannot fully explain the observed shift in M_{BH}/M_* (Pacucci & Loeb 2024; Jones et al. 2025), recent work by Ziparo et al. (2026) suggests that this evolution might not be a systematic shift in the mean $M_{\text{BH}} - M_*$ relation (which remains consistent with local values at $z > 4$), but rather a significant increase in its intrinsic scatter at higher redshifts.

Assuming certain colour selection criteria, approximately 10-30% (Hainline et al. 2025) of *JWST*-discovered BLAGN at $z \sim 5$ have been categorised as ‘Little Red Dots’ (LRDs; Matthee et al. 2024; Hviding et al. 2025). LRDs exhibit a number of characteristic features, including rest-frame optical compactness, extreme redness and a V-shaped continuum (e.g. Matthee et al. 2024; Kocevski et al. 2025; Setton et al. 2025). Other features commonly observed in LRDs (though not exclusive to the LRD population, see e.g. Brazzini et al. 2026) include exponential broad-line wings (e.g. Rusakov et al. 2026), deep Balmer line absorption (Juodžbalis et al. 2024a; Lin et al. 2026b; D’Eugenio et al. 2025a, 2026), prominent Balmer breaks (e.g. Labbe et al. 2024; Baggen et al. 2024; Setton et al. 2025; de Graaff et al. 2025b), and X-ray weakness (e.g. Yue et al. 2024b; Maiolino et al. 2025b, but see Hviding et al. 2026 for an exception).

The Balmer absorption commonly seen in LRDs indicates high gas column densities along the line of sight (Juodžbalis et al. 2024a; Inayoshi & Maiolino 2025; Chang et al. 2026), a feature which may

also be related to their observed X-ray weakness (Maiolino et al. 2025b). The number density of LRDs has been found to decline sharply at $z < 4$ (Kocevski et al. 2025; Inayoshi 2025), though there are some luminous instances at lower redshifts (e.g. $z \sim 2 - 3$; Juodžbalis et al. 2024a; Loiacono et al. 2025; Wang et al. 2025b; Euclid Collaboration et al. 2025; Ma et al. 2025; Rinaldi et al. 2025) and a few local analogues (e.g. $z \sim 0.1$; Lin et al. 2026c; Ji et al. 2026). This rapid evolution in LRD number density suggests that the obscuration and extreme gas column densities characteristic of LRDs may represent a transient phase of rapid black hole growth prevalent primarily in the early Universe.

A wide range of models have now been proposed for interpreting the peculiar observed properties of *JWST* AGN, with fundamental differences in their initial seeding mechanisms. These can be broadly divided into two categories: (i) light seeds ($10^2 - 10^3 M_{\odot}$), possibly originating from Pop III stellar remnants (e.g. Nandal et al. 2025; Cammelli et al. 2025; Sanati et al. 2025; Prole et al. 2025), and (ii) intermediate/heavy seeds ($10^4 - 10^6 M_{\odot}$) originating from the runaway merging of nuclear star clusters (e.g. Partmann et al. 2025; Rantala et al. 2025, possibly resulting from the ‘feedback-free starbursts’ proposed by Dekel et al. 2025a,b), or from the direct collapse of massive gas clouds (e.g. Natarajan et al. 2024; Jeon et al. 2025; Pacucci et al. 2026), or even primordial black holes (e.g. Zhang et al. 2025d; Dayal & Maiolino 2026). Crucially, black hole seeding mechanisms are influenced by the gas-phase metallicity of the host galaxy. Pristine, low-metallicity environments are often cited as a prerequisite for heavy seed formation (e.g. Begelman et al. 2006; Ferrara et al. 2014; Latif & Ferrara 2016; Maiolino et al. 2025a), as such environments suppress gas fragmentation which would otherwise lead to normal star formation (Inayoshi et al. 2020). By initiating growth from an already massive seed, a BH could outpace stellar assembly in its host galaxy even at early cosmic times ($z \gtrsim 4$; e.g. Natarajan et al. 2024; Pacucci et al. 2023). Conversely, higher metallicities could hinder heavy seed formation (Latif & Ferrara 2016) but may facilitate rapid growth via stellar mergers or super-Eddington accretion (Inayoshi et al. 2020).

If LRD BH masses are verified as overmassive relative to their host galaxies (see Juodžbalis et al. 2025, for example), they could be explained by invoking heavy seed models or even Primordial Black Holes (PBHs; see e.g. Dayal 2024; Dayal & Maiolino 2026; Zhang et al. 2025e; Ziparo et al. 2025). Alternatively, LRDs could represent an entirely new class of rapidly-accreting, low-mass BHs termed ‘black hole stars’ (Naidu et al. 2025; Begelman & Dexter 2026) with cool gas envelopes (Liu et al. 2025; Kido et al. 2025). It has even been proposed that LRDs could simply be super-Eddington accreting black holes, which otherwise have the same structure as normal AGN (Madau & Maiolino 2026; Madau et al. 2026). Distinguishing between these pathways is essential to understanding the formation and early growth of BHs (e.g., Jeon et al. 2026); in this regard, environmental metallicity may prove to be a valuable diagnostic for evaluating the plausibility of different seeding models.

A luminous, unlensed LRD at $z_{\text{spec}} = 3.548$, also known as *The Cliff*, was recently identified by the *Red Unknowns: Bright Infrared Extragalactic Survey* (RUBIES; de Graaff et al. 2025a,b). The survey obtained *JWST* Near Infrared Spectrograph (NIRSpec) microshutter assembly (MSA; Jakobsen et al. 2022; Ferruit et al. 2022) observations of *The Cliff* with both the low-resolution PRISM/Clear (0.6–5.3 μm) and medium-resolution G395M/F290LP (2.9–5.2 μm) disperser/filter combinations. These observations revealed broad H β and H α emission, typical of Type 1 AGN, along with the V-shaped continuum characteristic of LRDs. *The Cliff* also shows an exceptionally strong Balmer break, for which de Graaff et al. (2025b) exclude

a stellar origin. de Graaff et al. (2025b) also found *The Cliff* hosts an AGN with mass of $\log(M_{\text{BH}}/M_{\odot}) = 7.18_{-0.06}^{+0.07}$. Even with these medium-resolution observations they identify a significant absorption component in H α , indicative of absorption by dense gas.

Another key and interesting feature of the original PRISM and G395M observations is the non-detection of the [O III] λ 5007 line. This could indicate a weak Narrow Line Region (NLR), which has previously been found for powerful high-redshift quasars (Netzer et al. 2004). In this paper, we revisit *The Cliff*, now with recent high-resolution *JWST*/NIRSpec Integral Field Unit (IFU) observations (G235H/F170LP; PID 9433), which finally reveal constraints on the narrow line kinematics and metallicity from [O III] λ 5007. The high resolution spectrum shows a clear narrow component for both H β and [O III] λ 5007, revealing the presence of either a NLR or ISM ionised by star-formation. In this work we analyse the high-resolution spectral data to assess the [O III] λ 5007 weakness more comprehensively and argue it is tracing low-metallicity gas within which the BH is embedded. We additionally discuss the implications of such a low-metallicity, high- M_{BH} system existing at $z \sim 3.5$ in the context of chemical enrichment and black hole seeding.

The paper is organised as follows. In Section 2 we describe the spectroscopic data and data reduction used in this work. We present our analysis of the data in Section 3, the results of which are then described in Section 4. In Section 5, we discuss and interpret our findings in the wider context of other observations and simulations. Finally, our key observations and findings are summarised in Section 6. Throughout this work, unless stated otherwise, we assume a flat Λ CDM cosmology with $H_0 = 67.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.315$ (Planck Collaboration et al. 2020). At $z = 3.548$, this means an on-sky separation of $1''$ corresponds to 7.453 kpc . We additionally take solar metallicity to be defined by $[12 + \log(\text{O}/\text{H})]_{\odot} = 8.69$ (Asplund et al. 2009).

2 DATA

In this work, we utilise NIRSpec-IFS observations from the Cycle 4 *JWST* DDT programme 9433 (PIs R. Maiolino and F. D’Eugenio). The observations were taken with the high-resolution disperser/filter combination G235H/F170LP. We adopted a medium-cycling pattern with 10 dithers, 18 groups per integration and one integration per exposure, along with the NRSIRS2 readout mode (Rauscher et al. 2017). This yields a total on-source time of 3.7 h.

The data reduction uses the *JWST* Science calibration pipeline v1.20.2 with CRDS jwst_1464.pmap. To improve the quality of the datacube, we perform a number of further data processing steps, in addition to the default pipeline. The customised procedures used for flagging of residual cosmic ray snowballs, subtraction of pink noise ($1/f$ noise), and removal of failed open MSA shutters signal are described in detail in Perna et al. (2023, 2025). Stage 2 was executed using the bad-pixel self-calibration implemented in the standard pipeline, in which all exposures on a given detector are used to identify and flag bad pixels that may have been missed by the reference bad-pixel mask. For the final stage of the pipeline, the outlier rejection was performed using the standard *JWST* pipeline implementation. Before combination, the cube was resampled to a scale of $0.05''$ per spaxel; the final cubes were combined using the ‘drizzle’ method.

We note that the NIRSpec field of view (FoV) of *The Cliff* includes both a foreground galaxy at $z_{\text{spec}} = 3.05086 \pm 0.00002$ (CANDELS ID 24647; Grogin et al. 2011; Koekemoer et al. 2011), and a star (see de Graaff et al. 2025b and Fig. 1). To perform back-

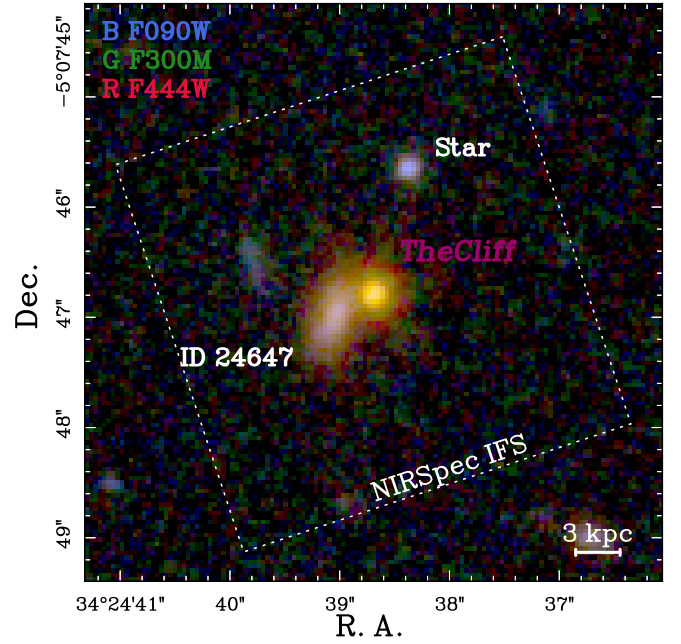


Figure 1. False-colour RGB image of *The Cliff*, illustrating the NIRSpec/IFS FoV. We label a foreground galaxy (ID 24647, at $z_{\text{spec}} = 3.05$) to the east and a star to the north west. We use publicly available imaging from PRIMER (PID 1837) and MINERVA (PID 7814; Muzzin et al. 2025).

ground subtraction, we created a mask of source-free regions in the FoV, carefully excluding *The Cliff*, the foreground galaxy and star. The background level is estimated independently for each spectral channel, performing σ -clipping with a threshold of $\sigma = 5$ from the median, with the local sky level then defined as the median of the resulting distribution. The final estimated background cube is smoothed with a median filter to avoid introducing additional noise, and is then subtracted from the flux data cube.

As noted by earlier studies (e.g. Übler et al. 2023; Rodríguez Del Pino et al. 2024), the ERR extension of the NIRSpec-IFS cubes is underestimated. Hence, to obtain realistic uncertainties on our aperture-integrated spectra while retaining crucial information about outliers and the correlated noise between channels, we adopt the same procedure as previous works: rescaling ERR to match the σ -clipped r.m.s. noise of the integrated spectrum in regions free of line emission.

3 DATA ANALYSIS

3.1 PSF and aperture determination

The optimal aperture from which an integrated spectrum of *The Cliff* is extracted is not perfectly circular, as it is necessary to account for the shape of the IFU observations’ point spread function (PSF). The PSF of IFU observations is non-trivial to determine, for two key reasons: (i) the observational PSF does not exactly follow the theoretical *JWST* PSF, and (ii) the PSF tends to be slightly elongated in the direction along the IFU slices. In the case of *The Cliff*, we leverage the star present in the IFU FoV as a reference point-source to obtain a direct measurement of the PSF for our observations (see Appendix A).

To measure the PSF shape and full-width at half maximum (FWHM), we utilise the morphology modelling tool PYSERSIC (Pasha & Miller 2023). We generate a set of square narrow-band image

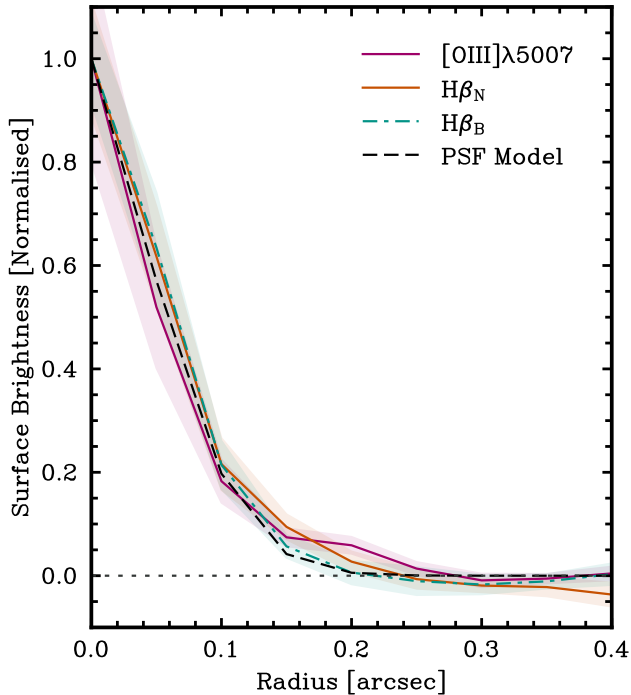


Figure 2. Radial surface brightness profiles (normalised to the peak) of the narrow $H\beta$ and $[O\text{ III}]\lambda 5007$ (the solid orange and pink lines, respectively) and broad $H\beta$ (dot-dashed green line) emission lines. The PSF model of *The Cliff* (see Section 3.1 and Appendix A) is shown by the black dashed line. The shaded areas correspond to 1σ uncertainties on the mean for each profile. The emission lines have compact Gaussian cores, and do not show any significant extension beyond the scope of the PSF. While there are hints of extension out to a radius of $0.2''$ in $[O\text{ III}]\lambda 5007$, beyond the scale of the PSF, it should be noted that the wings of the PSF are not well-constrained and therefore this cannot be interpreted as conclusive evidence for spatial extension of *The Cliff*. Overall, these profiles suggest *The Cliff* is spatially unresolved in NIRSpc/IFS.

cutouts centred on the star by median-stacking all wavelength slices within spectral windows spaced by $0.1\ \mu\text{m}$ in the NRS1 detector (where the star is brightest), and by $0.2\ \mu\text{m}$ in the NRS2 detector (which samples longer wavelengths). We model these synthetic images using a 2D Gaussian, implemented as a Sérsic profile with fixed index $n = 0.5$, along with a linear sky background. We adopt generous flat priors on all parameters, except for position angle (P.A.), which has a narrow flat prior within $\pm 5^\circ$ from the NIRSpc position angle. We find that in NRS1 no constraints on the position angle are necessary, and we find a position angle that is fully consistent with the NIRSpc position angle, i.e. P.A. = 198° . However, in NRS2 the position angle is loosely constrained, likely due to the combination of a mildly rounder shape and lower SNR. For consistency, therefore, we adopt the same P.A. for all wavelengths. As PYSERSIC also requires an input PSF for the fit, we use a Gaussian with FWHM = 0.1 spaxels, a very small width representing a delta distribution; using a much smaller FWHM would incur artefacts in the fitting procedure. The PSF parameters are estimated with a simple maximum-likelihood estimator. We find that for our main spectral region of interest, the $H\beta$ - $[O\text{ III}]\lambda 5007$ complex, the corresponding PSF is well-described by an ellipse with axis ratio $q = 0.86$ and position angle P.A. = 198° ; the direction of elongation of the PSF is along the IFU slices.

To determine the size of the optimal integrated spectrum extraction aperture, it is important to first investigate whether *The Cliff* shows




Figure 3. Map of narrow $H\beta$ emission, obtained by collapsing the three central channels of the continuum- and background-subtracted line (see Section 3.2.1 for the line shapes). This is overlaid with a white dashed ellipse illustrating the central aperture used for extracting the spectra shown in Figs. 4, 5 and 6. While this map may include some contribution from the broad $H\beta$ component, this has no impact on our results as it is not used for any quantitative measurements.

any evidence for spatially-extended emission. In Fig. 2, we present the radial surface brightness profiles of the $H\beta$ and $[O\text{ III}]\lambda 5007$ emission lines, compared to the profile of our best-fit PSF model (dashed black line; also refer to Appendix A). The profile of narrow $H\beta$ (henceforth denoted $H\beta_N$; solid orange line) is obtained by simply collapsing the three spectral channels around the peak of $H\beta_N$ (a range of $\pm 75\ \text{km s}^{-1}$); similarly, the profile of $[O\text{ III}]\lambda 5007$ (solid pink line) is also obtained by collapsing the narrow line over the three spectral channels around its peak. The profile of broad $H\beta$ (henceforth denoted $H\beta_B$; dot-dashed teal line) is extracted from the spectral regions $2.204\text{--}2.210\ \mu\text{m}$ and $2.214\text{--}2.221\ \mu\text{m}$ (i.e. including emission from both the blue and red wings). The comparison of the profiles demonstrates that the bulk of $H\beta$ and $[O\text{ III}]\lambda 5007$ emission comes from a very compact, unresolved region. Unlike in Abell2744-QSO1 (hereafter QSO1; Maiolino et al. 2025a), there is no conclusive evidence for an extended component of $H\beta$ or $[O\text{ III}]\lambda 5007$ emission in *The Cliff*, indicating that *The Cliff* is ultimately compact and spatially unresolved at our resolution. Further evidence for this conclusion arises from morphological fitting of the star in the IFU FoV: *The Cliff* appears to show the same spatial profile as the star, i.e. any apparent extension arises due to the PSF.

We determine the optimal spectral extraction aperture for *The Cliff*, using a curve of growth analysis, which allows us to maximise the signal-to-noise ratio (SNR) of the resulting integrated spectrum while accurately accounting for aperture losses. As $H\alpha$ is detected with the greatest SNR, we utilise the curve of growth of the broad $H\alpha$ emission to derive overall aperture corrections. To account for wavelength-dependent aperture losses, the $H\alpha$ correction is then scaled to determine the relevant aperture loss correction (ALC) factors for shorter-wavelength lines such as $H\beta$ and $[O\text{ III}]\lambda 5007$ (see Appendix C). The final extraction aperture is an ellipse defined by the same axis ratio and P.A. as the PSF, with major radius of $0.1''$, and centred on the continuum centroid of *The Cliff*. The final integrated




Figure 4. Full, background-subtracted NIRSpect-IFU G235H/F170LP integrated spectrum of *The Cliff*, extracted from the aperture shown in Fig. 3. The locations of key emission lines are labelled with coloured vertical lines (except [N II] $\lambda\lambda$ 6548,6583, to avoid overcrowding near H α). This spectrum has not been corrected for aperture losses.

spectrum extraction aperture used for *The Cliff* in this work is shown by the white dashed ellipse in Fig. 3, overlaid onto a map of H β_{N} emission. The full integrated spectrum of *The Cliff* extracted from this $r_{\text{maj}} = 0.1''$ aperture is presented in Fig. 4. For this aperture, the ALC factor for H α is found to be 1.88 ± 0.05 .

