

The JWST EXCELS Survey: A spectroscopic investigation of the ionizing properties of star-forming galaxies at $1 < z < 8$

R. Begley¹★, R. J. McLure¹, F. Cullen¹, A. C. Carnall¹, T. M. Stanton¹, D. Scholte¹,
D. J. McLeod¹, J. S. Dunlop¹, K. Z. Arellano-Córdova¹, C. Bondestam¹, C. T. Donnan²,
M. L. Hamadouche³, H.-H. Leung¹, A. E. Shapley⁴, S. Stevenson¹

¹*Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK*

²*NSF's National Optical-Infrared Astronomy Research Laboratory, 950 N. Cherry Ave., Tucson, AZ 85719, USA*

³*Department of Astronomy, University of Massachusetts, Amherst, MA 01003, USA*

⁴*Department of Physics & Astronomy, University of California, 430 Portola Plaza, Los Angeles, CA 90095, USA*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Charting the Epoch of Reionization demands robust assessments of what drives the production of ionizing photons in high-redshift star-forming galaxies, and requires better predictive capabilities from current observations. Using a sample of $N = 159$ star-forming galaxies (SFGs) at $1 \lesssim z \lesssim 8$, observed with ultra-deep medium-resolution spectroscopy from the JWST/NIRSpec EXCELS survey, we perform a statistical analysis of their ionizing photon production efficiencies (ξ_{ion}). We consider ξ_{ion} , accurately measured with spectroscopic Balmer line measurements, in relation to a number of key galaxy observables including: nebular emission line equivalent widths ($W_{\lambda}(\text{H}\alpha)$ and $W_{\lambda}([\text{O III}])$), UV luminosity (M_{UV}) and spectral slope (β_{UV}), as well as dust attenuation ($E(B - V)_{\text{neb}}$) and redshift. Implementing a Bayesian linear regression methodology, we fit ξ_{ion} against the principal observables while fully marginalising over all measurement uncertainties, mitigating against the impact of outliers and determining the intrinsic scatter. Significant relations between ξ_{ion} and $W_{\lambda}(\text{H}\alpha)$, $W_{\lambda}([\text{O III}])$ and β_{UV} are recovered. Moreover, the weak trends with M_{UV} and redshift can be fully explained by the remaining property dependencies. Expanding our analysis to multivariate regression, we determine that $W_{\lambda}(\text{H}\alpha)$ or $W_{\lambda}([\text{O III}])$, along with β_{UV} and $E(B - V)_{\text{neb}}$, are the most important observables for accurately predicting ξ_{ion} . The latter identifies the most common outliers as SFGs with relatively high $E(B - V)_{\text{neb}} \gtrsim 0.5$, possibly indicative of intense obscured star-formation or strong differential attenuation. Combining these properties enable $\xi_{\text{ion},0}$ to be inferred with an accuracy of ≈ 0.15 dex, with a population intrinsic scatter of $\sigma_{\text{int}} \approx 0.035$ dex.

Key words: galaxies: high-redshift - galaxies: evolution - cosmology: dark ages, reionization, first stars

1 INTRODUCTION

The rate of ionizing photon production (Lyman continuum; LyC, $h\nu \geq 13.6$ eV) and the subsequent escape fraction are critical quantities in determining the total number of ionizing photons available (Robertson et al. 2015, 2023) to drive and then maintain the ionized state of the Universe (Goto et al. 2021; Bosman et al. 2021). Sufficient evidence has now been collated to confidently suggest that it was star-forming galaxies (SFGs; Chary et al. 2016; Robertson et al. 2023), rather than quasars (Aird et al. 2015; Kulkarni et al. 2019; Matsuoka et al. 2023; Trebitsch et al. 2023), that forced this transition in the ionization state of the Universe, during the so-called Epoch of Reionization (EOR).

The population demographics of galaxies at redshifts $5 \leq z \leq 10$, during which the bulk of reionization takes place (Mason et al. 2018; Planck Collaboration et al. 2020; Robertson et al. 2023; Simmonds et al. 2024a; Roberts-Borsani et al. 2024; Tang et al. 2025),

is now increasingly well constrained. Determinations of the UV luminosity function at these epochs provide these key constraints (e.g., see McLure et al. 2013; Bouwens et al. 2015; Donnan et al. 2022; McLeod et al. 2024; Adams et al. 2025; Whitler et al. 2025), revealing large populations of low-luminosity sources ($M_{\text{UV}} \gtrsim -18$) consequential of a relatively steep faint-end slope (Finkelstein et al. 2015; Atek et al. 2023) as well as a surprising abundance of M_{UV} -bright sources (i.e., $M_{\text{UV}} \lesssim -20$ Bowler et al. 2015; Whitler et al. 2025).

Connecting the SFG demographics to models tracking the timeline and topology of reionization, then becomes a challenge of piecing together the relative ionizing photon budget contributions of different sub-populations, which in turn begs the question - which reionization-era SFGs are the most conducive to strong ionizing outputs?

The escape of ionizing photons (parameterised as the LyC escape fraction; $f_{\text{esc}}^{\text{LyC}}$) can only be directly constrained from the local Universe (e.g., see Izotov et al. 2018; Flury et al. 2022; Maulick et al. 2025) out to intermediate redshifts ($z \lesssim 4$) as a result of the increasingly opaque IGM at higher redshifts (Inoue & Iwata 2008).

★ E-mail:rbeg@roe.ac.uk

Photometric and spectroscopic studies at $z \simeq 3 - 4$, provide estimates of $f_{\text{esc}}^{\text{LyC}} \simeq 3 - 10$ per cent (Steidel et al. 2018; Pahl et al. 2021, 2023; Begley et al. 2022), in good agreement with predictions from simulations (e.g., see Ma et al. 2020; Trebitsch et al. 2023). Consequentially for current models of the reionization timeline utilising UV luminosity functions, these works (see also; Marchi et al. 2018; Fletcher et al. 2019) also find that higher $f_{\text{esc}}^{\text{LyC}}$ is more likely in M_{UV} -faint sources (however, see Naidu et al. 2020; Jung et al. 2023, for opposing arguments).