3.2 Spectral analysis

To analyse the emission lines, we adopt a Bayesian approach, using different models for different regions of the spectrum as discussed across the following subsections. To integrate the posterior distribution of each fit, we use the Markov Chain Monte Carlo (MCMC) method with the software *emcee* (Foreman-Mackey et al. 2013). We run 15000 steps for each chain, with 50 percent burn-in steps. To initialise the chains, we first identify the minimum- χ^2 solution using the relevant model and ordinary least-squares minimisation. For the MCMC fit, we adopt Gaussian priors informed by the minimum- χ^2 fit, with strong boundaries only set when physically motivated (e.g. to ensure non-negative emission line fluxes). All chains are visually inspected for convergence. The final best-fit parameters and their uncertainties are calculated as the median value and 68 percent confidence interval of the posterior distribution. To account for any uncorrelated uncertainties, all quantities derived from the spectral fitting are calculated from the posterior distribution. The key spectral properties determined from the model fits described in this section, together with their uncertainties, are summarised later in Table 1.

3.2.1 The H β -[O III] λ 5007 complex

Fig. 5 shows the integrated spectrum of *The Cliff*, now highlighting the H β and [O III] $\lambda\lambda$ 4959,5007 emission lines. This integrated spectrum reveals the narrow and broad components of H β emission, together with a broad and slightly-redshifted absorption component; the absorption in H β was not visible in the earlier PRISM observations presented by de Graaff et al. (2025b). We also note the detection of the weak and narrow [O III] $\lambda\lambda$ 4959,5007 lines in our observations, which were not significantly detected by de Graaff et al. (2025b) in the low resolution spectrum. In this work, the narrow [O III] λ 5007

emission provides valuable constraints for the kinematics of the narrow H β line in our spectral fits of the G235H spectrum, meaning it is possible to spectrally disentangle the H β_{B} and absorption components from the H β_{N} component.

To fit the H β -[O III] λ 5007 complex, we adopt an approach similar to that outlined by Ji et al. (2025) or Maiolino et al. (2025a). We fit a single Gaussian profile to describe the narrow components of each emission line, tying their redshift centroids and intrinsic FWHMs (z_{N} and FWHM_{N} , respectively) to common values to reduce the number of free parameters. We allow the amplitude of each Gaussian profile to vary freely, but the [O III] λ 5007/[O III] λ 4959 flux ratio is fixed to 2.99 (Dimitrijević et al. 2007). The broad line component is modelled as a Gaussian profile with its own independent redshift and intrinsic FWHM (z_{B} and FWHM_{B} , respectively). In order to retrieve intrinsic line FWHMs from the spectral data, we must account for the wavelength-dependent line spread function (LSF) of *JWST*/NIRSpect. The nominal, pre-launch resolution LSF (Jakobsen et al. 2022) from the JDOCS¹ is too conservative because *JWST* provides a lower than expected wavefront error. Hence, for our deconvolution, we assume the nominal LSF improved by a factor of 0.7 (roughly in agreement with e.g. Shajib et al. 2025). The intrinsic FWHM of any individual emission line is therefore given by:

$$\text{FWHM}_{\text{int}} = \sqrt{\text{FWHM}_{\text{obs}}^2 - (0.7 \times \text{FWHM}_{\text{pre-launch LSF}})^2}. \quad (1)$$

The continuum is modelled with a power-law, with the amplitude and exponent of the power law as additional free parameters.

Finally, we add a H β absorption component, associated with the broad-line region. The absorption component is described by 4 parameters: velocity offset v_{abs} relative to the narrow line, central optical depth τ_0 , covering fraction C_f and velocity dispersion σ_{abs} . The residual intensity at wavelength λ is given by

$$I(\lambda)/I_0(\lambda) = 1 - C_f + C_f \exp(-\tau(k;\lambda)), \quad \text{with} \quad (2)$$

$$\tau(k;\lambda) = \tau_0(k)f[v(\lambda)],$$

where $I_0(\lambda)$ is the spectral flux density before absorption (consist-

¹ Available at the [jwst-docs website](https://www.jwst-docs.com).

ing of both the continuum and $H\beta_B$ emission), $\tau_0(k)$ is the optical depth at the centre of the line ($k = H\beta$ here), and $f[v(\lambda)]$ is the velocity distribution of the absorbing atoms, assumed to be a Gaussian probability distribution (see e.g. D’Eugenio et al. 2025a) and parameterised by v_{abs} and σ_{abs} .

The middle panel of Fig. 5 shows the integrated spectrum, now with the broad $H\beta$, continuum and absorption components subtracted. The BLR-subtracted spectrum highlights the prominence of $H\beta_N$ and the relative weakness of $[O\text{ III}]\lambda 5007$. The inferred ratio between $[O\text{ III}]\lambda 5007$ and $H\beta_N$ is low, $F([O\text{ III}]\lambda 5007)/F(H\beta_N) = 1.59^{+0.26}_{-0.24}$, as reported in Table 3. While this ratio is not as low as in the case of QSO1 ($F([O\text{ III}]\lambda 5007)/F(H\beta_N) \sim 0.55$; Maiolino et al. 2025a) for example, we will discuss the $[O\text{ III}]\lambda 5007$ weakness and its implications for metallicity in Section 4.2.1.

Due to the relatively low signal to noise of the $H\beta$ - $[O\text{ III}]\lambda 5007$ complex, we note as a caveat that we cannot meaningfully perform statistical tests of different models for the broad $H\beta$ line profile (e.g. exponential wings, double Gaussian BLR, Lorentzian).

3.2.2 The $H\alpha$ profile

Fig. 6 shows the integrated spectrum of *The Cliff* around the $H\alpha$ emission line. The high-resolution IFU spectrum reveals the complexity of *The Cliff*’s $H\alpha$ profile: deep absorption, no visibly distinct narrow component, and potentially asymmetric broad wings. While a single- or double-Gaussian BLR model is often used to approximate the $H\alpha$ line breadth (as was the case for our $H\beta$ fit), such models result in a poor fit to the wings of the profile. In contrast, a ‘scattering model’ (see e.g. Rusakov et al. 2026) can account for exponential wings in the line profile, which are often attributed to Doppler broadening due to electron or resonant scattering (e.g. Laor 2006). Even with the higher resolution of the G235H spectrum (compared to earlier PRISM or G235M observations), we still lack the spectral resolution required to decisively disentangle the individual $H\alpha$ spectral components. As a result, we deem the exponential model to be a more appropriate parameterisation of the $H\alpha$ profile for this work (compared to a double Gaussian model, for example). For simplicity, and for comparison with some previous studies, we model the exponential profile in the context of the scattering scenario. However, we note that while the exponential profile is often ascribed to electron scattering by a ionised medium, Scholtz et al. (2026) and Madau et al. (2026) have recently shown that the same profile can be ascribed to normal BLR stratification, or to the overlap of multiple non-exponential components. A full exploration of the $H\alpha$ emission line profile, and the nature of its absorber, is beyond the scope of this work but has been presented in Scholtz et al. (2026).

To model the $H\alpha$ profile, we first define the narrow component of $H\alpha$ (henceforth $H\alpha_N$) and the $[N\text{ II}]\lambda\lambda 6548, 6583$ doublet as Gaussians with their kinematic parameters (i.e. z_N and FWHM_N) constrained to those from the $H\beta$ - $[O\text{ III}]\lambda 5007$ fit, with the $[N\text{ II}]\lambda 6548/[N\text{ II}]\lambda 6583$ flux ratio fixed to 0.327 (see Dojčinović et al. 2023). In line with the ‘scattering scenario’ parameterisation, we adopt a BLR profile comprised of an intrinsic Gaussian profile $G_{\text{BLR}}(\lambda)$ described by its own independent redshift (z_B) and FWHM (FWHM_B), convolved with an exponential kernel defined by

$$E(\lambda_0, W; \lambda) \propto \exp\left(-\frac{|\lambda - \lambda_0|}{W}\right), \quad (3)$$

where λ_0 is the central wavelength (assumed to be the same as that of the Gaussian BLR) and W is the exponential width (Laor 2006).

The overall BLR emission is therefore given by:

$$\text{BLR} = f_{\text{scatt}}E(\lambda) * G_{\text{BLR}}(\lambda) + (1 - f_{\text{scatt}})G_{\text{BLR}}(\lambda), \quad (4)$$

where f_{scatt} is the fraction of scattered light. f_{scatt} can be related to the optical depth of the scattering medium (τ_{thom}) by $1 - e^{-\tau_{\text{thom}}}$. The exponential width is related to both τ_{thom} and the temperature of the scattering gas (T) by

$$W = (428 \times \tau_{\text{thom}} + 370) \times (T/10^4 \text{ K})^{0.5}; \quad (5)$$

see Laor (2006) and Rusakov et al. (2026). Note that within our analysis T is treated as a nuisance parameter, and is not interpreted as a physical measurement of the system’s temperature. Once again, we model the continuum with a power law. Finally, we include an absorption component with the same C_f as that fitted to $H\beta$, but its other parameters ($v_{\text{abs}, H\alpha}$, $\tau_0(H\alpha)$ and $\sigma_{\text{abs}, H\alpha}$) allowed to vary independently.

Supplying constraints and Gaussian priors for this model based on results from the $H\beta$ - $[O\text{ III}]\lambda 5007$ fit ensures the resulting $H\alpha$ model is physically motivated. This is particularly important as the absorption component is only redshifted by approximately 2 spectral pixels, and there hence exists significant degeneracy between the narrow line and absorption components.

We investigate this degeneracy and its impact further by fitting two variations on our overall $H\alpha$ model. Firstly, in our *fiducial H α model* (top panel of Fig. 6), the $H\alpha_N/H\beta_N$ flux ratio is fixed to the typical Case B value of 2.86, appropriate for $T = 10,000$ K and $n_e \sim 10^2 - 10^3 \text{ cm}^{-3}$, and assuming no reddening (Osterbrock & Ferland 2006, here accounting for relative ALCs between $H\alpha$ and $H\beta$)². The second variation is our *secondary H α model* (middle panel of Fig. 6), in which the flux of $H\alpha_N$ is allowed to freely vary. We make a statistical comparison of the two models through the Bayesian Information Criterion (BIC; see e.g. Schwarz 1978), defined as

$$\text{BIC} = \chi^2 + k \ln n, \quad (6)$$

where χ^2 is the chi-squared fit statistic, k is the number of free parameters, and n is the number of data points. One model is generally considered to be statistically preferred over another if the relative change in BIC (ΔBIC) is < -10 . With the *secondary H α model*, we obtain a slightly better fit in terms of BIC, and measure $F(H\alpha_N)/F(H\beta_N) = 1.70 \pm 0.38$, much lower than the Case B ratio (though marginally consistent within 3σ). However, ΔBIC between the *fiducial* and *secondary H α models* is only ~ -5 , which is not statistically significant enough to argue the *fiducial* model should be ruled out. Hence, the observed $H\alpha$ profile could be regarded as consistent with a Case B Balmer decrement.

According to our best fit, $[N\text{ II}]\lambda\lambda 6548, 6583$ is undetected ($\text{SNR} < 0.3$), and we therefore determine the 3σ upper limit on the $[N\text{ II}]\lambda\lambda 6548, 6583$ flux to be $F([N\text{ II}]) < 0.18 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2}$ (see Table 2). We will utilise this constraint later in Section 4.2.1.

In principle, $H\alpha$ could have been included in a simultaneous fitting with the $H\beta$ - $[O\text{ III}]\lambda 5007$ complex. However, the high flux and severe blending of $H\alpha$ could have driven increased uncertainty in the $H\beta$ - $[O\text{ III}]\lambda 5007$ fit parameters and effectively reduced the SNR of these lines, introducing major caveats for our analysis. Instead, with our approach we prioritise securing an understanding of the narrow-line kinematics, applying these measurements as a prior on $H\alpha_N$ to better inform the deblending of the $H\alpha$ profile while ensuring confidence in our analysis of $H\beta_N$ and $[O\text{ III}]\lambda 5007$.

² We note that the Balmer decrement is insensitive to density variation of $10^2 - 10^6 \text{ cm}^{-3}$; see e.g. Storey & Hummer (1995).

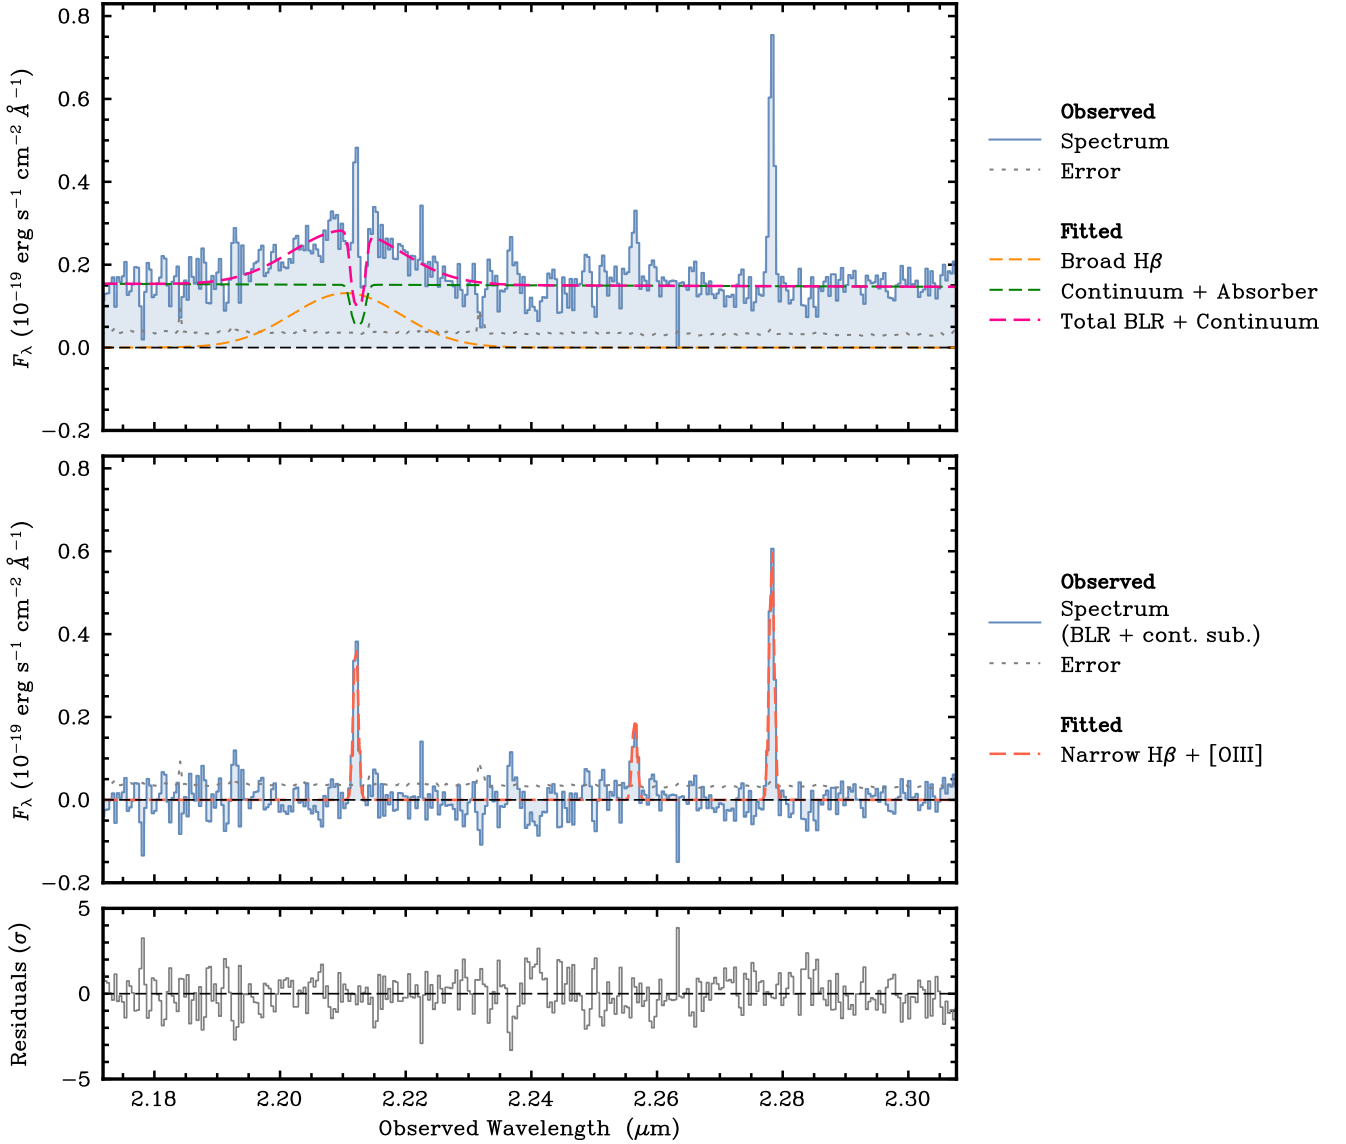


Figure 5. The integrated spectrum of *The Cliff* presented in Fig. 4, highlighting the H β -[O III] λ 5007 spectral region. **Top panel:** The integrated spectrum (blue solid line), with the fitted broad component of H β (orange dashed line), and the continuum and H β absorption (green dashed line). The pink dashed line shows the total fit to the BLR and continuum emission. **Middle panel:** The same spectrum, where the broad H β , continuum and absorption components have now been subtracted to highlight the narrow components of H β and [O III] λ 4959,5007. Note that the feature observed between H β and [O III] λ 4959 was not found to correspond to any known emission line (e.g. He I or [Fe II]) and is therefore deemed to be a noise feature. **Bottom panel:** Residuals from the fit, in units of σ .

3.2.3 Constraints on other emission lines

Some additional key lines for our study are the low-ionisation forbidden line doublets [O II] λ 3727,3729 and [S II] λ 6716,6731 (see Section 4.2.1), along with the temperature-sensitive auroral [O III] λ 4363 line. Fig. 7 shows integrated spectrum zoom-ins for each of these emission lines, illustrating a lack of any significant detection (all have SNR \lesssim 1.5; see Table 2). The upper limits on the fluxes of these emission lines provides valuable constraints on the inter-stellar medium (ISM) properties of *The Cliff* (see Section 4.2.1 for further discussion).