Significant effort has also been invested in constructing accurate and robust indirect methods of measuring LyC leakage, primarily from low-to-intermediate redshift galaxy observations (e.g., see Jaskot et al. 2019; Flury et al. 2022; Mascia et al. 2023; Saldana-Lopez et al. 2023). In principle, such indirect probes then allow key insights into the LyC escape and ionizing capabilities of galaxies beyond $z \sim 4$ and into the EOR (Chisholm et al. 2022; Begley et al. 2023; Mascia et al. 2024). However, it is worth highlighting there are still significant open questions regarding Lyman continuum escape. Firstly, the scatter in many of the derived relations between $f_{\text{esc}}^{\text{LyC}}$ and other observed properties is significant (e.g., see Flury et al. 2022; Jaskot et al. 2024; Jaskot 2025). Moreover, there are still significant gaps in the parameter spaces (e.g., M_{UV}) accessible from current simulations and observations, in addition to more complex temporal and small-scale geometrical factors thought to significantly impact LyC escape (Kimm et al. 2022; Maji et al. 2022; Katz et al. 2023; Choustikov et al. 2024) that are extremely challenging to observationally explore.

Nonetheless, taken together, the numerous population of faint galaxies, along with their apparently favourable $f_{\text{esc}}^{\text{LyC}}$ conditions, would suggest that they are likely the dominant contributors to the ionizing photon budget. However, such conclusions can only be maintained if these faint populations are also significant *producers* of ionizing photons. The ionizing photon production efficiency (ξ_{ion}), defined as $\xi_{\text{ion}} = N(\text{H}^0)/L_{\text{UV}}$ where $N(\text{H}^0)$ is the rate of ionizing photon production and L_{UV} is the intrinsic monochromatic UV luminosity ($\lambda_{\text{rest}} = 1500 \text{ \AA}$), is the most common parameterisation of ionizing photon production, and is typically required to be $\log_{10}(\xi_{\text{ion},0} / \text{erg s}^{-1} \text{ Hz}) \gtrsim 25.2$ on average to satisfy reionization modelling requirements (Robertson et al. 2015).

The first years of imaging and spectroscopy data from JWST enabled significant progress to be made characterising the ionizing photon production efficiencies of galaxies from cosmic noon ($z \sim 2$) through to the EOR (Simmonds et al. 2024a; Begley et al. 2025; Pahl et al. 2025b; Llerena et al. 2025; Pahl et al. 2025a; Papovich et al. 2025). Importantly, this collation of the latest literature suggests that the SFG population has sufficiently high ξ_{ion} on average to drive reionization to a timely completion, but with significant population scatter (e.g., $\langle \log_{10}(\xi_{\text{ion},0} / \text{erg s}^{-1} \text{ Hz}) \rangle \simeq 25.3$ and $\sigma \simeq 0.3$ dex, see Begley et al. 2025, where $\xi_{\text{ion},0}$ is computed implicitly assuming $f_{\text{esc}}^{\text{LyC}} = 0$, as common in the literature). This helps avoid a photon budget crisis (Muñoz et al. 2024) in which there are too many ionizing photons based on pre-JWST expectations (Finkelstein et al. 2019; Stefanon et al. 2022; Chen et al. 2024).

There remain however, critical impasses to fully understanding what the quintessential ‘reionization-driving galaxy’ is. Recent work by Endsley et al. (2022); Begley et al. (2025) and Pahl et al. (2025b) (see also; Simmonds et al. 2024a) find evidence for a $\xi_{\text{ion}} - M_{\text{UV}}$ anti-correlation during the EOR, in which M_{UV} -bright sources display larger ξ_{ion} , consistent with being observed in a recent burst phase (in addition to increasingly bursty star-formation towards higher redshifts). In contrast, other works suggest the fainter population has

elevated ξ_{ion} (Prieto-Lyon et al. 2023; Llerena et al. 2023; Simmonds et al. 2024b; Papovich et al. 2025). Differences in these inferred relations could be due to a multitude of factors, likely including sample bias effects from spectroscopically versus photometrically selected SFG populations and SED-modelling versus directly derived ionizing photon production rates.

These conflicting trends likely arise from a combinations of both the diversity in the true underlying physical conditions (e.g., stellar populations, metallicities, burstiness, initial mass function among others Conroy 2013; Eldridge et al. 2017; Robertson 2021; Roberts-Borsani et al. 2024; Hayes et al. 2025) as well as a range of possible observational and measurement effects (e.g., selection biases, spectroscopic versus photometric samples, treatment of dust correction, to name a few).

It is clear that a more detailed understanding of what dictates the production (and escape) of ionisation photons across the galaxy population, and how we can better calibrate this with the observed and derived physical properties of the SFG population is required. To this end, in this work we explore the physical drivers of ξ_{ion} using medium resolution spectroscopy of $N = 159$ SFGs observed in the JWST/NIRSpec Early eXtragalactic Continuum and Emission Line Science (EXCELS; GO 3543; PIs: Carnall, Cullen; Carnall et al. (2024)) program, spanning a wide swathe of cosmic time ($z \simeq 1 - 8$). By statistically investigating the factors associated with ξ_{ion} , we aim to decipher what facilitates increased ionizing photon production in star-forming galaxies.

The structure of the paper is as follows. In Section 2 we describe the EXCELS dataset and galaxy measurements, including the emission line fluxes, UV continuum slopes, and ionizing photon production efficiencies. In Section 3, we introduce our Bayesian Linear regression methodology, while Section 4 presents the results of fitting $\xi_{\text{ion},0}$ against the key observed properties (i.e., see Section 4.1). Section 4.2 expands this to a multivariate analysis, including discussions on the best predictors for $\xi_{\text{ion},0}$ and the impact of outliers. Lastly, in Section 5, we summarise the main conclusions of our analysis.

Throughout the paper we adopt the following cosmological parameters: $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-3}$, $\Omega_{\text{m}} = 0.3$, $\Omega_{\Lambda} = 0.7$ and all magnitudes are quoted in the AB system (Oke & Gunn 1983).

2 DATA AND SAMPLE PROPERTIES

2.1 The EXCELS survey

The spectroscopic sample that forms the basis of this work was observed in the JWST NIRSpec EXCELS survey. The primary goal of the EXCELS survey was to target massive $3 \leq z \leq 5$ quiescent galaxies with high signal-to-noise, medium resolution ($R = \lambda/\Delta\lambda \simeq 1000$) spectroscopy, in addition to hundreds of star-forming galaxies across $z \simeq 2 - 8$ observed in the remaining MSA slits. For complete details on the selection of EXCELS survey targets, observing configurations and data processing, readers are directed to Carnall et al. (2024) (see also Scholte et al. 2025; Stanton et al. 2025). Below we provide a brief overview.