Hence, we fit each of these emission lines with LMFIT (Newville et al. 2016) and simple single-Gaussian models to obtain upper limits on their integrated fluxes, again fixing the intrinsic FWHM_N of each line to the value from the H β fit. For the [O II] λ 3727,3729 and

[S II] λ 6716,6731 doublets, the flux ratios are dependent on ISM conditions (e.g. the [O II] λ 3727,3729 flux ratio depends on density), and their flux ratios are therefore allowed to vary over the permitted range identified with PYNEB (Luridiana et al. 2015). As we note a tentative detection of H γ (see middle panel of Fig. 7), we fit H γ and [O III] λ 4363 together, obtaining an upper limit on [O III] λ 4363 and a tentative detection of H γ (SNR \sim 2.9; see Section 4.1 for further discussion of H γ).

The integrated fluxes of all emission lines measured in this work, including their SNRs, are presented in Table 2. A summary of all flux ratios used in this work, together with their abbreviations and measured values or limits, is presented in Table 3.

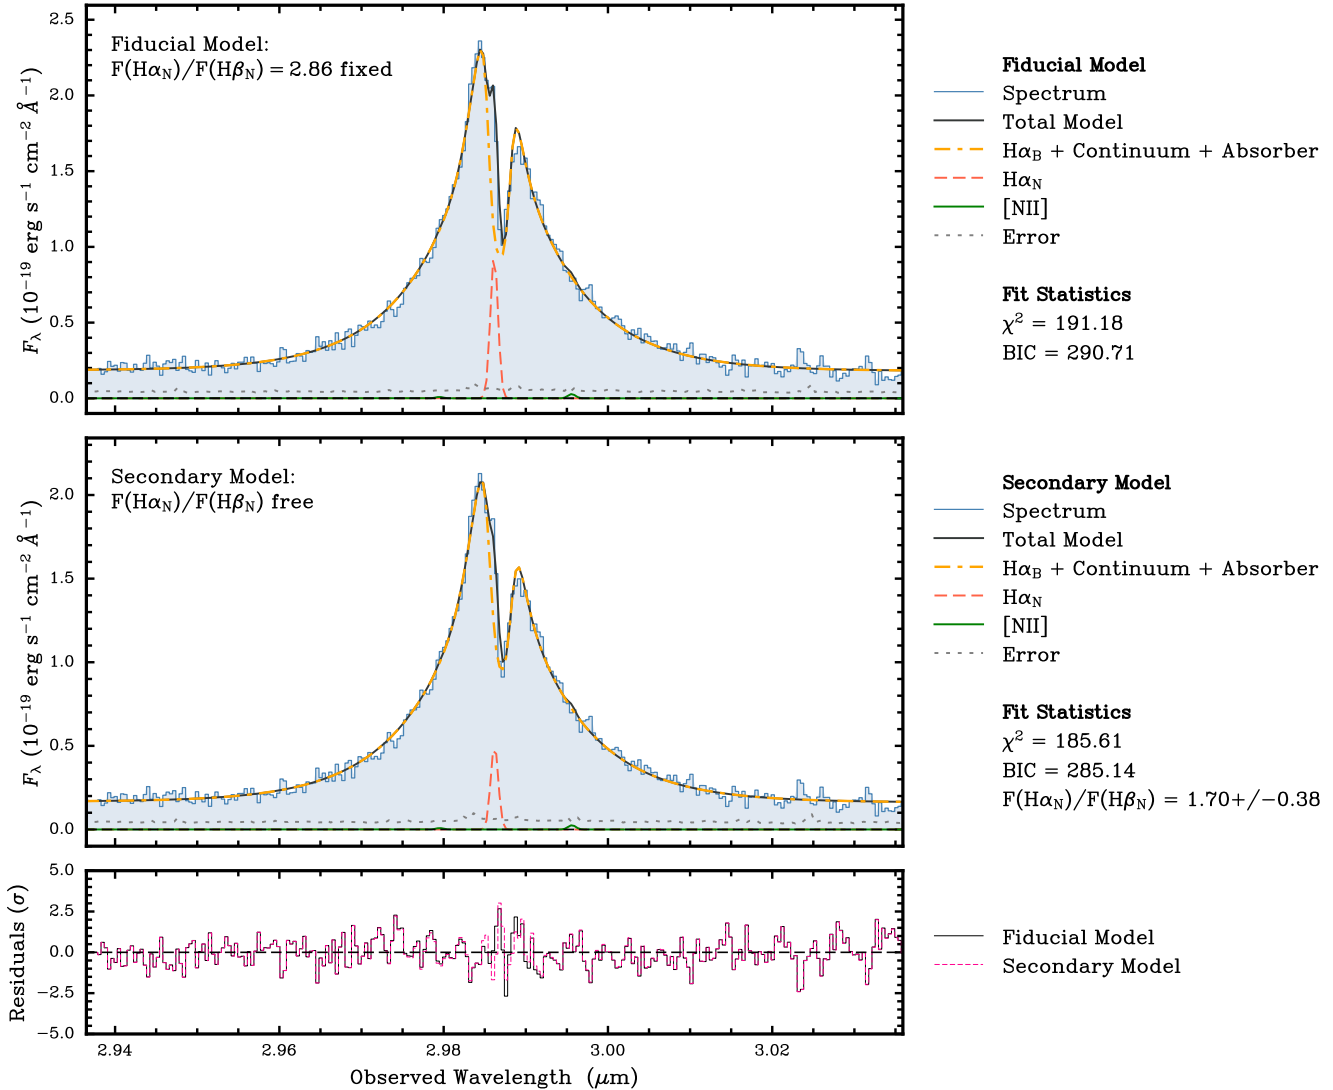


Figure 6. Integrated spectrum of *The Cliff* around $H\alpha$, fitted with two variations on the scattering model as described in Section 3.2.2. **Top panel:** Our *fiducial* $H\alpha$ model, with the $H\alpha_N/H\beta_N$ flux ratio restricted to Case B as part of the fit, i.e. $F(H\alpha_N)/F(H\beta_N) = 2.86$. Components shown are the integrated spectrum (blue solid line), the overall $H\alpha$ model (solid black line), the total fit to the BLR and continuum emission (dot-dashed orange line), the narrow component of $H\alpha$ (dashed red line) and the $[N\text{II}]$ doublet (solid green line). **Middle panel:** The same overall model, but now with the $H\alpha_N/H\beta_N$ flux ratio allowed to deviate from Case B (our *secondary* $H\alpha$ model). Individual components are shown with the same line styles as in the top panel. **Bottom panel:** A comparison of the fit residuals for the *fiducial* $H\alpha$ model (black solid line) and *secondary* $H\alpha$ model (dotted pink line) in units of σ .

4 RESULTS

4.1 Dust attenuation

The Cliff exhibits a complex $H\alpha$ profile, with strong degeneracy between the narrow component and absorber. Consequently, it was not possible to determine a well-constrained $H\alpha_N/H\beta_N$ Balmer decrement; indeed, detailed exploration of the $H\alpha$ and $H\beta$ profiles is beyond the scope of this paper. Consequently, within this work we adopt the Case B ratio of $F(H\alpha)/F(H\beta) = 2.86^{+0.48}_{-0.21}$ derived from our fiducial $H\alpha$ model, which assumes no dust.

Furthermore, we investigate the $H\gamma_N/H\beta_N$ Balmer decrement as a second check on the validity of Case B. In the middle panel of Fig. 7, we show a tentative detection ($\text{SNR} \approx 2.9$) of $H\gamma$ emission. Like the other Balmer lines, $H\gamma$ emission should have its own broad component, which would show absorption comparable in depth to that seen in $H\beta$ and deeper than that of $H\alpha$ (e.g., Naidu et al. 2025; D’Eugenio

et al. 2026). However, due to the low SNR in this region of the spectrum, it is only possible to meaningfully fit the narrow component of $H\gamma$ (henceforth $H\gamma_N$). As a consequence of the degeneracy between absorption and narrow components seen in the other Balmer lines, this suggests that our measurement of $H\gamma_N$ flux is likely an underestimate (unlike the other emission lines discussed in Section 3.2.3). However, our measurement yields $F(H\gamma)/F(H\beta) = 0.32^{+0.11}_{-0.10}$, consistent with the intrinsic Case B value of 0.466 (for a typical electron temperature of $T_e \sim 10^4$ K; Osterbrock & Ferland 2006) within 1.3σ .

Consequently, we cannot make a compelling claim for Case B violation, and henceforth adopt the assumption of no significant dust attenuation in *The Cliff*. This conclusion is supported by the previous studies which find no dust attenuation in the narrow lines of other LRDs (e.g. Lin et al. 2026c; Ji et al. 2026; D’Eugenio et al. 2025b; Nikopoulos et al. 2025).

We additionally note the $H\alpha_B/H\beta_B$ Balmer decrement is ~ 13

	Parameter	Value
H β -[O III] λ 5007	z_N	3.54903 ± 0.00004
	FWHM $_N$	90 ± 9 km s $^{-1}$
	z_B	3.546 ± 0.001
	FWHM $_B$	2770 ± 210 km s $^{-1}$
	v_{abs}	46^{+15}_{-14} km s $^{-1}$
	$\tau_0(\text{H}\beta)$	$3.22^{+2.26}_{-1.36}$
	C_f	$0.67^{+0.17}_{-0.13}$
	σ_{abs}	47 ± 21 km s $^{-1}$
H α	$z_{B,\text{H}\alpha}$	3.5486 ± 0.0001
	FWHM $_{B,\text{H}\alpha}$	480 ± 50 km s $^{-1}$
	T	$0.35^{+0.07}_{-0.06} \times 10^4$ K
	τ_{thom}	$2.22^{+0.19}_{-0.15}$
	$v_{\text{abs,H}\alpha}$	61 ± 4 km s $^{-1}$
	$\tau_0(\text{H}\alpha)$	$2.83^{+0.26}_{-0.24}$
	$\sigma_{\text{abs,H}\alpha}$	52 ± 6 km s $^{-1}$

Table 1. Summary of best-fit H β -[O III] λ 5007 and H α model parameters of *The Cliff*, reporting the median and 16th-84th percentile range of the posterior probability distributions. The ‘N’ and ‘B’ subscripts denote values associated with narrow and broad components, respectively. The free parameters for the emission line amplitudes and power-law continuum are not listed here. **Top segment:** Results from the H β -[O III] λ 5007 fit described in Section 3.2.1. **Bottom segment:** Results from the *Fiducial* H α fit described in Section 3.2.2. Parameters that were fixed to values from the H β -[O III] λ 5007 model as part of the H α fit are not repeated in this segment.

in this source. This high value may be associated with strong dust attenuation in the BLR ($A_V = 4$ mag assuming Case-B recombination and the SMC attenuation curve; Gordon et al. 2003), which can be achieved with relatively low amounts of dust concentrated on small spatial scales (Pacucci et al. 2026; Madau & Maiolino 2026). Alternatively, the high H α_B /H β_B flux ratio could arise from collisional excitation in high-density regions (see e.g. Yan et al. 2025). Regardless, the peculiar conditions of the BLR cannot apply to the narrow-line emitting gas, on account of its narrow spectral profile (FWHM $_N < 100$ km s $^{-1}$) and significantly lower Balmer decrement, implying the narrow-line emitting gas is physically distinct from the BLR.

4.2 Metallicity

4.2.1 Evidence for low metallicity in *The Cliff*

Following the standard approach for high-redshift sources within the literature, we infer the metallicity of *The Cliff* using strong-line metallicity diagnostics calibrated based on high- z star-forming galaxies (SFGs). In this work, we adopt the calibrations of Isobe et al. (in prep), which are based on stacks of SFGs over $1 < z < 10$ ($z_{\text{median}} \sim 3 - 5$), with linear extrapolations for metallicities $12 + \log(\text{O}/\text{H}) < 7$.




Figure 7. Portions of the integrated spectrum of *The Cliff*, illustrating, from top to bottom, the non-detections of the following emission lines: [O III] $\lambda\lambda$ 3727,3729, [O III] λ 4363, and [S II] $\lambda\lambda$ 6716,6731, respectively. The spectra presented here are continuum-subtracted to further emphasize the non detection of these lines. We do however note a marginal detection of H γ in the second panel (SNR = 2.9).

The [O III] λ 5007/H β_N flux ratio, denoted by R3, is often used to infer metallicity at high- z where other emission lines are not available (e.g. Curti et al. 2020; Sanders et al. 2025). Crucially, the R3 metallicity diagnostic has two branches, and a given R3 value could therefore correspond to both low-metallicity and high-metallicity solutions. Hence, to constrain the metallicity inferred from our measured R3 value, we must utilise the upper limits on the low-ionisation line fluxes determined earlier (see Section 3.2.3). The O32, N2, S2, O3N2 and O3S2 flux ratios (defined in Table 3) all have monotonic dependences on metallicity, and our observational limits on these flux ratios therefore provide upper limits on the metallicity of *The Cliff*. Furthermore, Isobe et al. (in prep.) have provided revised calibrations for these diagnostics, updating them for higher redshifts. We present these constraints from the low-ionisation line flux ratios in Fig. 8. Altogether, these diagnostics correspond to a metallicity upper limit of $Z \lesssim 0.25 Z_{\odot}$. While this upper limit is not particu-

Emission Line	Integrated Flux $10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2}$	SNR
H β_N	$0.65^{+0.09}_{-0.08}$	12.9
H β_B	$4.81^{+0.56}_{-0.45}$	18.8
[O III] λ 5007	1.02 ± 0.10	16.3
H α_N	1.85 ± 0.04	17.8
H α_B	$60.4^{+10.6}_{-9.3}$	133
H γ^*	0.21 ± 0.06	2.9
[O II] $\lambda\lambda$ 3727,3729	< 0.45	< 0.7
[O III] λ 4363	< 0.15	< 1.5
[N II] $\lambda\lambda$ 6548,6583	< 0.18	< 0.3
[S II] $\lambda\lambda$ 6716,6731	< 0.36	< 0.3

Table 2. Integrated fluxes and SNRs for the emission lines studied in this work. All fluxes have been corrected for aperture losses. H α fluxes refer to those determined from our fiducial Case B model (see Section 3.2.2); the low uncertainty arises because this flux is fixed based on the H β flux. H γ is marked with an asterisk as this represents a tentative detection; see Section 4.1 for further discussion. The lower portion of the table refers to undetected lines and therefore corresponds to upper limits on their fluxes, which are quoted at the 3σ upper limit.

larly constraining by itself, it is important for removing the potential degeneracy associated with the double-branched R3 diagnostic.

Fig. 9 shows the location of *The Cliff* on the R3 versus metallicity diagram (the filled magenta point), with the R3 calibration of Isobe et al. in prep shown by the solid blue line. The Isobe et al. (in prep.) R3 calibration provides a metallicity measurement of $0.017 \pm 0.004 Z_\odot$ for *The Cliff* (see Table 4). The upper branch solution would imply a super-solar metallicity, which would be implausible at such high redshift and for such a low mass system, but which is anyway ruled out by the weakness of the low-ionisation lines (see Fig. 8), as discussed above. The metallicity range excluded by the non-detection of the low ionisation lines is marked with a gray shaded region in Fig. 9.

We also note that the calibrations of Isobe et al. (in prep.) yield a higher metallicity for a given low-metal branch R3 value than other high-redshift strong-line metallicity calibrations (e.g. Cataldi et al. 2025; Sanders et al. 2024) by $\sim 0.2 - 0.3$ dex, so our choice of calibration provides the most conservative low-metallicity measurement.

The metallicity inferred from the R3 diagnostic is presented in Table 4, along with other properties of *The Cliff* measured later in this section. We note as a caveat that in making this metallicity measurement we are adopting an extrapolation of the calibration to metallicities $12 + \log(\text{O}/\text{H}) < 7$, similar to other studies of high-redshift, low-luminosity systems (e.g. Maiolino et al. 2025a). As an additional caveat, global measurements of metallicity at high redshift may be subject to systematic uncertainties driven by inhomogeneous ISM conditions (Belfiore et al. 2017; Katz et al. 2022; Moreschini et al. 2026). The measured emission line ratios can be skewed by localised high-excitation regions or dense clumps, potentially complicating the interpretation of galaxy-integrated scaling relations, and we cannot exclude the possibility this scenario applies to *The Cliff*.




Figure 8. Metallicity constraints on *The Cliff* inferred from upper limits on the non-detected emission lines ([O II] $\lambda\lambda$ 3727,3729, [S II] $\lambda\lambda$ 6716,6731 and [N II] $\lambda\lambda$ 6548,6583). The solid blue line in each panel illustrates the relevant diagnostic calibration from Isobe et al. (in prep.). The magenta symbols indicate the upper or lower limits obtained for flux ratios in *The Cliff* (see Table 3), which altogether exclude super-solar metallicity as an explanation for the measured [O III] λ 5007/H β_N flux ratio (see Fig. 9). The dashed vertical gray lines correspond to metallicities of $0.1 Z_\odot$ and Z_\odot , further emphasizing that a super-solar metallicity for *The Cliff* is inconsistent with our observations.

Flux Ratio	Abbreviation	Value
[for log(Ratio)]		
$\frac{[\text{O III}]\lambda 5007}{\text{H}\beta_{\text{N}}}$	R3	$1.59^{+0.26}_{-0.24}$
$\frac{\text{H}\alpha_{\text{N}}}{\text{H}\beta_{\text{N}}}$	-	$2.86^{+0.48}_{-0.36}$
$\frac{\text{H}\gamma_{\text{N}}}{\text{H}\beta_{\text{N}}}$	-	$0.32^{+0.11}_{-0.10}$
$\frac{[\text{O III}]\lambda 5007}{[\text{O II}]\lambda\lambda 3727, 3729}$	O32	> 2.3
$\frac{[\text{O III}]\lambda 4363}{[\text{O III}]\lambda 5007}$	-	< 0.15
$\frac{[\text{N II}]\lambda\lambda 6548, 6583}{\text{H}\alpha_{\text{N}}}$	N2	< 0.09
$\frac{[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}}{[\text{N II}]\lambda\lambda 6548, 6583/\text{H}\alpha_{\text{N}}}$	O3N2	> 18
$\frac{[\text{S II}]\lambda\lambda 6716, 6731}{\text{H}\alpha_{\text{N}}}$	S2	< 0.19
$\frac{[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}}{[\text{S II}]\lambda\lambda 6716, 6731/\text{H}\alpha_{\text{N}}}$	O3S2	> 8.3

Table 3. Flux ratios and their abbreviations for the narrow emission lines. Flux ratios are calculated based on aperture loss-corrected flux values (see Table 2). Uncertainties are bootstrapped from uncertainties on the fluxes and are quoted at the 68 percent confidence level. Upper or lower limits on flux ratios are quoted at the 3σ level. For further discussion of these calibrations, see e.g. Curti et al. 2020.

4.2.2 Accounting for potential AGN excitation

The extremely narrow line widths of $\text{H}\beta_{\text{N}}$ and $[\text{O III}]\lambda 5007$ measured in *The Cliff* ($\text{FWHM}_{\text{N}} \sim 90 \text{ km s}^{-1}$) hint that these narrow components are powered by weak star formation in the host galaxy (see e.g. D’Eugenio et al. 2025a). For the metallicity analysis above, we have therefore assumed that the narrow emission lines are entirely excited by star formation. In this section, we will now discuss the possible scenario in which the narrow lines are actually excited by the AGN and are therefore part of the NLR.

The calibrations obtained by Isobe et al. (in prep.) for high- z SFGs are not expected to change significantly in the case of ionisation by AGN rather than star formation. On average, high- z SFGs are characterised by high ionisation parameters ($\log U > -2.5$; see Section 4.2.4 for further discussion of ionisation parameter), and tend to have harder ionising spectra than local galaxies (Sanders et al. 2025; Cameron et al. 2023), primarily because of their reduced metal content (Cameron et al. 2024; Strom et al. 2017). High- z AGN are characterised by softer ionising spectra than local AGN, as indicated by the weakness of the high-ionisation lines (Lambrides et al. 2024; Zucchi et al. 2026). Overall, we expect the ionisation properties of the ISM in high- z SFGs and AGN NLRs to be similar (Kocevski et al. 2023; Übler et al. 2023; Juodžbalis et al. 2026), which is further evidenced by the significant overlap between high- z AGN and SFG populations in traditional diagnostic diagrams (e.g. the BPT diagrams; Übler et al. 2023; Maiolino et al. 2024b; Juodžbalis et al. 2026). Hence, the low metallicity inferred from the Isobe et al. (in prep.) calibrations may hold even in the case of AGN photoionisation of the narrow lines.

If anything, the harder-ionisation spectrum from an AGN would amplify the $[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}$ flux ratio for a given metallicity and ionisation parameter relative to SFGs (an illustration of this effect may be found in Appendix B of Maiolino et al. (2025a), for example). For the low-metallicity R3 branch, in the case of AGN photoionisation, a given $[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}$ flux ratio would yield an even lower

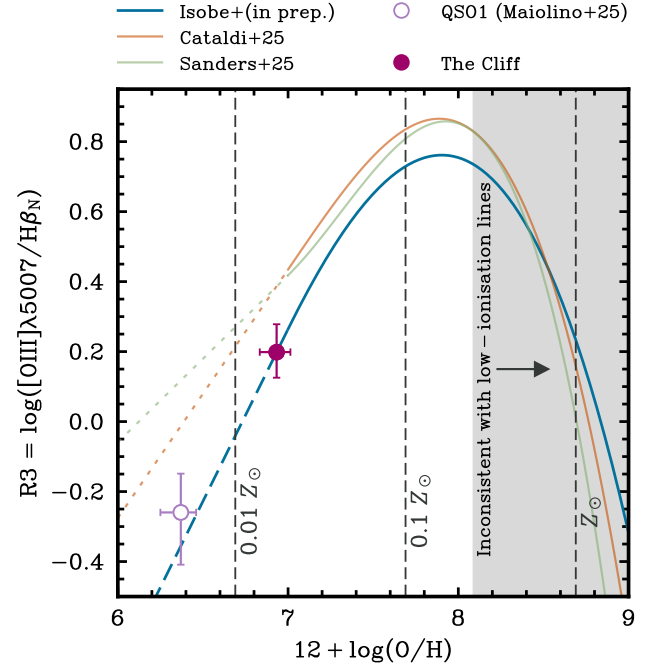


Figure 9. Metallicity of *The Cliff* as inferred from the measured $[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}$ flux ratio. The blue line illustrates the R3 calibration of Isobe et al. (in prep.), with the dashed segment of the line indicating the region for which the calibration is linearly extrapolated. The filled magenta point shows our favoured low-metallicity solution for *The Cliff* as derived from this calibration. The gray shaded area illustrates the high-metallicity branch ruled out by the low ionisation lines (Fig. 8). The turquoise point shows central region values for QSO1 from Maiolino et al. (2025a). The orange and green lines show the high-redshift strong-line metallicity calibrations of Cataldi et al. (2025) and Sanders et al. (2025) respectively, with dashed segments indicating where these calibrations are extrapolated. All error bars show uncertainties at the 1σ level.

metallicity than that inferred by assuming SF photoionisation (by $\sim 0.2 - 0.3$ dex). Once again, the assumptions we have adopted in our analysis yield the most conservative metallicity estimate.

4.2.3 Accounting for density

A potential alternative interpretation for the weakness of $[\text{O III}]\lambda 5007$ relative to $\text{H}\beta_{\text{N}}$ is that the gas density is so high that collisional de-excitation starts to affect the population of the $[\text{O III}]\lambda 5007$ levels.

The critical density of the $[\text{O III}]\lambda 5007$ line is about $5 \times 10^5 \text{ cm}^{-3}$. As discussed in Martinez et al. (2025) and Maiolino et al. (2025a), the emissivity of the $[\text{O III}]\lambda 5007$ line is only marginally affected at densities of 10^6 cm^{-3} (a decrease by a factor of only ~ 1.5), while it becomes prominent at densities of 10^7 cm^{-3} (suppressed by a factor of ~ 10). Although the densities of the ionised ISM are found to increase on average in high- z galaxies, they are still typically found to be $< 10^5 \text{ cm}^{-3}$ (Isobe et al. 2023; Topping et al. 2025).

The non-detection of the $[\text{O III}]\lambda 4363$ line (middle panel of Fig. 7), which has a significantly higher critical density than $[\text{O III}]\lambda 5007$, provides a key constraint on the density. Fig. 10 illustrates the variation of the $[\text{O III}]\lambda 4363/[\text{O III}]\lambda 5007$ flux ratio as a function of temperature and density, as estimated with PYNEB (Luridiana et al. 2015). For reasonable temperatures typical of photoionised gas emitting $\text{H}\beta$ ($T > 10^4 \text{ K}$, e.g. Osterbrock & Ferland 2006), the inferred upper limit of $F([\text{O III}]\lambda 4363)/F([\text{O III}]\lambda 5007) < 0.15$ yields a density upper




Figure 10. The $[\text{O III}]\lambda 4363/[\text{O III}]\lambda 5007$ flux ratio as a function of temperature (T) and density (n). The 3σ upper limit on the flux ratio, $F([\text{O III}]\lambda 4363)/F([\text{O III}]\lambda 5007) < 0.15$ (orange dot-dashed line), rules out densities $n > 6 \times 10^6 \text{ cm}^{-3}$ for reasonable temperatures $T > 10^4 \text{ K}$, as indicated by the gray shaded region. Note the values of $F([\text{O III}]\lambda 4363)/F([\text{O III}]\lambda 5007)$ quoted on each curve are logarithmic.

limit of $n < 7 \times 10^6 \text{ cm}^{-3}$. This is actually a stricter upper limit on the density than that found for QSO1 ($n < 10^7 \text{ cm}^{-3}$; Maiolino et al. 2025a), and excludes extreme suppression of $[\text{O III}]\lambda 5007$ as a consequence of collisional excitation. The 10^4 K limit is quite conservative, as most high- z galaxies are characterised by higher temperatures (Sanders et al. 2025; Cataldi et al. 2025), hence the upper limit on the density is likely close to 10^6 cm^{-3} .

Additional arguments against the high density scenario, following the methodology presented in Maiolino et al. (2025a), are discussed in Appendix D.

4.2.4 Excluding peculiar ionisation scenarios

Another potential scenario to explain the weakness of $[\text{O III}]\lambda 5007$ relative to $\text{H}\beta$, without invoking low density and metallicity, is to ascribe the effect to a gas ionisation state so low that O^{++} emission is suppressed. Quantitatively, photoionisation models indicate that to suppress the $[\text{O III}]/\text{H}\beta$ flux ratio to ~ 1 , the corresponding ionisation parameter would need to be $\log(U) \lesssim -2.5$ (e.g. Nagao et al. 2006; Baron & Netzer 2019; Cameron et al. 2023). This is in significant tension with established findings that high- z galaxies are typically characterised by high ionisation parameters ($\log(U) \gtrsim -2.0$; e.g. Cameron et al. 2023; Reddy et al. 2023; Tang et al. 2023). Of course, a single source could have lower than average $\log(U)$, but in this case the rarity of *The Cliff* would suggest a link between the spectral properties of the AGN and host galaxy. Assuming densities below the critical density of $[\text{O II}]\lambda\lambda 3727, 3729$ (i.e. $n = 10^4 \text{ cm}^{-3}$), we utilise our measurements of R3 and O32 and the AGN photoionisation models of Isobe et al. (2025) to lay a joint constraint on U and Z , as illustrated in Fig. 11. The lower limit we have obtained for O32 provides a lower limit on the ionisation parameter U as $\log(U) \gtrsim -3$.




Figure 11. Joint constraints on Z and the ionisation parameter U from the measured R3 ratio and lower limit on O32. The solid and dotted lines correspond to the results of AGN photoionisation models from Isobe et al. 2025, calculated for an assumed density of $n = 10^4 \text{ cm}^{-3}$. The blue lines illustrate curves of constant Z , decreasing from $\log(Z/Z_\odot) = 0$ (at the top of the plot) to $\log(Z/Z_\odot) = -2$ (at the bottom). Similarly, the dotted gray lines illustrate curves of constant U , increasing from $\log(U) = -3.5$ to $\log(U) = -0.5$ from left to right across the plot. The position of *The Cliff* is illustrated by the magenta circle, with the lower limit on O32 (see Table 3) providing a constraint on the minimum ionisation parameter.

Furthermore, *The Cliff* sits just below the model of $\log(Z/Z_\odot) = -1$ (i.e. $12 + \log(\text{O}/\text{H}) \lesssim 7.69$), providing a further conservative upper limit on the metallicity of *The Cliff*, which notably still sits on the low-metallicity branch of the R3 diagnostic regardless of high- z strong line calibration adopted (see Fig. 9).

A similar scenario, without invoking low metallicity or low ionisation parameter, is a lack of photons energetic enough to even ionise O^+ ($\sim 35 \text{ eV}$; see e.g. Wenaker 1990; Martin et al. 1993; Drainé 2011). However, this scenario would also not explain the absence of $[\text{O II}]\lambda\lambda 3727, 3729$ from the spectrum, as oxygen and hydrogen have similar ionisation potentials (Osterbrock & Ferland 2006). Similarly, $[\text{N II}]\lambda\lambda 6548, 6583$ and $[\text{S II}]\lambda\lambda 6716, 6731$ are typically expected to be strongly detected in such low-ionisation environments (Rhea et al. 2025). A potential caveat of this interpretation is that the abundance of nitrogen, and thus the strength of its associated emission lines, drops significantly faster than oxygen in low-metallicity regimes due to their different enrichment channels (Nicholls et al. 2017). Nevertheless, as previously established (Sections 3.2.2 and 3.2.3), neither of these doublets are significantly detected in the spectrum of *The Cliff*, let alone comparable to $\text{H}\beta$ or $[\text{O III}]\lambda 5007$, further disavouring a simple low-ionisation interpretation.

Finally, we consider a scenario in which the ionisation parameter is so high that oxygen primarily exists as O^{+++} . Such a scenario has never been observed, not even in the most extreme AGN (Dors et al. 2020). Such a high level of ionisation would also result in other high ionisation lines, such as C IV and He II (Wang et al. 2025a), appearing very strong; however, there is no evidence for these lines in the *The Cliff* (de Graaff et al. 2025b).

Galaxy Property		Value
Metallicity	$12 + \log(\text{O}/\text{H})$	$6.93^{+0.07}_{-0.09}$
	$\log(Z/Z_{\odot})$	$-1.75^{+0.07}_{-0.09}$
$\log(M_{\text{dyn}}/M_{\odot})$		8.41 ± 0.12
$\log(M_{*}/M_{\odot})$		< 8.41
$\log(M_{\text{BH}}/M_{\odot})$	Fiducial (from $\text{H}\beta$)	7.35 ± 0.24
	Scattering Scenario	6.32 ± 0.22
$\log(M_{\text{BH}}/M_{*})$	Fiducial (from $\text{H}\beta$)	> -1.06
	Scattering Scenario	> -2.07

Table 4. Summary of the properties of *The Cliff* as derived in this work. The metallicity was calculated using the strong-line R3 calibration of Isobe et al. in prep. (see Section 4.2.1). The fiducial BH mass is determined from the average of three different single-epoch $\text{H}\beta$ -based virial calibration BH mass estimates. M_{dyn} is taken as a conservative upper limit on M_{*} . Uncertainties are quoted at the 68 percent confidence level.

By ruling out both extreme high- and low- ionisation parameter scenarios, we conclude that the low $[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}$ flux ratio in *The Cliff* most likely originates from a metal-poor ISM.

4.3 Potential low metallicity satellite

Analysis of the $\text{H}\alpha$ emission in the vicinity of *The Cliff* reveals a candidate $\text{H}\alpha$ -emitting satellite located to the South, as illustrated in the left panel of Fig. 12. Its $\text{H}\alpha$ emission is redshifted by about $+380 \text{ km s}^{-1}$ relative to the narrow $\text{H}\alpha$ and $[\text{O III}]\lambda 5007$ components of *The Cliff* (Fig. 12; right panel). Notably, this satellite lacks detectable $[\text{O III}]$ or other emission features within the spectral coverage, indicative of low metallicity gas. The elongated morphology and redshifted velocity of the satellite may trace a tidally stretched, accreting system, potentially associated with the redshifted absorber of *The Cliff* identified in Section 3.2.2.

The detection of this potential satellite suggests that *The Cliff* resides in an environment characterised by the accretion of metal-poor gas (e.g. preceding a ‘compaction’ event; see Dekel et al. 2023; Mannucci et al. 2010), and this may explain the low metallicity of the galaxy.

However, while the $\text{H}\alpha$ identification is the most plausible explanation of the observed feature, we cannot exclude the possibility that it originates from a system at a different redshift which is not physically associated with *The Cliff*. Furthermore, the narrow emission lines of *The Cliff* exhibit a remarkably low velocity dispersion ($\sigma_{\text{N}} \sim 38 \text{ km s}^{-1}$), providing no clear evidence for the turbulence typically associated with an ongoing inflow. Higher-sensitivity, deeper follow-up observations are required to robustly confirm the physical association of the satellite with *The Cliff* and further constrain its nature.

4.4 Dynamical mass

In this work, we do not attempt to directly measure the stellar mass of *The Cliff*. M_{*} has proven difficult to measure across the wider LRD population, for reasons including but not limited to: degeneracies between AGN and stellar light (e.g. Labbe et al. 2024; Greene et al.

2024), non-stellar Balmer breaks (e.g. Ji et al. 2025), and systematic uncertainties in SED fitting (e.g. de Graaff et al. 2025b; Durodola et al. 2025). Additionally, stellar masses inferred from the optical continuum of LRDs may be unreliable, as some LRDs have been found to exhibit a time-variable rest-frame optical continuum (Ji et al. 2025; Naidu et al. 2025), suggesting the optical continuum may be AGN- rather than SF-dominated. In fact, in the case of *The Cliff*, de Graaff et al. (2025b) have already demonstrated the difficulties in measuring M_{*} by fitting the spectral energy distribution (SED) of this object, as SED models favour an implausibly high stellar mass, and stellar emission is insufficient to explain the strength of the observed Balmer break. Furthermore, attempts to measure M_{*} of *The Cliff* from its spectral continuum may be affected by contamination from the nearby foreground galaxy (see Fig. 1). Consequently, in this work we adopt the dynamical mass of *The Cliff* as a strong upper constraint on its stellar mass.

As discussed earlier, the small velocity dispersion inferred from $[\text{O III}]\lambda 5007$ and the narrow component of $\text{H}\beta$ (see Section 3.2.1 and Table 1) indicates that the dynamical mass of *The Cliff*’s host galaxy is likely fairly low. To constrain the dynamical mass, we adopt the same approach outlined in other high-redshift studies (e.g. Ji et al. 2025; Übler et al. 2023; Maiolino et al. 2024b; Ivey et al. 2026) of estimating the dynamical mass through the calibration of van der Wel et al. (2022):

$$M_{\text{dyn}} = \beta(n)K(q)\frac{\sigma_{*}^2 R_e}{G}, \quad (7)$$

where $\beta(n) = 8.87 - 0.831n + 0.0241n^2$ with Sérsic index n following Cappellari et al. (2006), and $K(q) = [0.87 + 0.38e^{-3.71(1-q)}]^2$ with axis ratio q following van der Wel et al. (2022). Given we lack information about q beyond the fit to the PSF, we adopt $q = 1$ as this yields a conservative estimate of the upper limit on the dynamical mass (taking $q = 0$ would reduce M_{dyn} by a factor of ~ 2). For the Sérsic index we assume $n = 1$ as is the case for the majority of high- z SFGs (Ormerod et al. 2024; Danhaive et al. 2026a), including AGN-host galaxies (Maiolino et al. 2024b; Juodžbalis et al. 2026). R_e is the effective radius, already estimated by de Graaff et al. (2025b) from morphological fitting of *JWST*/NIRCam F200W imaging to be $38.6^{+7.4}_{-6.9} \text{ pc}$, which we adopt here as we have determined *The Cliff* to be spatially unresolved. Finally, in Equation 7, σ_{*} is the stellar velocity dispersion. To estimate σ_{*} , we use our constraint on the narrow line width to measure the gas velocity dispersion, $\sigma_{\text{gas}} \approx 38 \text{ km s}^{-1}$, and then apply a correction of $\Delta \log(\sigma_{\text{gas}}/[\text{km s}^{-1}]) = +0.1$ following Übler et al. (2023), since galaxies with low integrated ionised gas velocity dispersion tend to have a higher integrated stellar velocity dispersion (Bezanson et al. 2018)³. By adopting these assumptions, we obtain $2.6 \times 10^8 M_{\odot}$ as a conservative upper limit on the dynamical mass of *The Cliff*. Hence, we obtain an upper limit on the stellar mass of $M_{*} < 2.6 \times 10^8 M_{\odot}$ (Table 4). We note that dark matter and gas fractions are likely to be high (see e.g. Danhaive et al. 2026a; McClymont et al. 2026), hence this is a very conservative upper limit.

Inayoshi & Maiolino (2025) have recently proposed that dense gas clouds in the vicinity of an AGN accretion disc can give rise to strong Balmer series absorption, including both line and bound-free absorption, i.e. the same phenomena happening in the atmosphere of hot stars ($T_{\text{eff}} \sim 10^4 \text{ K}$). Ji et al. (2025) have verified that such

³ Tests of this relation at high-redshift (e.g. Carnall et al. 2023; D’Eugenio et al. 2024; Pescalau et al. 2026a) indicate that it does appear to hold at high redshift, though this result refers to systems much more massive than *The Cliff*.