2.1.1 Observing configuration and data reduction

The EXCELS survey consists of four NIRSpec/MSA pointings within the UDS (Ultra-Deep Survey) footprint of the Cycle 1 JWST/NIRCam survey PRIMER (Public Release IMaging for Extra-galactic Research; GO 1837; PI: Dunlop). Each pointing is observed with the G140M/F100LP, G235M/F170LP and G395M/F290LP




Figure 1. Example spectroscopy available for a galaxy in the EXCELS sample (ID: EXCELS-59720). The G235M/F170LP and G395M/F290LP spectra from EXCELS for this SFG are shown (*left* and *right*, respectively), highlighting the plethora of rest-frame optical emission lines (including but not limited to; [O II], H α , H β , H γ , H δ ; see insets) potentially observable with the deep NIRSpec spectroscopy.

gratings (with unique masks in each) providing continuous wavelength coverage over $\lambda_{\text{obs}} \approx 1 - 5 \mu\text{m}$, and employing a 3-shutter slitlet configuration and a 3-point dither pattern. The sources are observed for ≈ 4 hours in the G140M and G395M gratings, and ≈ 5.5 hours in the G235M grating, reaching 3σ line flux limits of $\approx 5 \times 10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2}$ (see [Carnall et al. 2024](#), for full details).

Unique NIRSpec/MSA masks were designed for each grating configuration and pointing, optimised to maximise the number of targets observed with wavelength coverage spanning key spectral diagnostics, specifically at $\lambda_{\text{rest}} \approx 3500 - 7000$ (i.e., the rest-frame near-UV to optical regime). The EXCELS target selection was primarily based on the VANDELS survey catalogues ([McLure et al. 2018b](#); [Garilli et al. 2021](#)), in addition to galaxies selected from PRIMER/UDS photometric catalogues (see [Carnall et al. 2024](#), and additionally [McLeod et al. 2024](#); [Begley et al. 2025](#); [McLeod et al.](#) in preparation for further details).

The level 1 NIRSpec data products were reduced using the JWST reduction pipeline, including the advanced snowball rejection (v1.15.1 with the CRDS_CTX = jwst_1258.pmap JWST Calibration Reference Data System files). Following this, the level 2 and level 3 pipeline procedures are implemented in their default configuration. A custom 1D optimal extraction of the 2D spectra is implemented, whereby the extraction centroid is set as the flux-weighted mean position ([Horne 1986](#)) of the source within the NIRSpec/MSA slitlet.

Lastly, a two-stage flux calibration procedure is implemented to ensure accurate absolute flux calibrations for the EXCELS spectra accounting for wavelength-dependent slit losses arising from the broadening PSF at longer wavelengths, as well as small source centroid offsets between gratings. In the first stage, a relative flux calibration between the gratings is performed by scaling the G140M and/or G395M spectra to match the G235M spectra using the median flux of the overlapping spectra regions. Subsequently, an absolute spectrophotometric calibration is applied by scaling the spectra to align with the available broadband photometry described in Section 2.5 (see also [Rogers et al. 2024](#), following an analogous method used in the CECILIA survey). Each EXCELS spectrum is integrated through the available set of filters overlapping in wavelength, with a linear interpolation applied to derive a wavelength-dependent absolute flux calibration. The calibration is verified using a subsample of EXCELS sources with high signal-to-noise continuum detections and/or emis-

sion lines in overlapping wavelength ranges between the gratings (e.g., see Fig. 1), indicating that the flux calibrations are accurate to $\lesssim 5 - 8$ per cent.

2.2 Supplementary photometric catalogues

The EXCELS spectroscopic sample in this work benefits from having deep multi-wavelength imaging (as a result of its original selection). Specifically, all EXCELS sources are covered by the PRIMER/UDS footprint and therefore have imaging in eight JWST/NIRCam filters (F090W, F115W, F150W, F200W, F277W, F356W, F410M, F444W), in addition to deep optical imaging from HST/ACS in three filters (F435W, F606W, F814W) from CANDELS (Cosmic Assembly Near-IR Deep Extragalactic Legacy Survey; [Grogin et al. 2011](#); [Koekemoer et al. 2011](#)). Sources are detected in the PRIMER NIRCam/F356W image using SOURCE EXTRACTOR ([Bertin & Arnouts 1996](#)), with photometry measured using $0.5 - \text{arcsec}$ diameter apertures (enclosing ≈ 85 per cent of the total flux assuming a point source) in dual-image mode on the full set of available images after they have been PSF homogenised to the F444W image (see [McLeod et al. 2024](#)).

The associated photometric uncertainties are deduced from local depth maps, where each object flux error in a given band is calculated as the scaled median absolute deviation ($\sigma_{\text{local}} = 1.483 \times \sigma_{\text{MAD}}$) of the fluxes measured in the nearest 150 – 200 empty-sky apertures (e.g., see [McLeod et al. 2021](#)). A more comprehensive description of the PRIMER photometric catalogues and their creation can be found in [Begley et al. \(2025\)](#) (see also [McLeod et al. 2024](#), [McLeod et al. 2025](#), in preparation, for further methodology details).

2.3 EXCELS emission line measurements

To measure robust line fluxes and equivalent widths for each source, a two-step fitting procedure is implemented - first estimating the continuum flux, and then fitting the continuum-subtracted (line only) spectrum with a linear combination of Gaussian line profiles. The methodology used to obtain the emission line and continuum models for the calibrated 1D EXCELS spectra, as shown in Fig. 1, are described comprehensively in [Scholte et al. \(2025\)](#) (see also [Stanton et al. 2025](#)), with a summary of the key elements included here.