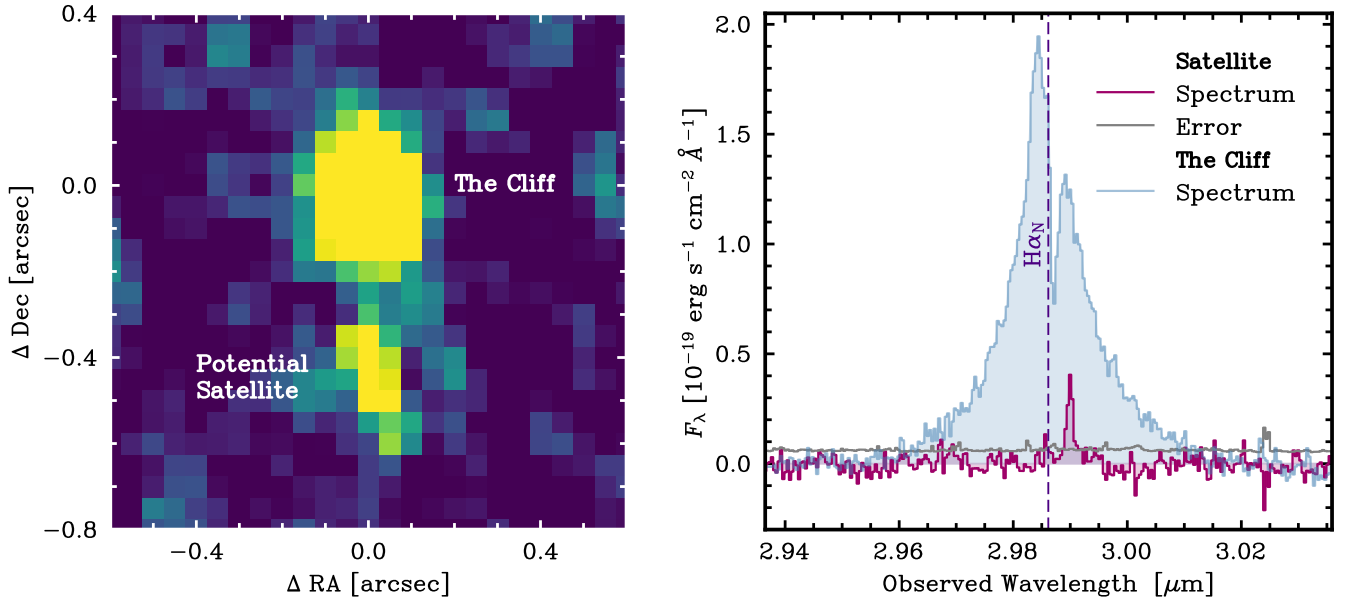


Figure 12. Image and integrated spectrum of the potential southern satellite of *The Cliff*. **Left panel:** The continuum-subtracted data cube collapsed over the 2.9888 – 2.9907 μm spectral channels, highlighting both *The Cliff* and the potential southern $\text{H}\alpha$ satellite (both labelled). **Right panel:** Comparison of continuum-subtracted integrated spectra of *The Cliff* and the satellite. The dashed vertical line indicates the wavelength of the $\text{H}\alpha_{\text{N}}$ component in *The Cliff*, demonstrating the $\text{H}\alpha$ emission associated with the satellite has a redshifted offset. Hence, this satellite could correspond to an inflow of pristine gas towards *The Cliff*, and may therefore be associated with the observed low metallicity of this system.




Figure 13. Constraints on the BH mass of *The Cliff* from several different calibrations, based on both $\text{H}\alpha$ and $\text{H}\beta$ (Greene & Ho 2005; Vestergaard & Peterson 2006; Parlanti et al. 2025; Reines & Volonteri 2015). Our $\text{H}\beta$ BH mass measurements and associated 1σ uncertainties are shown by the topmost three points; these measurements are combined to derive our fiducial BH mass (the magenta point), which is quoted in Table 4. We also show the ‘electron scattering’ scenario BH mass as determined from our $\text{H}\alpha$ fit, following the argument of Rusakov et al. (2026); this represents a more conservative estimate of the BH mass, and is also quoted in Table 4. The dashed gray line and shaded region show the BH mass estimate from de Graaff et al. (2025b) along with its 1σ uncertainty.

a model can indeed produce the strong Balmer break and Balmer absorption lines such as those we observe in *The Cliff* and in the case of QSO1, without the need to invoke massive, evolved stellar populations. Although QSO1 has a weaker Balmer break than *The Cliff*, a similar argument certainly applies here, meaning a small M_* is entirely reasonable.

4.5 Black hole mass

The inferred metallicity provides important constraints for models and simulations, but becomes more powerful when accompanied by measurements of BH mass.

In this work, we estimate the black hole mass of *The Cliff* by assuming the local virial relations between BH mass, and the broad line luminosity (e.g. Greene & Ho 2005; Vestergaard & Peterson 2006; Reines & Volonteri 2015), which are calibrated either through reverberation mapping or direct measurements. It is not obvious that these locally-calibrated relations apply at high redshift. In fact, we note that BH mass measurement at high- z remains a contentious topic, with arguments that *JWST*-measured BH masses are inaccurately estimated by these calibrations, for example due to super-Eddington accretion (e.g. Marconi et al. 2008, 2009; Lambrides et al. 2024; Lupi et al. 2024) or electron scattering (e.g. Rusakov et al. 2026; Naidu et al. 2025; Chang et al. 2026).

The black hole mass of *The Cliff* was previously estimated by de Graaff et al. (2025b) as $\log(M_{\text{BH}}/M_{\odot}) = 7.18^{+0.07}_{-0.06}$, based on the profile of broad $\text{H}\alpha$. However, the medium-resolution spectra result in a lack of strong constraints on the narrow line kinematics or the Balmer line absorption component, which have now been revealed by the G235H observations.

In this work, as we have the best constraints on the narrow and broad components of $\text{H}\beta$ (compared to $\text{H}\alpha$), we primarily utilise the results of our $\text{H}\beta$ fit to determine an estimate of M_{BH} . We take




Figure 14. The location of *The Cliff* on the $M_{\text{BH}}-M_*$ plane, for both the fiducial $\text{H}\beta$ -measured BH mass (magenta circle) and the scattering scenario BH mass (black circle). The grey points correspond to measurements from other *JWST* observations of low mass AGN (Juodžbalis et al. 2026, 2025; Harikane et al. 2023; Kokorev et al. 2023; Furtak et al. 2024; Maiolino et al. 2024a; Carnall et al. 2023) and quasars (Stone et al. 2024; Yue et al. 2024a; Ding et al. 2023), along with the sample of $3 < z < 7$ broad-line AGN from Jones et al. (2025) and the mean stacks of low-mass AGN from Geris et al. (2026). The solid purple line shows the local scaling relation from Reines & Volonteri (2015), with its scatter indicated by the shaded region. The other purple lines indicate constant M_{BH}/M_* ratios. Regardless of the BH mass adopted, *The Cliff* is overmassive relative to the local relation, a conclusion made stronger by noting that M_* presented here is a conservative upper limit estimated from M_{dyn} (see Section 4.4).

FWHM_{B} as measured from our spectral fit to the $\text{H}\beta$ -[O III] λ 5007 complex, and calculate the luminosity of broad $\text{H}\beta$. We utilise 3 different $\text{H}\beta$ -based single epoch calibrations to determine the BH mass: Greene & Ho (2005); Vestergaard & Peterson (2006) and Parlanti et al. (2025). For the Vestergaard & Peterson (2006) and Parlanti et al. (2025) calibrations, when calculating uncertainties we incorporate the quoted intrinsic scatters of 0.43 and 0.40 dex, respectively.

There is a spread of ~ 0.2 dex across the M_{BH} values obtained from the various $\text{H}\beta$ calibrations, but all of the measurements are consistent within 1σ , and this spread is lower than the $\sim 0.3 - 0.4$ dex scatter on the BH mass calibrations overall. Hence, to obtain a single $\text{H}\beta$ -determined value for M_{BH} , we combine these three measurements with an uncertainty-weighted average and propagate their uncertainties, to obtain the *Mean* BH mass measurement of $\log(M_{\text{BH}}/M_{\odot}) = 7.35 \pm 0.24$. This *fiducial* BH mass measurement will be used in discussions and analysis throughout the rest of this work.

Noting that the $\text{H}\alpha$ line profile was parameterised by a model compatible with the proposed electron scattering scenario (Rusakov et al. 2026), we obtain an alternative, more conservative estimate of BH mass, by using $\text{FWHM}_{\text{B,H}\alpha} = 480 \pm 50 \text{ km s}^{-1}$ as associated with the intrinsic Gaussian. For this calculation, we utilise the single-epoch $\text{H}\alpha$ calibrations of Reines & Volonteri (2015). This yields a BH mass an order of magnitude smaller than our fiducial estimate, $\log(M_{\text{BH}}/M_{\odot}) = 6.32 \pm 0.22$. However, we emphasise that exponential profiles are not necessarily tracing electron scattering and that simple models of the (virialised) BLR can also naturally result

in exponential profiles (Scholtz et al. 2026; Madau et al. 2026). In particular, Madau et al. (2026) have recently shown that the black hole masses derived by modelling LRD exponential profiles through BLR stratification, without invoking electron scattering, are fully consistent with standard single epoch estimates.

Finally, one could measure BH mass from $\text{H}\alpha$ using the calibration of Reines & Volonteri (2015), now by using the overall FWHM of the broad component as estimated from its overall profile, $920 \pm 110 \text{ km s}^{-1}$ (which is still significantly lower than the FWHM of $\text{H}\beta_{\text{N}}$, $2770 \pm 210 \text{ km s}^{-1}$). This approach yields a measurement of $\log(M_{\text{BH}}/M_{\odot}) = 6.91 \pm 0.23$, which is consistent with both our fiducial, mean $\text{H}\beta$ BH mass measurement and the earlier estimate of de Graaff et al. (2025b) (within 2σ). For clarity, in this work we will limit further discussion to the $\text{H}\beta$ -measured and $\text{H}\alpha$ ‘scattering scenario’ masses as these correspond to the extreme scenarios and therefore enable a better exploration of the parameter space.

In Fig. 13, we present an overview of the BH masses of *The Cliff* determined in this work (the various coloured points); we also compare to the earlier measurement from de Graaff et al. (2025b) (the dotted gray line, with a shaded region illustrating the 1σ uncertainty). Our fiducial, mean $\text{H}\beta$ -estimated BH mass (the magenta circle) is consistent with the value from de Graaff et al. (2025b) within 1σ . We will discuss the implications of our BH mass estimates in Section 5.1.

5 DISCUSSION

5.1 An overmassive black hole

Beyond the challenges for stellar mass measurements discussed in the previous section (Section 4.4), as, discussed, BH mass measurements for LRDs are also subject to critical systematic uncertainties. While virial, locally-calibrated single-epoch SMBH mass estimators are widely adopted to measure BH masses for *JWST*-discovered AGN, their reliability is limited by extensive systematic uncertainties in BLR geometry, ionisation and kinematics. Furthermore, single-epoch BH mass estimates can vary by up to ~ 0.5 dex depending on the line used (e.g. $H\alpha$, $H\beta$) and the adopted calibration (e.g. Dalla Bontà et al. 2025). These systematics are particularly important when considering LRDs, which are effectively a new class of AGN. Direct measurements of BH mass are essential for independently validating the single epoch estimators for high redshift AGN and LRDs. However, while a few direct BH mass measurements exist outside of the local Universe (e.g. Abuter et al. 2024; Shen et al. 2024; Liao et al. 2025; Gravity+ Collaboration et al. 2026), the only direct M_{BH} measurement in an LRD to date is that of QSO1 at $z \approx 7.04$ by Juodžbalis et al. (2025). M_{BH} of QSO1 was shown to be consistent with that from the single-epoch mass estimates, but this single measurement is insufficient for making any broad conclusions about whether the single epoch estimators are applicable to the wider early AGN population.

With these caveats on virial BH masses in mind, we compare the upper limit on M_* of *The Cliff* (Section 4.4) to M_{BH} inferred from the virial relations (Section 4.5). We interpret this comparison in the context of both local BH-galaxy scaling relations and of other high-redshift *JWST*-observed AGN. Fig. 14 illustrates the relationship between BH mass and host galaxy stellar mass. We include *The Cliff* in this figure (the magenta circle), together with other high- z *JWST* AGN sources from the literature whose masses were also estimated with single-epoch virial calibrations (or direct dynamical measurement in the case of QSO1; Juodžbalis et al. 2025). We find that even when adopting the lower of our two BH mass estimates (the BH mass derived assuming the scattering scenario, shown by the black circle), *The Cliff* remains overmassive relative to the local scaling relations. If we instead adopt the fiducial BH mass inferred from $H\beta$, *The Cliff* appears even more overmassive ($M_{\text{BH}}/M_* \gtrsim 0.1$), though we note that both measured BH masses render *The Cliff* consistent with the wider population of high- z *JWST* AGN shown in this figure. It is crucial to note that as we have adopted an upper limit on M_* , the black hole of *The Cliff* is likely even more overmassive than our measurements suggest, as indicated by the arrows on the plot. This argument is further strengthened by recent findings that $H\alpha$ profiles with exponential wings are not direct evidence of electron scattering (Brazzini et al. 2025, 2026), because they can alternatively be explained by stratification of the BLR (Scholtz et al. 2026; Madau et al. 2026).

5.2 Metal-poor LRDs and deviation from the MZR

As we have noted across the previous sections, *The Cliff* exhibits a number of observational similarities to QSO1. This raises a number of questions: Do the formation pathways producing objects such as *The Cliff* operate across cosmic time? Is low metallicity characteristic of unusual LRDs and therefore an indicator of rare pathways for BH seeding and evolution?

To further investigate the connection between *The Cliff* and other extreme LRDs, Fig. 15 illustrates comparison of the mass-metallicity

relation (MZR) to ‘unusual’ LRDs across three different redshift regimes: $z \sim 0$, $3 < z < 4$ and $6.5 < z < 7.5$. We note as a caveat for this discussion that this comparison requires extrapolation of MZRs to lower stellar masses. The center panel of Fig. 15, presents a comparison of *The Cliff* to the $3 < z < 4$ MZR determined from the calibrations of Isobe et al. (in prep.). Adopting the same calibrations for both the MZR and the strong-line metallicity measurement ensures a self-consistent comparison. Visually, *The Cliff* is clearly offset from the MZR at this redshift; its deviation of $\sim 11\sigma$ from the best-fit MZR statistically confirms its status as an outlier. Indeed, marginal consistency with the MZR (i.e. a 3σ deviation) would require a stellar mass low enough to attain $M_{\text{BH}}/M_* \sim 9$, a ratio larger than the lower limit determined for QSO1 ($M_{\text{BH}}/M_* > 2$), which has been identified be a ‘nearly naked’ black hole in an early stage of its evolution (Juodžbalis et al. 2025). Even in the case of the electron scattering black hole mass, such a scenario results in $M_{\text{BH}}/M_* \sim 1$. This hints at a fundamental tension between the MZR and $M_{\text{BH}} - M_*$ relation for *The Cliff*.

In the left panel of Fig. 15, we present a comparison of the so-called ‘local LRDs’ ($z \sim 0.1 - 0.22$ Lin et al. 2026c) to the local MZR from Curti et al. (2020); these are 3 of a mere handful of local LRDs found to date (Lin et al. 2025, 2026c) and therefore, like *The Cliff*, represent a rare population. We adopt the direct metallicity measurements of Lin et al. (2026c) for this comparison. We additionally estimate the dynamical masses of the local LRDs of Lin et al. (2026c) as an upper limit on their stellar masses using Equation 7, adopting the narrow-line FWHMs measured by Lin et al. (2026c) and the tightest available constraint on R_e for each LRD (i.e. from Lin et al. 2026c; Ji et al. 2026). With these upper limits on stellar mass, even the highest-metallicity of the three local LRDs deviates from the median $z \sim 0$ MZR by more than 5σ . The deviation from the MZR is yet more significant for the other two local LRDs.

Finally, in the rightmost panel of Fig. 15, we present a comparison of QSO1 ($z \sim 7$) to the $6.5 < z < 7.5$ MZR (also from Isobe et al. in prep.). We re-estimated the metallicity of QSO1 based on the Maiolino et al. (2025a) measurement of $F([\text{O III}]\lambda 5007)/F(H\beta_{\text{N}}) = 0.55 \pm 0.16$, now using the Isobe et al. (in prep.) R3 calibration to obtain a self-consistent comparison to the MZR. At the upper M_* limit afforded by M_{dyn} , we find QSO1 is also a clear outlier from the $6.5 < z < 7.5$ MZR by $\gtrsim 7\sigma$.

Taken together, these comparisons reveal a fundamental tension between extreme LRDs and the established galaxy-black hole scaling relations. The statistical significance with which *The Cliff*, QSO1 and local LRDs deviate from the MZR suggests that standard frameworks for galaxy assembly do not apply to these systems. Indeed, these LRDs may follow a unique evolutionary trajectory, in which the processes governing black hole growth are decoupled from the metal enrichment and stellar assembly of the host galaxy.

5.3 The $M_{\text{BH}} - M_* - Z$ tension

Fig. 16 illustrates the key tensions highlighted by the previous sections. The dynamical mass of *The Cliff* represents an upper limit on its stellar mass and a lower M_* would therefore align *The Cliff* more closely with the $3 < z < 4$ MZR (Fig. 16; left panel). However, reducing M_* simultaneously pushes *The Cliff* further into the ‘overmassive’ regime on the $M_{\text{BH}} - M_*$ plane. This holds true even when adopting a more conservative BH mass estimate calculated based on an electron scattering scenario, increasing the tension with the local $M_{\text{BH}} - M_*$ relation (Fig. 16; centre panel). Consequently, under standard assumptions, objects like *The Cliff* cannot simultaneously follow both scaling relations. Possible resolutions to this predic-

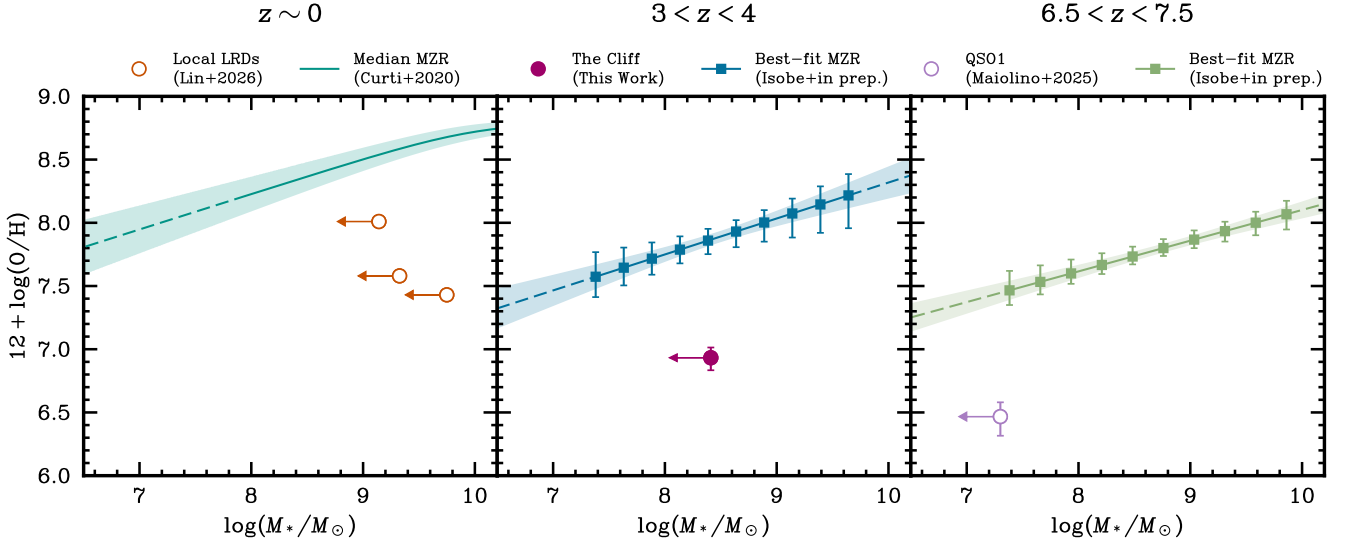


Figure 15. The mass-metallicity relation (MZR) for three different redshift ranges. For each plotted MZR, the dashed segments show the regions for which the relation is extrapolated, and the shaded area indicates the 1σ uncertainty of the relation. **Left panel:** The median MZR for the local Universe (teal line), derived by Curti et al. (2020) based on galaxies from the Sloan Digital Sky Survey (SDSS). The unfilled orange circles illustrates the position of the local LRDs (Lin et al. 2026c) on the mass-metallicity plane, with M_{dyn} taken as an upper limit on M_* , following Equation 7. Here we plot the direct- T_e method metallicities derived by Lin et al. (2026c); the corresponding uncertainties are 0.01 – 0.03 dex, smaller than is visible on the scale of this plot. **Center panel:** The MZR for galaxies at $3 < z < 4$ (blue line), as derived by Isobe et al. (in prep.). The position of *The Cliff* ($z \sim 3.5$) is shown by the filled magenta point, with 1σ uncertainties on its metallicity and M_{dyn} plotted as an upper limit on M_* . **Right panel:** The MZR for galaxies at $6.5 < z < 7.5$ (green line), as derived by Isobe et al. (in prep.). QSO1 ($z \sim 7$) is shown by the unfilled purple point, using the stellar mass upper limit from Juodžbalis et al. (2025), and a metallicity derived from the $[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}$ flux ratio measured by Maiolino et al. (2025a) via the Isobe et al. (in prep.) R3 calibration.




Figure 16. A sketch illustrating the key tensions discussed in this work (see Section 5.2 for a full discussion). The black stars illustrate ‘real measurements’ of observational properties for an LRD such as *The Cliff*, with the hollow stars corresponding to hypothetical scenarios. In brief, from left to right: for an object such as *The Cliff*, using M_{dyn} as an upper limit on M_* means that if the true M_* is substantially lower, the system will be in closer agreement with the corresponding MZR. However, this lower M_* would also increase M_{BH}/M_* , so matching the MZR would make the BH even more overmassive relative to the host galaxy. Matching both the MZR and the local $M_{\text{BH}}-M_*$ relations requires the virial M_{BH} to be overestimated by ~ 4 dex, a factor much larger than can be attributed to even the most extreme cases of electron scattering. A black hole mass this much smaller would increase the inferred Eddington ratio to $\lambda_{\text{Edd}} \geq 100$, resulting in a system too exotic to obey the typical BH scaling relations, and indicating that systems such as *The Cliff* fundamentally deviate from the local scaling relations.

ment involve invoking extreme density conditions (Section 4.2.3), or assuming black hole masses are not measured accurately by the virial relations.

If the virial BH masses are indeed overestimated, the tension with the local $M_{\text{BH}}-M_*$ relation is reduced - however, this would require the BH mass to be overestimated by four orders of magnitude, much higher than then even the most extreme electron scattering scenarios. Additionally, this would require a extremely high Eddington ratios

($\lambda_{\text{Edd}} \geq 100$; right panel of Fig. 16). Such extreme Eddington ratios are virtually absent even from large-scale quasar surveys, which show that even the most rapidly accreting BHs rarely exceed $\lambda_{\text{Edd}} \sim 1-3$, and ratios significantly above 10 remain extreme outliers (e.g. Kollmeier et al. 2006; Jin et al. 2012).

In any case, if one accepts the possibility of such high λ_{Edd} , there is no compelling reason to expect such objects would obey local scaling relations, invalidating the assumptions that led us to this extreme

picture. Consequently, our findings suggest that these extreme LRDs, and potentially other high- z *JWST* AGN, inherently deviate from the local $M_{\text{BH}} - M_*$ relation. This deviation is likely driven by selection biases which systematically miss lower M_{BH} sources (e.g. Geris et al. 2026; Ziparo et al. 2026), while simulations indicate that these early overmassive black holes eventually evolve onto the local relation over time (Hu et al. 2026).

As highlighted by Fig. 15, *The Cliff* is not the only object for which this set of tensions exists, pointing to a formation channel which can produce massive BHs in metal-poor environments across cosmic times. While low metallicity does appear to be characteristic of certain LRDs such as *The Cliff* or QSO1, the evolutionary pathways that produce (or indeed maintain) such metal-poor states remain effectively unconstrained. These systems could have transitioned from an earlier metal-rich phase to their observed metal-poor state due to the efficient ejection of metals by galactic outflows; alternatively, the ISM could have been diluted by significant pristine gas inflows (e.g. in the case of *The Cliff*, by accretion from a metal-poor satellite; see Section 4.3). Evaluating the feasibility of such pathways is limited by uncertainties associated with factors including metal transport and outflow efficiencies, for example. Consequently, numerical simulations of black hole formation and evolution provide a vital framework for interpreting the chemical histories and BH seeding pathways of LRDs such as *The Cliff*.

5.4 Comparison to hydrodynamical simulations

In many scenarios of BH-galaxy co-evolution, it is unlikely for a BH to grow to a mass $> 10^7 M_{\odot}$ in a metal-free environment. This follows from the fact that the conditions conducive to rapid black hole growth (i.e. abundant cold gas accretion) would also catalyse star formation, resulting in rapid metal enrichment. Hence, the findings of this work (and those of Maiolino et al. 2025a and Jones et al. 2026, in the case of QSO1) are indicative of a route that could allow a black hole to form and grow within a chemically unevolved stellar system, seemingly at odds with any simple picture of galaxy–black hole co-evolution.

Due to the complexity and stochasticity of the processes involved in co-evolution, any quantitative approach to the question of whether black holes can grow significantly in pristine environments requires a simulation-based approach. In this Section, we compare our findings for *The Cliff* to predictions from the AESOPICA simulations (Koudmani et al., in prep.; see Appendix E for a full description). We stress that these simulations incorporate stellar and BH feedback capable of gas and metal removal, as well as merger and gas accretion capable of ISM dilution. For the purposes of this work, we utilise two different simulation modes: (i) fiducial Eddington-limited accretion following FABLE (*‘Fiducial’*; Henden et al. 2018) and (ii) boosted accretion with super-Eddington accretion up to ten times the Eddington limit (*‘SuperEdd+BoostAcc’*).

Fig. 17 shows the location of *The Cliff* on the $M_{\text{BH}} - Z$ plane (first column of panels) and $M_{\text{BH}}/M_* - Z$ plane (second column). The large symbols show the location of *The Cliff* as inferred from our analysis: both the single-epoch virial BH mass (magenta circle) and the BH mass inferred in the electron-scattering scenario (black circle), plotted against our measurement of the metallicity. Each panel shows the results of four different simulation runs, each associated with a different BH seed mass; the range of seed masses explored in these simulations is $10^2 - 10^5 M_{\odot}$.

In the $M_{\text{BH}} - Z$ plane, even when adopting the lower electron-scattering M_{BH} estimate, only simulations utilising the highest seed masses ($10^5 M_{\odot}$) successfully approach the observed parameter space of *The Cliff*; this discrepancy is exacerbated when adopting the

fiducial M_{BH} estimate. In particular, the mismatch relative to lower-mass seeds is so large that it is unlikely to be salvaged by earlier seeding. While a small subset of the highest seed mass simulations recover metallicities and BH masses comparable to *The Cliff*, the simulation runs incorporating boosted accretion exhibit greater metallicity scatter in this plane, suggesting that super-Eddington growth phases may be a key evolutionary component for extreme objects like *The Cliff*. Tracking the specific inflow and outflow histories of the simulated outliers could provide key insights into the mechanisms governing the evolution of *The Cliff*.

The $M_{\text{BH}}/M_* - Z$ plane proves to be less constraining as *The Cliff* overlaps more broadly with the parameter space of the simulation results, although heavy seeds ($10^4 - 10^5 M_{\odot}$) remain favoured. Crucially, as the stellar mass presented here corresponds to a strict upper limit, the true position of *The Cliff* in this plane likely shifts toward the highest seed mass simulation runs or indeed the extreme outliers of the simulated population, depending on the BH mass estimate adopted. Collectively, these constraints point toward a rare evolutionary pathway for *The Cliff*, defined by heavy black hole seeding and periods of super-Eddington growth.

To evaluate whether gas inflows are predominantly responsible for the low- Z outliers in the simulations, tracking the evolutionary history of individual outliers is required. Regardless of the underlying driver of these outliers, the left-hand panels of Fig 17 demonstrate that *The Cliff* remains ~ 1 dex below the simulation $M_{\text{BH}} - Z$ relation, even for heavy seed simulation runs. Adopting an earlier seeding prescription would not resolve this tension, as the $M_{\text{BH}} - Z$ relation appears to evolve towards higher metallicities as M_{BH} increases.

To gain context for the rareness of objects such as *The Cliff* in the simulations compared to the rareness of *The Cliff* itself, it is key to account for the survey and simulation volumes. The AESOPICA simulation boxes have side length 60 cMpc, giving a total simulation volume of $2.16 \times 10^5 \text{ cMpc}^3$ for the snapshot at $z \sim 3.5$. Adopting the same cosmology used by AESOPICA (see Henden et al. 2018) and the survey area of $\sim 150 \text{ arcmin}^2$ (de Graaff et al. 2025a), the RUBIES survey volume over the redshift range $3 < z < 4$ is $5.06 \times 10^5 \text{ cMpc}^3$, i.e. about 2.3 times larger than the simulation volume. *The Cliff* is an extremely unique object within the RUBIES survey, and seems similarly rare within the simulations, and we conclude it is fair to evaluate the underlying seeding and evolutionary mechanisms as similarly rare.

One alternative scenario involving primordial black hole seeds (black holes theorised to form shortly after the Big Bang) may naturally produce such rare systems with overmassive BHs in chemically unevolved environments (for a review, see Carr & Kühnel 2020). Recent hydrodynamical simulations by Zhang et al. (2025b,e) have explored galaxy formation around massive PBH seeds with masses of $\sim 10^6 - 5 \times 10^7 M_{\odot}$ (a mass scale motivated by Carr et al. 2021a, for example), which act as deep gravitational potential wells for halo formation and subsequent baryonic inflow. In these models, strong accretion-driven feedback heats the surrounding gas and suppresses or delays stellar assembly. Once star formation begins, feedback-driven outflows expel metal-enriched gas, while continued baryonic inflow replenishes the system with pristine material. Such simulations reproduce several qualitative properties inferred for objects like *The Cliff* or QSO1, including extremely low metallicities ($Z \lesssim 0.01 Z_{\odot}$), compact morphologies ($R_e \lesssim 50 \text{ pc}$), and BH-dominated mass budgets ($M_{\text{BH}}/M_* \gtrsim 0.1$). Although these simulations currently stop at higher redshift, $z \sim 7 - 9$, and are carried out in a small volume of $\sim 1 \text{ cMpc}^3$ (targeting a single halo around the PBH), they nonetheless illustrate that PBH seeds could provide a viable evolu-

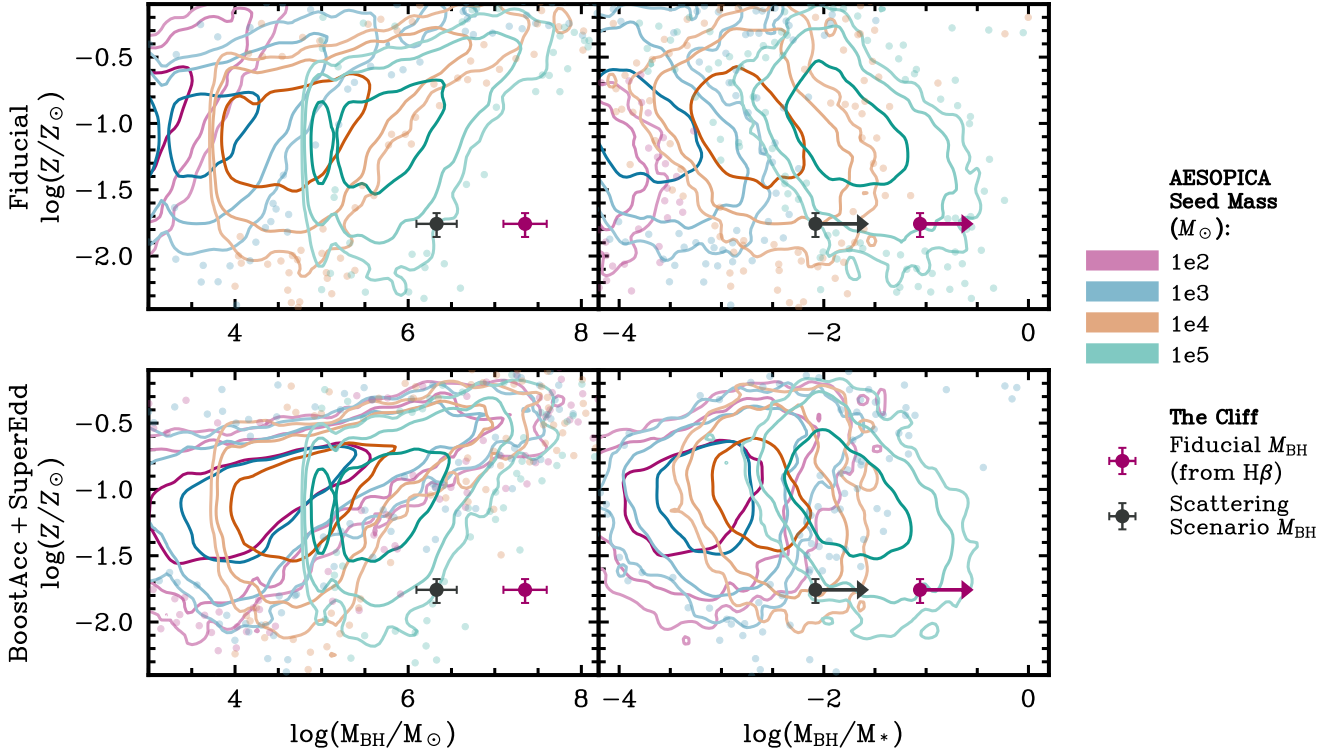


Figure 17. Comparison of *The Cliff* to results from AESOPICA hydrodynamical simulations run down to $z \sim 3.5$ on the $M_{\text{BH}} - Z$ (left) and $M_{\text{BH}}/M_* - Z$ (right) diagrams. Further descriptions of the simulations are provided in Section 5.4 and Appendix E. Contours enclose 68%, 95%, and 99% of the simulations; simulations outside of the 99% contours are plotted as individual points. The large magenta and black circles illustrate the properties of *The Cliff* for our fiducial and scattering scenario BH masses, respectively.

tionary pathway for near-pristine LRDs.⁴ In fact, sources like *The Cliff* could present an opportunity to empirically constrain PBH abundance as a fraction of dark matter density, a key cosmological parameter (see e.g. Carr et al. 2021b) which to date remains poorly understood. For simplicity, if one assumes that the comoving number density of such sources is approximately the inverse survey volume, i.e. $\sim 2 \times 10^{-6} \text{ cMpc}^{-3}$, then the corresponding PBH seed density fraction is roughly $f_{\text{PBH,seed}} \sim 3.2 \times 10^{-11} (M_{\text{PBH}}/10^6 M_{\odot})$. This estimate is well below current observational upper limits over the mass range $\sim 10^6 - 5 \times 10^7 M_{\odot}$. If only a fraction of such objects are seeded by PBHs, the implied constraint would become correspondingly tighter.

Due to the high computational costs and variety of relevant physical processes involved in running simulations down to low redshift, very few simulations targeted at BH formation and evolution in the early Universe reach $z \sim 3.5$, so we caution that the models presented here represent a limited comparison and are certainly not intended to be exhaustive. As an additional caveat for AESOPICA, we must stress that the simulations are not necessarily developed with assumptions aimed at reproducing BHs similar to *The Cliff*, nor do they directly model the multi-phase ISM structure which would be crucial to fully unravelling the origin of these pristine galaxies with overmassive BHs. Finally, it is crucial to note that simulations and observations do not infer metallicities in the same way, which can result in significant systematic differences between observational measurements

and measurements derived from simulations. For example, with forward modelling of metallicity, Payoor Vijayan et al. (2026) note a discrepancy between simulation-inferred and strong-line metallicities of up to 0.3 dex, which they find is driven by dust attenuation and properties of the underlying stellar population. Ultimately, future models and simulations will cover a wider parameter space, pushing to higher resolution and lower redshifts, offering a more comprehensive comparison with rare objects such as *The Cliff* and offering better insight into the underlying physics.

6 CONCLUSIONS

We have presented new, high-resolution *JWST*/NIRSpec-IFU observations of *The Cliff*, a remarkable Little Red Dot at $z \sim 3.55$. These observations have offered valuable new insights into the ISM properties and chemical enrichment of *The Cliff*. In this work we have focused on the narrow components of the emission lines, $\text{H}\beta$ and $[\text{O III}]\lambda 4959,5007$ in particular, with the goal of obtaining constraints on BH mass and metallicity which could be combined with simulations to investigate potential BH growth and seeding pathways. The main findings of our analysis are as follows:

- The $\text{H}\beta$ and $[\text{O III}]\lambda 5007$ emission are spatially compact and unresolved beyond the scale of the PSF.
- The high-resolution spectral observations reveal deep, redshifted absorption components in both $\text{H}\alpha$ and $\text{H}\beta$, indicative of inflowing gas.
- We additionally identify a potential satellite south of *The Cliff* detected only in $\text{H}\alpha$, possibly tidally stretched, which could trace

⁴ For details of the numerical setup and analysis, the reader should refer to Zhang et al. (2025e,b).

pristine gas inflows onto *The Cliff*. However, we note that confirming this satellite and performing a detailed characterisation will require deeper follow-up observations.

- In the galaxy-integrated spectrum of *The Cliff*, the $[\text{O III}]\lambda 5007$ line is weak relative to the narrow component of $\text{H}\beta$ emission compared to the bulk of the galaxy population at the same redshift, with a flux ratio $F([\text{O III}]\lambda 5007)/F(\text{H}\beta_{\text{N}}) = 1.59^{+0.26}_{-0.24}$.

- Adopting the updated strong-line metallicity calibrations of Isobe et al. (in prep.), we estimate the metallicity of *The Cliff* as $Z \approx 0.017 \pm 0.004 Z_{\odot}$.

- This finding is reinforced by the non-detection of key low-ionisation lines ($[\text{O II}]\lambda\lambda 3727, 3729$, $[\text{N II}]\lambda\lambda 6548, 6583$ and $[\text{S II}]\lambda\lambda 6716, 6731$). Alternative explanations for the low $[\text{O III}]\lambda 5007/\text{H}\beta_{\text{N}}$ ratio, such as collisional saturation or extreme ionisation parameters, are disfavoured.

- We identify *The Cliff* as a significant outlier from the $3 < z < 4$ MZR (also derived from the calibrations of Isobe et al. in prep.), appearing exceptionally metal-poor for its stellar mass limit ($\log(M_*/M_{\odot}) < 8.41$).