The continuum flux is estimated from the running mean of the flux values within a $\Delta_{\text{rest}} = 350 \text{ \AA}$ top hat function, after clipping values outside the 16th – 84th percentile range to mitigate the impact of strong spectral features or noise spikes on the continuum measurement (Stoughton et al. 2002). Given the underestimated uncertainties in the pixel flux measurements produced from the JWST data reduction pipeline, an empirical correction based on the residuals of the continuum subtracted spectra was applied. This is implemented as a multiplicative factor to the pipeline flux uncertainties, given as $c_F = \frac{1}{2}(R_{84} - R_{16})/\bar{\sigma}_F$ where R_X is the X^{th} percentile of the residuals spectrum (R) and $\bar{\sigma}_F$ is the median of the pipeline flux uncertainty. Typical values of the empirical correction factor are $c_F \approx 1.5 - 2.0$ (with median $\bar{c}_F \approx 1.6$), in agreement with similar independent analyses (Maseda et al. 2023; Carnall et al. 2024).

The (continuum-subtracted) emission line spectrum is fit using a linear combination of Gaussian line profiles, where each individual line amplitude is a free parameter. The total width of the line profile (σ_{tot}) is given as the convolution of the intrinsic line width (σ_{intr}) and the instrumental broadening ($\sigma_g(\lambda)$) of the NIRSpec grating used: $\sigma_{\text{tot}}^2 = \sigma_{\text{intr}}^2 + \sigma_g^2(\lambda)$. A common line velocity (relative to the input z_{spec}) and intrinsic line width is used for each galaxy spectral fit resulting in $2 + n_{\text{line}}$ free parameters (where n_{line} is the number of emission line profiles included in the model). The fitting is performed using a weighted least-squares minimisation via SciPy (Branch et al. 1999; Virtanen et al. 2020a). The line fluxes are calculated from the line profile weighted integral over the individual wavelength pixels in each fitted spectrum, with the uncertainties propagated from the same weighted pixel fluxes (e.g., see Scholte et al. 2025).

To account for any additional uncertainties arising from the absolute flux calibration employed, an additional ≈ 8 per cent error is included in quadrature when absolute flux values, or line flux ratios between lines spanning multiple gratings, are used (Scholte et al. 2025, see also Section 2.1.1).

2.4 Sample Selection

Observations taken in the EXCELS survey enabled spectroscopic redshifts to be derived for $N = 349$ sources, of which $N = 341$ were assigned high quality redshift flags of 3, 4 or 9 according to the framework outlined in Le Fèvre et al. (2005) (see also; Pentericci et al. 2018). The most secure quality flag, 4 indicating $\gtrsim 99$ per cent confidence in the derived redshift, was assigned to $N = 309$ sources representing ~ 90 per cent of the full sample. Sources without a high quality redshift flag are not considered in this study.

Accurately inferring the ionizing photon production efficiency spectroscopically from nebular emission lines (see Section 2.5) requires robust determination of the *intrinsic* line fluxes corrected for the effects of dust attenuation. To calculate the nebular dust attenuation we use the Balmer decrement (see Section 2.5.1 and Fig. 2), and thus select galaxies with at least one Balmer line pair from $H\alpha$, $H\beta$, $H\gamma$ and $H\delta$. We additionally impose that the emission lines in a given pair are significantly detected, requiring $S/N \geq 3$ for each. From this selection, we obtain a final sample of $N = 159$ star-forming galaxies (noting that none of the targets classified as quiescent according to the priority classes defined in Carnall et al. 2024 are selected in the defined criteria). The sample contains galaxies across a wide redshift range ($1 \lesssim z_{\text{spec}} \lesssim 8$), with a median redshift $\langle z_{\text{spec}} \rangle = 3.24$ and ≈ 73 per cent ($N = 116$) of the sample between $z_{\text{spec}} \approx 2 - 6$. The spectroscopic redshift distribution of the EXCELS sample is shown in Fig. 3 alongside physical properties calculated in Section 2.5 below. We note here that given the target selection from the VANDELS and PRIMER parent photometric catalogues (McLure et al. 2018b;

Carnall et al. 2024) used in EXCELS, the sample assembled here is consistent with being drawn from the star-forming main sequence (i.e., as opposed to a narrow-band selected catalogue preferentially targeting emission line galaxies).

2.5 Physical properties

2.5.1 Nebular dust attenuation

To correct the nebular emission line fluxes for the effects of dust attenuation, we compare the observed Balmer decrements in each source to the expected intrinsic line ratios. In each case, we use all the available Balmer line pairs from $H_X = \{H\alpha, H\beta, H\gamma, H\delta\}$ where the lines in each pair are robustly detected ($S/N \geq 3$). The theoretical intrinsic line ratios are calculated with PYNEB (v1.1.27 Luridiana et al. 2015; Morisset et al. 2020), assuming case B recombination (Storey & Hummer 1995) with electron temperatures and densities of $T_e = 12,000 \text{ K}$ and $n_e = 300 \text{ cm}^{-3}$. These temperature and density assumptions are based on the average values found in EXCELS galaxies (e.g., see Scholte et al. 2025; Stanton et al. 2025; Arellano-Córdova et al. 2024), and produce only marginally different intrinsic line ratios compared to the default case B assumptions commonly used (i.e., $T_e = 10,000 \text{ K}$ and $n_e = 100 \text{ cm}^{-3}$; with $\Delta(H\alpha/H\beta)_{\text{intr}} \approx 0.03$ and $\Delta(H\gamma/H\beta)_{\text{intr}} \approx 0.002$). In Fig. 2, we show distributions of the observed Balmer decrements for $H\alpha/H\beta$, $H\gamma/H\beta$ and $H\delta/H\beta$ for our EXCELS sample, with the theoretical intrinsic limits under our case B recombination assumptions also marked. The vast majority of the sample (≈ 92 per cent) have observed Balmer decrements consistent with Case B recombination assumptions. Only 13 sources show discrepancies at the $> 2\sigma$ level - for these sources we assume $E(B-V)_{\text{neb}} = 0.0$ (as commonly adopted in the literature; e.g., Pahl et al. 2025b), with a more detailed treatment of different recombination cases left to future work (but see Scarlata et al. 2024; Yanagisawa et al. 2024; Arellano-Córdova et al. 2024; McClymont et al. 2025, for recent relevant discussions).

The reddening from each galaxy, $E(B-V)_{\text{neb}}$, is calculated as the variance weighted mean of values from each available Balmer line pair $E(B-V)_{\text{neb}, i}$, where:

$$E(B-V)_{\text{neb}, i} = 2.5 \frac{\log_{10} \left(\frac{(f_{H_X}/f_{H_Y})_{\text{obs}}}{(f_{H_X}/f_{H_Y})_{\text{intr}}} \right)}{k_\lambda(\lambda_{H_Y}) - k_\lambda(\lambda_{H_X})} \quad (1)$$

where $(f_{H_X}/f_{H_Y})_{\text{obs}}$ and $(f_{H_X}/f_{H_Y})_{\text{intr}}$ are the observed and intrinsic Balmer flux ratio, respectively, and k_λ parameterises the dust attenuation law. We assume the nebular dust attenuation follows a Cardelli et al. (1989) attenuation curve, having been shown to be a good approximation for high-redshift star-forming galaxies (Reddy et al. 2020). In reality, the dust attenuation law likely varies across the galaxy population (e.g., Salim & Narayanan 2020; Sanders et al. 2024; Reddy et al. 2025), however, inferring the dust attenuation law on an individual galaxy-by-galaxy basis would require larger dynamic range in wavelength for each source than is currently available. Nevertheless, detailed analysis of the individual dust attenuation laws for a subset of EXCELS galaxies provides supporting evidence in favour of the Cardelli et al. (1989) – like prescription on average (Stanton et al. in prep.).

We take a Monte Carlo approach to estimate the uncertainties on $E(B-V)_{\text{neb}}$ by repeating the calculation $\sim 10^3$ times and perturbing the line flux measurements by their errors in each iteration, to construct a $E(B-V)_{\text{neb}}$ distribution for each EXCELS galaxy. The fiducial $E(B-V)_{\text{neb}}$ values and errors are given as the median and




Figure 2. The observed Balmer decrement (see Section 2.5.1 for details) distributions for the EXCELS galaxy sample ($N = 159$) subsets with $H\alpha/H\beta$ ($N = 105$; **top panel**), $H\gamma/H\beta$ ($N = 121$; **centre panel**) and $H\delta/H\beta$ ($N = 112$; **bottom panel**). The inset panels show the distributions of measurement errors (σ) for each Balmer decrement pair $(H_X/H_Y)_{\text{obs}}$, with typical errors of $\approx +0.24$ for $H\alpha/H\beta$, and of $\approx +0.06$ for both $H\gamma/H\beta$ and $H\delta/H\beta$. The dashed lines and preceding grey shaded regions illustrate the intrinsic Balmer line ratios under the Case B recombination assumptions employed here.

$\pm 1\sigma$ (16th – 84th percentiles) of the distributions¹, noting however, that a Monte Carlo approach is taken for all subsequent calculations based on the intrinsic nebular line fluxes using the measured $E(B - V)_{\text{neb}}$ probability distribution functions. Lastly, we note that the probability distribution functions are truncated with a floor of $E(B - V)_{\text{neb}} \geq 0.0$, setting unphysical reddening values to zero.

2.5.2 UV continuum slope

The UV continuum slope, β_{UV} , where $f_\lambda \propto \lambda^{\beta_{\text{UV}}}$, is measured following the method described in Begley et al. (2025) (see also Cullen

¹ Computing the median $E(B - V)_{\text{neb}}$ after removing MCMC draws that are inconsistent with case B recombination would increase $E(B - V)_{\text{neb}}$ in the low attenuation tail ($E(B - V)_{\text{neb}} < 0.1$) by ≈ 0.05 . This corresponds to a ≈ 0.02 dex shift in $\xi_{\text{ion},0}$, which is well within our individual galaxy uncertainties.

et al. 2024). Briefly here, only photometry covering rest-frame wavelengths $\lambda_{\text{rest}} \leq 3000 \text{ \AA}$ is considered (specifically, selecting filters with $\lambda_{95\text{th}} \leq 3000 \text{ \AA}$, where $\lambda_{95\text{th}}$ is the 95th percentile wavelength of the filters cumulative transmission). We then adopt a power-law functional form to model the flux redward of $\text{Ly}\alpha$ ($f_\lambda = a \cdot \lambda^{\beta_{\text{UV}}}$; $\lambda > 1216 \text{ \AA}$), with the addition of the Inoue et al. (2014) IGM prescription at shorter wavelengths. We note that the Inoue et al. (2014) IGM attenuation prescription is that of the *average* transmission function at a given redshift, and so to account for the increasing IGM stochasticity with redshift we include a scaling parameter ϕ , which we marginalise over in our fitting. The final fitted model is then $f_\lambda = a \cdot T_\lambda(z, \phi) \cdot \lambda^{\beta_{\text{UV}}}$, where $T_\lambda = 1$ at $\lambda > 1216 \text{ \AA}$, and $T_\lambda = \phi \cdot g_{\text{Inoue14}}(\lambda, z)$, otherwise.

The best-fitting model parameters are estimated from their posterior distributions, sampled using the MCMC ensemble sampler EMCEE (Foreman-Mackey et al. 2013). During the sampling, we fix the model redshift to the EXCELS spectroscopic redshift ($z = z_{\text{spec}}$), and adopt uniform priors across the remaining parameters with the UV continuum slope allowed to vary in the range $-10 \leq \beta_{\text{UV}} \leq 10$. We note that the ϕ parameter, which captures the effect of the stochasticity of the IGM transmission, is treated as a dummy parameter and does not appreciably affect our measurements.

The β_{UV} measurements are based on the available photometric catalogues (see Section 2.2), which provide robust HST+JWST photometry over the observed wavelength range $0.36 \mu\text{m} \lesssim \lambda_{\text{obs}} \lesssim 5.0 \mu\text{m}$. As such, galaxies in our sample at $z \gtrsim 2.95$ will have robust β_{UV} constraints owing to sufficient photometric coverage redward and blueward of the Lyman break (typically with 3 – 4 bands redward of $\lambda_{\text{rest}} = 1216 \text{ \AA}$, and up to 7 bands when also considering those blueward.). On the other hand, lower redshift sources in our sample lack the photometry fully redward of $\lambda_{\text{Ly}\alpha} = 1216 \text{ \AA}$ and have more uncertain UV continuum slopes as a consequence.