- By fitting the Balmer lines, we derive two separate black hole mass measurements: a fiducial virial mass of $\log(M_{\text{BH}}/M_{\odot}) = 7.35 \pm 0.24$, and a scattering scenario mass of $\log(M_{\text{BH}}/M_{\odot}) = 6.32 \pm 0.22$ (from $\text{H}\alpha$).

- In both scenarios, the black hole of *The Cliff* is overmassive relative to the local $M_{\text{BH}} - M_*$ scaling relation. Reconciling *The Cliff* with the MZR would require an even lower stellar mass, further exacerbating its deviation from the $M_{\text{BH}} - M_*$ relation.

- The observed tension between the MZR and the $M_{\text{BH}} - M_*$ relation could be resolved by a black hole mass overestimated by four orders of magnitude and an extreme super-Eddington growth phase, but the typical galaxy-black hole scaling relations would not be expected to hold for such a scenario. It is more likely that MZR deviation and overmassive BHs are in some way characteristic of the evolutionary pathway underpinning *The Cliff*, meaning *The Cliff* inherently deviates from these relations.

- This tension is not limited to *The Cliff*, but is also shown to apply to the local LRDs (Lin et al. 2026c) and Abell2744-QSO1 (Maiolino et al. 2025a; Juodžbalis et al. 2025), pointing to a potential rare evolutionary pathway shared by these unusual LRDs which operates across cosmic time.

- While some simulations of BH growth from heavy seeds ($10^4 - 10^5 M_{\odot}$) with periods of super-Eddington accretion can capture the high M_{BH} and M_{BH}/M_* , or low metallicity, of *The Cliff*, these outcomes are rare. The existence of such a metal-poor, overmassive-BH system at $z \sim 3.55$ remains difficult to capture within existing evolutionary frameworks. While a Primordial Black Hole origin for *The Cliff* may be a possibility, this scenario requires further, more quantitative exploration.

The constraints on the metallicity and black hole mass of *The Cliff*, when combined with simulations of BH growth and evolution, suggest a rare BH seeding and growth pathway underpins this system. Without the context of a larger sample of similar systems, it remains unclear whether the MZR deviation and overmassive BH of *The Cliff* are typical of other LRDs. Future simulations incorporating a resolved, multi-phase ISM will be essential to disentangle the mechanisms fuelling the growth of massive black holes within such metal-poor environments. Until then, *The Cliff* represents a critical benchmark for understanding the formation and early evolution of Little Red Dots.

ACKNOWLEDGEMENTS

LRI, FDE, RM, XJ, GCJ and JS acknowledge support by the Science and Technology Facilities Council (STFC), European Research Council (ERC) Advanced Grant 695671 “QUENCH” and the UKRI Frontier Research grant RISEandFALL. RM also acknowledges funding from Royal Society research professorship. YI is supported by JSPS KAKENHI grant number 24KJ0202. IJ acknowledges support by the Huo Family Foundation through a P.C. Ho PhD Studentship. SK is supported by a Junior Research Fellowship from St Catharine’s College, Cambridge and a Research Fellowship from the Royal Commission for the Exhibition of 1851. MP acknowledges support through the grants PID2021-127718NB-I00, PID2024-159902NA-I00, and RYC2023-044853-I, funded by the Spain Ministry of Science and Innovation/State Agency of Research MCIN/AEI/10.13039/501100011033 and El Fondo Social Europeo Plus FSE+. SZ, VB and BL acknowledge the Texas Advanced Computing Center (TACC) for providing HPC resources under allocation AST23026. AJB acknowledges funding from the “FirstGalaxies” Advanced Grant from the ERC under the European Union’s Horizon 2020 research and innovation program (Grant agreement No. 789056). SC acknowledges support by EU HE ERC Starting Grant No. 101040227 - WINGS. KI acknowledges support from the National Natural Science Foundation of China (12573015, 1251101148, 12233001, 12473037), and the China Manned Space Program (CMS-CSST-2025-A09). BL acknowledges the funding of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC 2181/1 - 390900948 (the Heidelberg STRUCTURES Excellence Cluster). RP acknowledges funding support from a STFC PhD studentship. PR and BR acknowledge support from JWST/NIRCam contract to the University of Arizona NAS5-02105. ST acknowledges support from the Royal Society Research Grant G125142.

This work is based on observations made with the NASA/ESA/CSA James Webb Space Telescope. The data were obtained from the Mikulski Archive for Space Telescopes (MAST) at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST. These observations are associated with JWST program ID 9433.

DATA AVAILABILITY

The datasets used in this work were derived from a source in the public domain. The JWST/NIRSpec data can be downloaded from the MAST (<https://mast.stsci.edu/portal/Mashup/Clients/Mast/Portal.html>) under PID 9433. Photometry was obtained from PID 1837 and PID 7814.

REFERENCES

- Abuter R., et al., 2024, *Nature*, 627, 281
 Adamo A., et al., 2025, *Nature Astronomy*, 9, 1134
 Ananna T. T., Bogdán Á., Kovács O. E., Natarajan P., Hickox R. C., 2024, *ApJ*, 969, L18
 Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481
 Baggen J. F. W., et al., 2024, *ApJ*, 977, L13
 Baron D., Netzer H., 2019, *MNRAS*, 486, 4290
 Begelman M. C., Dexter J., 2026, *ApJ*, 996, 48
 Begelman M. C., Volonteri M., Rees M. J., 2006, *MNRAS*, 370, 289
 Belfiore F., et al., 2017, *MNRAS*, 469, 151
 Bezanson R., et al., 2018, *ApJ*, 868, L36

- Brazzini M., D'Eugenio F., Maiolino R., Juodžbalis I., Ji X., Scholtz J., Chang S.-J., 2025, *MNRAS*, 544, L167
- Brazzini M., et al., 2026, *arXiv e-prints*, p. [arXiv:2601.22214](#)
- Cameron A. J., et al., 2023, *A&A*, 677, A115
- Cameron A. J., Katz H., Witten C., Saxena A., Laporte N., Bunker A. J., 2024, *MNRAS*, 534, 523
- Cammelli V., Monaco P., Tan J. C., Singh J., Fontanot F., De Lucia G., Hirschmann M., Xie L., 2025, *MNRAS*, 536, 851
- Cappellari M., et al., 2006, *MNRAS*, 366, 1126
- Carnall A. C., et al., 2023, *Nature*, 619, 716
- Carr B., Kühnel F., 2020, *Annual Review of Nuclear and Particle Science*, 70, 355
- Carr B., Clesse S., García-Bellido J., Kühnel F., 2021a, *Physics of the Dark Universe*, 31, 100755
- Carr B., Kohri K., Sendouya Y., Yokoyama J., 2021b, *Reports on Progress in Physics*, 84, 116902
- Cataldi E., et al., 2025, *A&A*, 703, A208
- Chang S.-J., Gronke M., Matthee J., Mason C., 2026, *MNRAS*, 545, [staf2131](#)
- Chen C.-H., Ho L. C., Li R., Zhuang M.-Y., 2025, *ApJ*, 983, 60
- Chisholm J., et al., 2024, *MNRAS*, 534, 2633
- Curti M., Mannucci F., Cresci G., Maiolino R., 2020, *MNRAS*, 491, 944
- D'Eugenio F., et al., 2024, *Nature Astronomy*, 8, 1443
- D'Eugenio F., et al., 2025a, *arXiv e-prints*, p. [arXiv:2503.11752](#)
- D'Eugenio F., et al., 2025b, *arXiv e-prints*, p. [arXiv:2510.00101](#)
- D'Eugenio F., et al., 2026, *MNRAS*, 545, [staf2117](#)
- Dalla Bontà E., et al., 2025, *A&A*, 696, A48
- Danhaive A. L., et al., 2026a, *MNRAS*,
- Danhaive A. L., et al., 2026b, *MNRAS*, 546, [stag119](#)
- Dayal P., 2024, *A&A*, 690, A182
- Dayal P., Maiolino R., 2026, *A&A*, 706, A72
- Dekel A., Sarkar K. C., Birnboim Y., Mandelker N., Li Z., 2023, *MNRAS*, 523, 3201
- Dekel A., et al., 2025a, *arXiv e-prints*, p. [arXiv:2511.07578](#)
- Dekel A., Stone N. C., Chowdhury D. D., Gilbaum S., Li Z., Mandelker N., van den Bosch F. C., 2025b, *A&A*, 695, A97
- Dimitrijević M. S., Popović L. Č., Kovačević J., Dačić M., Ilić D., 2007, *MNRAS*, 374, 1181
- Ding X., et al., 2023, *Nature*, 621, 51
- Dojčinović I., Kovačević-Dojčinović J., Popović L. Č., 2023, *Advances in Space Research*, 71, 1219
- Dors O. L., Maiolino R., Cardaci M. V., Hägele G. F., Krabbe A. C., Pérez-Montero E., Armah M., 2020, *MNRAS*, 496, 3209
- Draine B. T., 2011, *Physics of the Interstellar and Intergalactic Medium*. Princeton University Press
- Durodola E., Pacucci F., Hickox R. C., 2025, *ApJ*, 985, 169
- Euclid Collaboration et al., 2025, *arXiv e-prints*, p. [arXiv:2503.15323](#)
- Ferrara A., Salvadori S., Yue B., Schleicher D., 2014, *MNRAS*, 443, 2410
- Ferruit P., et al., 2022, *A&A*, 661, A81
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Furtak L. J., et al., 2024, *Nature*, 628, 57
- Furtak L. J., et al., 2025, *A&A*, 698, A227
- Geris S., et al., 2026, *MNRAS*, 545, [staf1979](#)
- Gordon K. D., Clayton G. C., Misselt K. A., Landolt A. U., Wolff M. J., 2003, *ApJ*, 594, 279
- Gravity+ Collaboration et al., 2026, *A&A*, 706, A99
- Greene J. E., Ho L. C., 2005, *ApJ*, 630, 122
- Greene J. E., et al., 2024, *ApJ*, 964, 39
- Grogin N. A., et al., 2011, *ApJS*, 197, 35
- Habouzit M., 2025, *MNRAS*, 537, 2323
- Hainline K. N., et al., 2025, *ApJ*, 979, 138
- Harikane Y., et al., 2023, *ApJ*, 959, 39
- Henden N. A., Puchwein E., Shen S., Sijacki D., 2018, *MNRAS*, 479, 5385
- Hu H., Yanagisawa H., Nishigaki M., Kiyota T., Ishiyama T., Ohsuga K., 2026, *arXiv e-prints*, p. [arXiv:2602.14496](#)
- Hviding R. E., et al., 2025, *A&A*, 702, A57
- Hviding R. E., et al., 2026, *arXiv e-prints*, p. [arXiv:2601.09778](#)
- Inayoshi K., 2025, *ApJ*, 988, L22
- Inayoshi K., Maiolino R., 2025, *ApJ*, 980, L27
- Inayoshi K., Visbal E., Haiman Z., 2020, *ARA&A*, 58, 27
- Isobe Y., Ouchi M., Nakajima K., Harikane Y., Ono Y., Xu Y., Zhang Y., Umeda H., 2023, *ApJ*, 956, 139
- Isobe Y., et al., 2025, *MNRAS*, 541, L71
- Ivey L. R., et al., 2026, *MNRAS*, 546, [stag094](#)
- Jakobsen P., et al., 2022, *A&A*, 661, A80
- Jeon J., Bromm V., Liu B., Finkelstein S. L., 2025, *ApJ*, 979, 127
- Jeon J., et al., 2026, *ApJ*, 998, 148
- Ji X., et al., 2024, *MNRAS*, 535, 881
- Ji X., et al., 2025, *MNRAS*, 544, 3900
- Ji X., et al., 2026, *MNRAS*, 545, [staf2235](#)
- Jin C., Ward M., Done C., 2012, *MNRAS*, 425, 907
- Jones B. L., et al., 2025, *arXiv e-prints*, p. [arXiv:2510.07376](#)
- Jones G. C., et al., 2026, *MNRAS*, 546, [stag115](#)
- Juodžbalis I., et al., 2024a, *MNRAS*, 535, 853
- Juodžbalis I., et al., 2024b, *Nature*, 636, 594
- Juodžbalis I., et al., 2025, *arXiv e-prints*, p. [arXiv:2508.21748](#)
- Juodžbalis I., et al., 2026, *MNRAS*,
- Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
- Katz H., et al., 2022, *arXiv e-prints*, p. [arXiv:2211.04626](#)
- Kido D., Ioka K., Hotokezaka K., Inayoshi K., Irwin C. M., 2025, *MNRAS*, 544, 3407
- Kocevski D. D., et al., 2023, *ApJ*, 954, L4
- Kocevski D. D., et al., 2025, *ApJ*, 986, 126
- Koekemoer A. M., et al., 2011, *ApJS*, 197, 36
- Kokorev V., et al., 2023, *ApJ*, 957, L7
- Kokubo M., Harikane Y., 2025, *ApJ*, 995, 24
- Kollmeier J. A., et al., 2006, *ApJ*, 648, 128
- Koudmani S., Sijacki D., Smith M. C., 2022, *MNRAS*, 516, 2112
- Labbe I., et al., 2024, *arXiv e-prints*, p. [arXiv:2412.04557](#)
- Lambrides E., et al., 2024, *arXiv e-prints*, p. [arXiv:2409.13047](#)
- Laor A., 2006, *ApJ*, 643, 112
- Latif M. A., Ferrara A., 2016, *Publ. Astron. Soc. Australia*, 33, e051
- Li J., et al., 2025, *ApJ*, 981, 19
- Liao M., et al., 2025, *arXiv e-prints*, p. [arXiv:2504.13409](#)
- Lin R., et al., 2025, *ApJ*, 980, L34
- Lin R., et al., 2026a, *arXiv e-prints*, p. [arXiv:2603.01473](#)
- Lin X., et al., 2026b, *ApJ*, 996, 93
- Lin X., et al., 2026c, *ApJ*, 997, 364
- Liu H., Jiang Y.-F., Quataert E., Greene J. E., Ma Y., 2025, *ApJ*, 994, 113
- Loiacono F., et al., 2025, *A&A*, 703, A36
- Lupi A., Trinca A., Volonteri M., Dotti M., Mazzucchelli C., 2024, *A&A*, 689, A128
- Luridiana V., Morisset C., Shaw R. A., 2015, *A&A*, 573, A42
- Ma Y., et al., 2025, *arXiv e-prints*, p. [arXiv:2504.08032](#)
- Madau P., Maiolino R., 2026, *arXiv e-prints*, p. [arXiv:2602.22386](#)
- Madau P., Maiolino R., Scholtz J., D'Eugenio F., 2026, *arXiv e-prints*, p. [arXiv:2604.04216](#)
- Maiolino R., et al., 2024a, *Nature*, 627, 59
- Maiolino R., et al., 2024b, *A&A*, 691, A145
- Maiolino R., et al., 2025a, *arXiv e-prints*, p. [arXiv:2505.22567](#)
- Maiolino R., et al., 2025b, *MNRAS*, 538, 1921
- Mannucci F., Cresci G., Maiolino R., Marconi A., Gnerucci A., 2010, *MNRAS*, 408, 2115
- Marconi A., Axon D. J., Maiolino R., Nagao T., Pastorini G., Pietrini P., Robinson A., Torricelli G., 2008, *ApJ*, 678, 693
- Marconi A., Axon D. J., Maiolino R., Nagao T., Pietrini P., Risaliti G., Robinson A., Torricelli G., 2009, *ApJ*, 698, L103
- Martin W. C., Kaufman V., Musgrove A., 1993, *Journal of Physical and Chemical Reference Data*, 22, 1179
- Martinez Z., et al., 2025, *ApJ*, 995, 204
- Matthee J., et al., 2024, *ApJ*, 963, 129
- Mazzolari G., et al., 2024, *arXiv e-prints*, p. [arXiv:2412.04224](#)
- Mazzolari G., et al., 2025, *A&A*, 700, A12
- McClummont W., et al., 2026, *MNRAS*, 545, [staf2092](#)
- Mezcua M., Pacucci F., Suh H., Siudek M., Natarajan P., 2024, *ApJ*, 966, L30
- Moreschini B., et al., 2026, *arXiv e-prints*, p. [arXiv:2601.08939](#)