Lastly, we measure the absolute rest-frame UV magnitude ($M_{\text{UV}} \equiv M_{1500}$) by first sampling 1000 realisations from the posterior UV continuum models of each source (described above). Each model is then integrated through a $\Delta\lambda = 100 \text{ \AA}$ top-hat filter centered on $\lambda_{\text{rest}} = 1500 \text{ \AA}$ (Begley et al. 2023) to produce a distribution of M_{UV} for each galaxy. The β_{UV} and M_{UV} for the EXCELS galaxies are then given as the median of their respective distributions, with the $\pm 1\sigma$ errors denoted as the 16th – 84th percentiles. The median absolute UV magnitude for our final sample of $N = 159$ EXCELS galaxies is $\langle M_{\text{UV}} \rangle = -19.5$, with a 16th – 84th percentile range of $M_{\text{UV}} \in [-20.3, -18.4]$ and reaching a faint (bright) extreme of $M_{\text{UV}} \approx -16.5$ (-21.2). Similarly, across our EXCELS sample we find a median UV continuum slope of $\langle \beta_{\text{UV}} \rangle = -1.50$, and a 16th – 84th percentile range of $\beta_{\text{UV}} \in [-2.21, -0.88]$. The EXCELS sample spans a large dynamic range in UV continuum slopes, with sources as blue (red) as $\beta_{\text{UV}} \approx -3.0$ ($+1.0$).

2.5.3 Ionizing photon production efficiency

In this analysis, the ionizing photon production efficiency ($\xi_{\text{ion},0} / \text{erg} \cdot \text{Hz}^{-1}$) complies with the common definition of: $\xi_{\text{ion},0} = N(\text{H}^0) / L_{\text{UV}}^{\text{intr}}$, where $N(\text{H}^0) / \text{s}^{-1}$ is the production rate of hydrogen-ionizing photons ($h\nu \geq 13.6 \text{ eV}$) and $L_{\text{UV}}^{\text{intr}} / \text{erg} \cdot \text{s}^{-1} \cdot \text{Hz}^{-1}$ is the UV continuum luminosity at $\lambda_{\text{obs}} = 1500 \text{ \AA}$ after correcting for the effects of reddening from attenuating dust grains in the diffuse ISM.²

Firstly, we calculate the intrinsic $H\alpha$ luminosity ($L_{\text{H}\alpha}^{\text{intr}}$) by dust-correcting the observed line fluxes assuming a Cardelli et al. (1989)

² Note, this is equivalent to the definition of $\xi_{\text{ion}}^{\text{HII}}$ used in Tang et al. (2019).




Figure 3. A corner plot of the six key galaxy observables $\mathbf{X}$ (where $X_i = \{W_\lambda(\text{H}\alpha), W_\lambda([\text{O III}]), \beta_{\text{UV}}, M_{1500}, E(B - V), z_{\text{spec}}\}$, shown in order as columns left to right, and as rows top to bottom), from which we derive relationships with $\xi_{\text{ion},0}$ in Section 3. Along the diagonal axes, we show the 1D marginal histograms for each covariate, highlighting the large dynamic range of physical properties spanned by our spectroscopic sample. Markers throughout are coloured by the inferred ionizing photon production efficiency ($\xi_{\text{ion},0}$; here specifically derived from $\text{H}\alpha$, where available), with the histogram bars along the diagonal axes indicating the median $\xi_{\text{ion},0}$ in the given bin. In each panel, the star symbol represents the property medians and 16th – 84th percentile range. Additionally, we present the Spearman rank correlation coefficient (ρ_s) and associated p -value between parameter pairs, alongside the 95 per cent confidence intervals on ρ_s , computed with PINGOUIN (e.g., see Vallat 2018, for details). As shown in Section 3, the apparent strength of the predictive power for a single parameter X_i on $\xi_{\text{ion},0}$ can significantly weaken as a result of the collinearity between the observables in the multivariate case (see also Table 1).

dust curve and adopting the $E(B - V)_{\text{neb}}$ values calculated from Balmer decrements as outlined in Section 2.5.1. The ionizing photon production rate can then be deduced from the intrinsic $\text{H}\alpha$ luminosity under the assumption of Case B recombination (Osterbrock & Ferland 2006) using: $N(\text{H}^0) = 7.35 \times 10^{-13} \cdot L_{\text{H}\alpha}^{\text{intr}}$ (noting that we have implicitly dropped a factor of $1/(1 - f_{\text{esc}}^{\text{LyC}})$ as our definition of $\xi_{\text{ion},0}$ assumes $f_{\text{esc}}^{\text{LyC}} = 0$).

To compute the intrinsic UV luminosity ($L_{\text{UV}}^{\text{intr}}$) we take advantage of the robust UV continuum slope measurements β_{UV} (calculated in Section 2.5.2). As detailed in Begley et al. (2022) (see also McLure et al. 2018a), the UV dust attenuation can be derived from the observed UV continuum slope as: $A_{1500} = \kappa \cdot (\beta_{\text{UV}}^{\text{obs}} - \beta_{\text{UV}}^{\text{intr}})$, where $\beta_{\text{UV}}^{\text{intr}}$ is intrinsic UV spectral slope and $\kappa = dA_{1500} / d\beta$ as specified by the underlying assumed dust attenuation law. Here, we adopt a fiducial value of $\beta_{\text{UV}}^{\text{intr}} = -2.4$ which was found to be representative of SFGs




Figure 4. The distribution of ionizing photon production efficiencies ($\xi_{\text{ion},0}$) for the sample of $N = 159$ EXCELS galaxies (light grey), calculated from the intrinsically brightest available Balmer line from $\text{H}\alpha$, $\text{H}\beta$, $\text{H}\gamma$ (see Section 2.5.3 for details). From the total sample, $N = 112$ are computed from $\text{H}\alpha$, with a remaining $N = 34$ from $\text{H}\beta$ and $N = 13$ from $\text{H}\gamma$ (with each plotted individually in purple, blue and green, respectively). The insets show the associated full $\xi_{\text{ion},0}$ distributions from each Balmer line (in galaxies where they are measurable), highlighting the excellent consistency across each. The coloured thick dotted lines indicate the distribution medians, and the black dashed vertical lines mark the canonical $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) \approx 25.2$ value often cited as the benchmark “typical” limit for galaxies to have driven reionization completely (Robertson et al. 2013)

at $z \approx 3 - 4$ in Begley et al. (2022)³. Additionally, we assume a dust law intermediate in greyness between a relatively flat Calzetti et al. (2000)–like dust attenuation law and a steeper SMC dust extinction curve (Gordon et al. 2003), using the Salim et al. (2018) dust curve parameterisation with $\delta = -0.25$.

From the full sample of $N = 159$ galaxies studied in this work, we compute $\text{H}\alpha$ –based $\xi_{\text{ion},0}$ for $N = 112$. For the remaining $N = 47$ sources without $\text{H}\alpha$ measurements (mainly due to EXCELS observing different sources with different combinations of NIRSpec gratings and a small number of sources where $\text{H}\alpha$ is redshifted beyond the observed wavelength coverage), we calculate $\xi_{\text{ion},0}$ from the next available higher-order Balmer line ($N = 34$ from $\text{H}\beta$ and $N = 13$ from $\text{H}\gamma$). The full sample $\xi_{\text{ion},0}$ distribution, alongside the distributions calculated from $\text{H}\alpha$, $\text{H}\beta$, and $\text{H}\gamma$ individually, are shown in Fig. 4. The median ionizing photon production efficiency for the full EXCELS sample is $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = 25.22 \pm 0.01$, with a population scatter (quoted as the σ_{MAD}) of $\Delta\xi_{\text{ion},0} \approx 0.42$ dex (and 16th – 84th percentile range of $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) \in [24.82, 25.74]$). We also highlight that the medians of the $\xi_{\text{ion},0}$ distributions calculated from each of the Balmer series lines individually (shown in the inset panels of Fig. 4) agree within $\Delta\xi_{\text{ion},0} \lesssim 0.1$ dex. Comparisons between the individual $\xi_{\text{ion},0}$ distributions reveal a marginal skew to increased $\xi_{\text{ion},0}$ values when inferred from higher order Balmer lines. However, this trend is minimal ($\lesssim 0.1$ dex) and is consistent with the selec-

tion criteria requiring at least two Balmer series lines with $S/N \geq 3$ in conjunction with the fainter intrinsic luminosities of higher-order transitions.

3 ANALYSIS

In the following section, we aim to establish relations between the ionizing photon production efficiency and a number of key galaxy observables $\mathbf{X}$; the equivalent widths of the strong optical nebular emission lines $\text{H}\alpha$ and $[\text{O III}]\lambda 5007$, the UV continuum slope and UV absolute magnitude, the nebular dust attenuation and redshift (that is $\mathbf{X}_i = \{W_\lambda(\text{H}\alpha), W_\lambda([\text{O III}]), \beta_{\text{UV}}, M_{1500}, E(B-V), z_{\text{spec}}\}$). These properties are chosen for a number of reasons, including they are: (i) routinely measurable in currently available JWST imaging and spectroscopic datasets, (ii) directly related to the physical drivers of ionizing photon production, and (iii) commonly considered in discussions related to the properties of the star-forming galaxies driving reionization and their evolution.

3.1 Correlations between predictor variables

Strong covariance between data in multivariate regression can degrade the quality of fits and inflate the uncertainties. Therefore, before introducing our fitting analysis in Section 3.2 below and to assist in interpretation of the regression models, here we first examine the pairwise and conditional correlations between the six predictors in this analysis.

Fig. 3 shows a corner plot of the six galaxy observables used in this work, with the Spearman rank coefficients listed in each panel⁴. An important first point to highlight, as evident from the figure, is that the majority of observed properties show at least mild-to-moderate correlations $|\rho_S| \approx 0.2 - 0.6$ with robust significance ($p \lesssim 10^{-3}$). We further verify these correlations from computing the Skipped correlation coefficient (otherwise known as the *robust* Spearman correlation), which has reduced sensitivity to outliers. These are illustrated in the upper triangle panels of Fig. 5, and generally suggest the standard correlation coefficients are not being driven by outliers.

These indicate that at some level, any single-observable fit (e.g., see Section 4.1 below) will implicitly fold in evolution in other covariates - a fact that should be held in consideration when being used for inferring ionizing photon production efficiencies. This would also imply that the covariances significantly influence the standard (Spearman) correlation coefficients. To test this, we also compute the partial correlation coefficient (controlling for all other variables), shown in the lower triangle of the heat map in Fig. 5. A number of key conclusions can be drawn from comparisons between the partial and standard correlation coefficients. Firstly, most of the observed correlations are considerably weaker as well as less significant based on the partial correlations. This trend is generally anticipated given all the variables show at least mild correlation with all other variables. A number of covariates show flipped correlations when controlling for those remaining (e.g., $W_\lambda(\text{H}\alpha) - \beta_{\text{UV}}$, or $W_\lambda([\text{O III}]) - M_{1500}$), and others become practically uncorrelated (e.g., $W_\lambda(\text{H}\alpha) - M_{1500}$), which is indicative of the predictor bringing unique information with which to constrain $\xi_{\text{ion},0}$.

³ Different assumptions for $\beta_{\text{UV}}^{\text{intr}}$ will systematically shift our $\xi_{\text{ion},0}$ measurements. Changes of $\Delta(\beta_{\text{UV}}^{\text{intr}}) = \pm 0.1$ will impact $L_{\text{UV}}^{\text{intr}}$ by ± 15 per cent, corresponding to $d\xi_{\text{ion},0}/d\beta_{\text{UV}}^{\text{intr}} \approx +0.5$. Specifically, adopting an alternative ≈ 0.1 bluer intrinsic slope of $\beta_{\text{UV}}^{\text{intr}} = -2.5$ will decrease $\xi_{\text{ion},0}$ by ≈ 0.05 dex

⁴ We note here that a Pearson correlation coefficient more directly measures a linear monotonic relationship, and as such the correlation coefficients tend to be weaker (and marginally less significant) when compared to the Spearman correlation coefficient.




Figure 5. Heat map of the Spearman rank correlation (ρ) matrix for the six key predictors used in this work: $W_\lambda(\text{H}\alpha)$, $W_\lambda([\text{O III}])$, β_{UV} , M_{1500} , $E(B - V)$, z_{spec} . The lower panels (below the diagonal; denoted with a $*$) show the *partial- ρ* which removes the affect of all other covariates, while the upper panels (above the diagonal; denoted with a †) display the ‘Skipped Spearman’ robust- ρ , which mitigates against the impact of outliers. All quantities, including the p -values displayed below the coefficient at the centre of each panel, are computed using PINGOUIN (Vallat 2018). In particular, the heat map brings attention to the mild-to-moderate (and in some cases very strong, e.g., between $\text{H}\alpha$ and $[\text{O III}]$ equivalent width) collinearity between the predictors, which may impact the regression parameters.