- Muzzin A., et al., 2025, *arXiv e-prints*, p. [arXiv:2507.19706](#)
- Nagao T., Maiolino R., Marconi A., 2006, *A&A*, 459, 85
- Naidu R. P., et al., 2025, *arXiv e-prints*, p. [arXiv:2503.16596](#)
- Nandal D., Buldgen G., Whalen D. J., Regan J., Woods T. E., Tan J. C., 2025, *A&A*, 701, A262
- Natarajan P., Pacucci F., Ricarte A., Bogdán Á., Goulding A. D., Cappelluti N., 2024, *ApJ*, 960, L1
- Netzer H., Shemmer O., Maiolino R., Oliva E., Croom S., Corbett E., di Fabrizio L., 2004, *ApJ*, 614, 558
- Newville M., Stensitzki T., Allen D. B., Rawlik M., Ingargiola A., Nelson A., 2016, Lmfit: Non-Linear Least-Square Minimization and Curve-Fitting for Python, Astrophysics Source Code Library, record ascl:1606.014 (ascl:1606.014)
- Nicholls D. C., Sutherland R. S., Dopita M. A., Kewley L. J., Groves B. A., 2017, *MNRAS*, 466, 4403
- Nikopoulos G. P., Watson D., Sneppen A., Rusakov V., Heintz K. E., Witstok J., Brammer G., 2025, *arXiv e-prints*, p. [arXiv:2510.06362](#)
- Ormerod K., et al., 2024, *MNRAS*, 527, 6110
- Osterbrock D. E., Ferland G. J., 2006, Astrophysics of gaseous nebulae and active galactic nuclei
- Pacucci F., Loeb A., 2024, *ApJ*, 964, 154
- Pacucci F., Nguyen B., Carniani S., Maiolino R., Fan X., 2023, *ApJ*, 957, L3
- Pacucci F., Ferrara A., Kocevski D. D., 2026, *arXiv e-prints*, p. [arXiv:2601.14368](#)
- Parlanti E., et al., 2025, *arXiv e-prints*, p. [arXiv:2512.14844](#)
- Partmann C., Naab T., Lahén N., Rantala A., Hirschmann M., Hislop J. M., Petersson J., Johansson P. H., 2025, *MNRAS*, 537, 956
- Pascalau R. G., et al., 2026a, *MNRAS*,
- Pascalau R. G., et al., 2026b, *arXiv e-prints*, p. [arXiv:2603.00232](#)
- Pasha I., Miller T. B., 2023, *The Journal of Open Source Software*, 8, 5703
- Payoor Vijayan A., et al., 2026, *The Open Journal of Astrophysics*, 9, 57554
- Perna M., et al., 2023, *A&A*, 679, A89
- Perna M., et al., 2025, *A&A*, 696, A59
- Perrin M. D., Long J., Sivaramakrishnan A., Lajoie C.-P., Elliot E., Pueyo L., Albert L., 2015, WebbPSF: James Webb Space Telescope PSF Simulation Tool, Astrophysics Source Code Library, record ascl:1504.007 (ascl:1504.007)
- Planck Collaboration et al., 2020, *A&A*, 641, A6
- Prole L. R., et al., 2025, *arXiv e-prints*, p. [arXiv:2511.09640](#)
- Rantala A., Lahén N., Naab T., Escobar G. J., Iorio G., 2025, *MNRAS*, 543, 2130
- Rauscher B. J., et al., 2017, *PASP*, 129, 105003
- Reddy N. A., Topping M. W., Sanders R. L., Shapley A. E., Brammer G., 2023, *ApJ*, 952, 167
- Reines A. E., Volonteri M., 2015, *ApJ*, 813, 82
- Rhea C. L., et al., 2025, *AJ*, 169, 203
- Rinaldi P., et al., 2025, *arXiv e-prints*, p. [arXiv:2507.17738](#)
- Rodríguez Del Pino B., et al., 2024, *A&A*, 684, A187
- Rusakov V., et al., 2026, *Nature*, 649, 574
- Sanati M., et al., 2025, *MNRAS*, 542, 1532
- Sanders R. L., Shapley A. E., Topping M. W., Reddy N. A., Brammer G. B., 2024, *ApJ*, 962, 24
- Sanders R. L., et al., 2025, *arXiv e-prints*, p. [arXiv:2508.10099](#)
- Scholtz J., et al., 2025, *A&A*, 697, A175
- Scholtz J., et al., 2026, *arXiv e-prints*, p. [arXiv:2603.22277](#)
- Schwarz G., 1978, *Annals of Statistics*, 6, 461
- Setton D. J., et al., 2025, *ApJ*, 995, 118
- Shajib A. J., Treu T., Melo A., Roberts-Borsani G., Knabel S., Cappellari M., Frieman J. A., 2025, *A&A*, 702, L12
- Shen Y., et al., 2024, *ApJS*, 272, 26
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575
- Smith A., Bromm V., 2019, *Contemporary Physics*, 60, 111
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Stone M. A., Lyu J., Rieke G. H., Alberts S., Hainline K. N., 2024, *ApJ*, 964, 90
- Storey P. J., Hummer D. G., 1995, *MNRAS*, 272, 41
- Strom A. L., Steidel C. C., Rudie G. C., Trainor R. F., Pettini M., Reddy N. A., 2017, *ApJ*, 836, 164
- Tang M., et al., 2023, *MNRAS*, 526, 1657
- Taylor A. J., et al., 2025a, *ApJ*, 986, 165
- Taylor A. J., et al., 2025b, *ApJ*, 989, L7
- Topping M. W., et al., 2025, *MNRAS*, 541, 1707
- Trefoloni B., et al., 2025, *A&A*, 700, A203
- Übler H., et al., 2023, *A&A*, 677, A145
- Vestergaard M., Peterson B. M., 2006, *ApJ*, 641, 689
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M., et al., 2014, *MNRAS*, 444, 1518
- Wang B., et al., 2025a, *arXiv e-prints*, p. [arXiv:2508.18358](#)
- Wang B., et al., 2025b, *ApJ*, 984, 121
- Wenaker I., 1990, *Phys. Scr.*, 42, 667
- Wiersma R. P. C., Schaye J., Smith B. D., 2009a, *MNRAS*, 393, 99
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, *MNRAS*, 399, 574
- Yan Z., Inayoshi K., Chen K., Guo J., 2025, *arXiv e-prints*, p. [arXiv:2512.11050](#)
- Yue M., et al., 2024a, *ApJ*, 966, 176
- Yue M., Eilers A.-C., Ananna T. T., Panagiotou C., Kara E., Miyaji T., 2024b, *ApJ*, 974, L26
- Zhang Z., et al., 2025a, *arXiv e-prints*, p. [arXiv:2512.05180](#)
- Zhang S., Liu B., Bromm V., Kühnel F., 2025b, *arXiv e-prints*, p. [arXiv:2512.14066](#)
- Zhang Z., Jiang L., Liu W., Ho L. C., 2025c, *ApJ*, 985, 119
- Zhang S., Liu B., Bromm V., Jeon J., Boylan-Kolchin M., Kühnel F., 2025d, *ApJ*, 987, 185
- Zhang S., Liu B., Bromm V., Jeon J., Boylan-Kolchin M., Kühnel F., 2025e, *ApJ*, 987, 185
- Ziparo F., Gallerani S., Ferrara A., 2025, *J. Cosmology Astropart. Phys.*, 2025, 040
- Ziparo F., Carniani S., Gallerani S., Trefoloni B., 2026, *arXiv e-prints*, p. [arXiv:2603.04358](#)
- Zucchi G., Ji X., Madau P., Maiolino R., Juodžbalis I., D’Eugenio F., Geris S., Isobe Y., 2026, *A&A*, 707, A52
- de Graaff A., et al., 2025a, *A&A*, 697, A189
- de Graaff A., et al., 2025b, *A&A*, 701, A168
- van der Wel A., et al., 2022, *ApJ*, 936, 9

APPENDIX A: MODELLING THE NIRSPEC/IFS POINT SPREAD FUNCTION

Our investigation of *The Cliff* exploits the presence of a star within the instrument FoV (see Fig. 1). To provide independent validation of the radial surface brightness profiles presented in Fig. 2, we create a synthetic median image of the star in a 0.1- μm spectral window centred on *The Cliff*’s [O III] λ 5007 emission (Fig. A1a). We then model this image with the Bayesian modelling tool `PYSERSIC` (Pasha & Miller 2023), with the setup outlined in D’Eugenio et al. (2026), and using a Sérsic profile with fixed index $n = 0.5$ (i.e. representing a Gaussian). We tested a range of n values and convolution methods to assess the stability of this model.

Our results are presented in Fig. A1, where we show the PSF model (panel b), the relative residuals (panel c), and the circularised radial profile (panel d).




Figure A1. Modelling the star near *The Cliff* with `pYSERVIC`. From left to right, from top to bottom, the panels show the data, best-fit model, relative residuals, and the circularised radial profile, using `asinh` scaling, to highlight the PSF wings. See Fig. 2 for a comparison of the radial profile of this best-fit model to the radial profile of various emission lines of *The Cliff*.




Figure B1. Narrow-line $[\text{O III}]\lambda 5007$ relative velocity map of *The Cliff*, presented with a SNR cut of 5. The best-fit kinematic PA is shown by the black dashed line, while the instrumental PA is shown by the dotted magenta line. Even though the observed velocity gradient is perpendicular to the slices (as expected from an instrument artefact), its sign is opposite to the instrumental effect we infer from modelling the interloper ID 24647.

APPENDIX B: MODELLING THE NIRSPEC/IFS INSTRUMENT GRADIENT

At the time of writing this article, reconstructed NIRSpec/IFS datacubes are known to display an artificial velocity gradient⁵. This gradient is thought to be introduced by slice-to-slice astrometric drift, therefore any kinematic signal aligned with the NIRSpec position angle (P.A.) must be interpreted with caution (as discussed in e.g., Juodžbalis et al. 2025). *The Cliff* itself was observed with a NIRSpec P.A. of 198.31° , and we measure a velocity gradient of -15 km s^{-1}

⁵ for more detail, see <https://jwst-docs.stsci.edu/known-issues-with-jwst-data/nirspec-known-issues/nirspec-ifu-known-issues/> and e.g. Pascalau et al. 2026b.




Figure B2. Portions of integrated spectra extracted from a $0.15''$ radius circular aperture centred on ID 24647 (see Fig. 1), showing emission lines detected in this foreground galaxy. **Top panel:** The $\text{H}\beta$ - $[\text{O III}]\lambda 5007$ spectral region. **Bottom panel:** The $\text{H}\alpha$ spectral region, also showing a detection of $[\text{N II}]\lambda\lambda 6548, 6583$.

along the east–west direction, i.e. extremely close to the instrument P.A. (Fig. B1). Given this alignment and the compactness of *The Cliff*, there is therefore a non-negligible chance that the observed gradient may be artificial.

To assess this possibility, we exploit the serendipitous presence of ID 24647, an inclined galaxy within the NIRSpec FoV (Fig. 1). This source displays strong $\text{H}\beta$, $[\text{O III}]\lambda\lambda 4959, 5007$, $\text{H}\alpha$ and $[\text{N II}]\lambda\lambda 6548, 6583$ emission (see Fig. B2), and is therefore suitable for kinematic modelling. We here present a simple model to assess whether the velocity gradient observed in *The Cliff* is artificial and if galaxy kinematics can be used to constrain the artificial gradient. For this demonstration we combine tilted-disc kinematics, an arctan rotation curve and a constant instrument gradient; we do not take into account the non-uniform light profile of ID 24647, nor the instrument PSF. While both effects are important, their careful modelling is beyond the exploratory scope of this appendix.

We consider a fourteen parameter model: seven parameters describe the emission-line map (centre, total flux, half-light semi-major axis, projected axis ratio, Sérsic index, and position angle), five are for the galaxy kinematics (kinematic position angle, inclination, intrinsic axis ratio, maximum velocity v_{max} , and turnover radius), and

σ deviation from the prior). In Model 3 we free both $\nabla v_{\text{instr.}}$ and $P.A._{\text{instr.}}$. This model finds $\nabla v_{\text{instr.}} = +7.7 \pm 0.5 \text{ km s}^{-1} \text{ spaxel}^{-1}$ and $P.A._{\text{instr.}} = 159^{+5}_{-3}^\circ$, different from the NIRSpec P.A. (8σ). While statistically consistent with the data (and achieving $\Delta \text{BIC} = -9$ and -1 for [O III] and H α , respectively), Model 3 is optimised by collapsing v_{max} to only $7^{+17}_{-6} \text{ km s}^{-1}$, which we deem unlikely for a galaxy with $\log(M_*/M_\odot) = 9.6 \pm 0.1$ (Muzzin et al. 2025); indeed, the inferred rotation is so small that the implied virial mass is over 30 times smaller than M_* (Fig. B4, bottom panels). More to the point, the magnitude and orientation of the NIRSpec gradient inferred from Model 3 are still opposite to what seen in *The Cliff*, supporting the hypothesis that the observed gradient may be not only real, but actually even stronger than what is measured. Still, substantial uncertainties remain. Aside from the limits of our toy galaxy model, the instrument gradient may vary spatially, and in any case a single disc galaxy may be out of stationary equilibrium, as evidenced by residual tumbling motions even after subtracting the instrument model (e.g., Fig. B3c). Before drawing scientific conclusions from applying our proposed correction, more observations and exploring more complex models are required.

APPENDIX C: APERTURE LOSS CORRECTION FACTORS

In Fig. C1 we present the curve used to calculate ALC factors for each emission line in this study, based on a H α curve of growth analysis as discussed in Section 3.1. We derive this from the empirical ALC curve calculated from the PSF derived from STPSF⁶ for our aperture (Fig. 3), which is rescaled to match the empirically measured H α ALC factor of 1.88 ± 0.05 . For further detail regarding ALCs derived from the PSF, see Jones et al. (2026).

APPENDIX D: ANOTHER ARGUMENT AGAINST HIGH DENSITY

In this section, we provide a supplemental argument against high density as an explanation for [O III] λ 5007 weakness, as outlined for QSO1 in Maiolino et al. (2025a).

Firstly, we assume the density follows a radial distribution described by a power law,

$$n(r) = n_0 \left(\frac{r}{r_0} \right)^{-\alpha}, \quad (\text{D1})$$

where n_0 is the density at some reference radius r_0 . Here we take $r_0 = 40 \text{ pc}$, i.e. the effective radius of *The Cliff* as inferred by de Graaff et al. (2025b); this is smaller than the aperture described in Section 3.1, but a smaller r_0 is a more conservative assumption for obtaining a bound on the gas density. Furthermore, we assume that the clouds within this density profile have a filling factor f defined by

$$f = \frac{V_{\text{clouds}}}{V_{\text{total}}} = \frac{N_{\text{clouds}} R_c^3}{R_{\text{max}}^3}, \quad (\text{D2})$$

where V_{total} is the total volume enclosed by the aperture, assumed spherical with radius R_{max} . V_{clouds} is the volume occupied by ionised

⁶ <https://stpsf.readthedocs.io/en/latest/>; the successor to WEBBPSF (Perrin et al. 2015). STPSF accounts for many properties of the NIRSpec instrument (e.g. geometric distortion, detector charge transfer effects), and therefore may be used to derive physical models of PSF wavelength dependence.




Figure C1. Aperture loss correction factors as function of observed wavelength for the aperture shown in Fig. 3. This was derived using the full PSF model from STPSF, rather than a Gaussian approximation. Values have been rescaled to match our empirically measured ALC factor for H α found through a curve of growth analysis. The wavelengths of some key emission lines have been indicated with vertical dashed lines and labelled.

gas clouds of radius R_c , and N_{clouds} is the number of clouds. Assuming a medium comprising only hydrogen for simplicity, the mass of gas (M_g) contained within the aperture is therefore given by

$$M_g = \int_{R_{\text{min}}}^{R_{\text{max}}} 4\pi f m_p n(r) r^2 dr, \quad (\text{D3})$$

where m_p is the proton mass. We adopt $R_{\text{min}} = 5 \text{ pc}$, as smaller radii would be entirely within the black hole sphere of influence and would therefore result in broader line emission⁷. Hence, we plot the inferred gas mass as a function of density for different filling factors (Fig. D1), and determine the implied maximum cloud size (determined by assuming the extreme, very conservative case of $N_{\text{cloud}} = 1$) associated with each filling factor. In reality, not all of the gas would be contained within a single cloud, so the clouds would be much smaller.

We can compare the above estimate with the ionised gas mass, obtained from the luminosity of H β_N

$$L(\text{H}\beta_N) = \int_{R_{\text{min}}}^{R_{\text{max}}} 4\pi f \varepsilon(T) n(r)^2 r^2 dr, \quad (\text{D4})$$

where $\varepsilon(T)$ is the temperature-dependent emissivity of narrow H β . The total mass of ionised gas can then be inferred from the measured $L(\text{H}\beta_N)$, and is given by

$$\frac{M_{\text{ion}}}{M_\odot} \approx 10^9 \frac{L(\text{H}\beta_N)}{10^{43} \text{ erg s}^{-1}} \frac{100 \text{ cm}^{-3}}{n_0} \frac{3 - 2\alpha}{3 - \alpha} \frac{R_{\text{max}}^{3-\alpha} - R_{\text{min}}^{3-\alpha}}{R_{\text{max}}^{3-2\alpha} - R_{\text{min}}^{3-2\alpha}} r_0^{-\alpha}, \quad (\text{D5})$$

where we have assumed a typical temperature $T = 10^4 \text{ K}$.

Assuming $\alpha = 2$, in order for the gas to have a density greater

⁷ As in Maiolino et al. (2025a), we note that adopting a smaller inner radius would simply strengthen the results.




Figure D1. Constraints on gas density for narrow-line emitting gas. The orange lines indicate the gas mass M_g enclosed within the aperture (Fig. 3) as a function of gas density n_0 at $R = 40$ pc, for a power-law radial distribution (Equation D1) with indices $\alpha = 2$ (solid) and $\alpha = 0$ (dashed). Each line is labelled by its corresponding filling factor and implied upper limit on single-cloud size. The blue lines indicate the mass of ionised gas constrained by the $H\beta_N$ luminosity for the two power indices. The vertical dotted line indicates the density where collisional de-excitation of $[O\text{ III}]\lambda 5007$ begins to affect its flux; the vertical dot-dashed line indicates the density required to suppress $[O\text{ III}]\lambda 5007$ flux by a factor of 10. Explaining the weakness of $[O\text{ III}]\lambda 5007$ via collisional de-excitation would require extremely low filling factors and small cloud sizes. The horizontal pink line indicates the upper limit on dynamical mass derived in Section 4.4.

than 10^6 cm^{-3} , at which $[O\text{ III}]\lambda 5007$ starts to be collisionally suppressed by a factor of 1.5, the ionised clouds would have extremely small filling factors ($f \lesssim 10^{-8} N_{\text{clouds}}^{-1}$), and sizes < 0.09 pc. Such extreme filling factors are unrealistic, as no gas phase is known to have such extreme values, especially the ionized phase (not even for the BLR of AGN). Even an extreme scenario with $\alpha = 0$, i.e. uniform filling, would still yield $f \lesssim 10^{-7} N_{\text{clouds}}^{-1}$. Adopting the same CLOUDY models as Maiolino et al. (2025a), we infer that to suppress $[O\text{ III}]\lambda 5007$ by a factor of ten, the ISM would need to reach densities $> 10^7 \text{ cm}^{-3}$. This would necessitate even lower filling factors and yet smaller clouds. We also note that such high density would imply an extremely low mass of ionized gas, less than $10 M_\odot$.

APPENDIX E: THE AESOPICA HYDRODYNAMICAL SIMULATIONS

AESOPICA is a new suite of large-volume cosmological simulations (Koudmani et al., in prep), built upon the FABLE galaxy formation model (Henden et al. 2018); the FABLE sub-grid models are themselves largely based on the Illustris galaxy formation model (Vogelsberger et al. 2014). In FABLE, some aspects are unchanged from Illustris, namely the models for star formation (Springel & Hernquist 2003), radiative cooling (Katz et al. 1996; Wiersma et al. 2009a) and chemical enrichment (Wiersma et al. 2009b). However, the models of stellar feedback (Vogelsberger et al. 2013) and AGN feedback

(Sijacki et al. 2015) have been updated, incorporating thermal stellar feedback and AGN duty cycles as part of the simulations.

AESOPICA introduces targeted updates for modelling the growth of infant SMBHs in the early Universe. In particular, AESOPICA explores three key modifications to fiducial galaxy formation models: enabling efficient accretion in the low-mass regime (Koudmani et al. 2022), incorporating super-Eddington accretion, and examining a broad range of seed masses ($10^2 M_\odot$ to $10^5 M_\odot$) following seed evolution from early cosmic epochs ($z \sim 20$); while this is earlier than the standard threshold in FABLE, we note that seeding could occur even earlier.

APPENDIX F: SPECTRAL FITTING POSTERIORIS

In this section, we present the posteriors from the best-fit spectral models described in Section 3.2.1 and 3.2.2. Fig. F1 shows the posteriors from fitting the $H\beta$ - $[O\text{ III}]\lambda 5007$ complex (Section 3.2.1), while Fig. F2 and F3 show the posteriors from the Fiducial and Secondary $H\alpha$ model fits, respectively. We only include posteriors for the key fitted parameters to provide clarity about potential correlations, so for the sake of brevity fixed parameters (e.g. FWHM_N for $H\alpha$) or parameters not key to our analysis (e.g. continuum slopes) have been omitted from these figures.

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




Figure F1. Posterior distributions of the fitted parameters in the $H\beta$ - $[O\text{ III}]\lambda 5007$ model, including the best fit values and their associated errors. Contours enclose 68%, 95%, and 99% of the distributions.




Figure F2. Posterior distributions of the fitted parameters in the *Fiducial H α* model, including the best fit values and their associated errors. Contours enclose 68%, 95%, and 99% of the distributions.




Figure F3. Posterior distributions of the fitted parameters in the *Secondary* H α model, including the best fit values and their associated errors. Contours enclose 68%, 95%, and 99% of the distributions.