We note here, that some of the correlations will be influenced by the data limitations and sample selection adopted in the analysis. For example, an apparent correlation between redshift and absolute UV magnitude is expected and seen, given our faintest population at $M_{1500} \gtrsim -18$ is predominantly at $z \lesssim 2$.

The most strongly correlated observed properties, including when accounting for other controlled variables with the partial correlation coefficient, are the equivalent widths of $\text{H}\alpha$ and $[\text{O III}]$. Both these nebular emission lines trace the presence of intense, recent star-formation, leading to a strong correlation being expected (e.g., see Tang et al. 2019). Practically for the regression modelling, the strong relationship motivates careful consideration when fitting together, including which is a more robust predictor, which we discuss at a later point in Section 4.2.

3.2 Bayesian linear regression

To model the relation between $\xi_{\text{ion},0}$ and one or more of the predictor variables ($X_n = \{W_\lambda(\text{H}\alpha), W_\lambda([\text{O III}]), \beta_{\text{UV}}, M_{1500}, E(B - V), z_{\text{spec}}\}$), we adopt a Bayesian linear regression methodology. Specifically, we aim to determine the regression parameters, including robust inferences of the intrinsic scatter, while marginalising over the per-galaxy uncertainties.

Firstly, for each galaxy i , the regression model can be written:

$$\mu_i(\theta) = \alpha + \sum_{k=1}^K \beta_k \cdot g_k(x_{i,k}^{\text{true}}) \quad (2)$$

where K is the number of predictor variables in the model each with a characteristic gradient β_i , resulting in $K + 1$ model parameters θ , including the intercept (i.e., $\theta = \{\alpha, \beta_1, \dots, \beta_K\}$). Here, $g_k(\dots)$ represents any transformation of the parameter x_i (e.g., taking $\log_{10}$ in the

case of W_λ , or centering the variable). To robustly account for errors in the observed data, we introduce latent ‘true’ x_i variables which are related through a measurement model as: $x_i^{\text{obs}} \sim \mathcal{N}(x_i^{\text{true}}, \sigma_i)$.

To better account for the possibility of significant outliers in the linear relationships (which are not captured by a Gaussian distribution; e.g., see Fig. 6), we generalise the typical Normal distribution likelihood to a Student- t likelihood in our Bayesian formalism, given as:

$$p(y | \mu, \sigma_{\text{int}}, \nu) = t_\nu(y | \mu, \sigma_{\text{int}}^2) = \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2}) \cdot \sqrt{\nu\pi} \cdot \sigma_{\text{int}}} \left[1 + \frac{1}{\nu} \left(\frac{y - \mu}{\sigma_{\text{int}}} \right)^2 \right]^{-\frac{\nu+1}{2}} \quad (3)$$

where σ_{int} is the intrinsic scatter in the model (that is, explicitly stating our model is therefore $\mu_i(\theta) \Rightarrow \mu_i(\theta) + \epsilon_i$, and ν is the degrees of freedom in the Student- t distribution. We fit for both σ_{int} and ν as free parameters, and note that the standard Normal likelihood is recovered when $\nu \rightarrow \infty$ (in practice $\nu \gg 10$). However, the flexibility offered by a Student- t model reduces the fitting sensitivity to outliers (unlike a Gaussian configuration) and yields more robust and accurate estimates of the parameter posterior distributions.

As discussed in Section 2.5.3, for each object we obtain a $\xi_{\text{ion},0}$ distribution from our Monte Carlo error propagation (see examples shown in Fig. A3). Let D_i denote the upstream data for object i (photometry / line fluxes used to derive ξ_{ion}) and write $y \equiv \log_{10} \xi_{\text{ion},0}$. To accurately fold in the constraints provided by the per-object $\xi_{\text{ion},0}$ distributions we begin from the exact marginal likelihood of the upstream data under the regression model:

$$p(D_i | \theta) = \int p(D_i | y_i) p(y_i | \theta) dy_i, \quad (4)$$

where $p(y | \theta)$ is the Student- t predictive density introduced in Eq. 3 evaluated at $\mu_i(\theta)$ (i.e. $p(y | \theta) \equiv t_\nu(y | \mu, \sigma_{\text{int}}^2)$). In practice we do not evaluate $p(D_i | y_i)$ directly, instead treating the Monte Carlo draws $y_{i,1}, \dots, y_{i,J}$ from the per-object posterior $p(y | D_i)$ as an empirical measurement-uncertainty cloud and score the regression predictive density against it:

$$\mathcal{L}_i(\theta) = \int p(y_i | \theta) p(y_i | D_i) dy_i \approx \frac{1}{J} \sum_{j=1}^J p(y_{i,j} | \theta). \quad (5)$$

We evaluate this in log space using the numerically stable log-sum-exp identity:

$$\begin{aligned} \log p(D_i | \theta) &\approx \log \left(\frac{1}{J} \sum_{j=1}^J p(y_{i,j} | \theta) \right) \\ &= \text{logsumexp}(\{\log p(y_{i,j} | \theta)\}_{j=1}^J) - \log J. \end{aligned} \quad (6)$$

Finally, the joint marginalised log-likelihood is the sum over galaxies, $\log p(\{D_i\} | \theta) = \sum_{i=1}^N \log p(D_i | \theta)$.

We note that Eq. 5 is formally equivalent (up to a θ -independent constant) to the exact marginal likelihood in Eq. 4 when the per-object posterior is data-dominated. In our analysis the $y_{i,j}$ are produced by direct Monte Carlo propagation of the measured fluxes and Balmer-decrement dust corrections (i.e. they are not sampled assuming some strongly informative prior); we therefore adopt the estimator in Eq. 5 as our pragmatic choice for the likelihood.

3.3 Implementation using PyMC

We implement our Bayesian linear regression formalism using the PyMC framework (Salvatier et al. 2015), and sample the posterior




Figure 6. Figure demonstrating the relationships between the ionizing photon production efficiency, $Y = \xi_{\text{ion},0}$, and selected key observable covariates, $X_n = \{W_\lambda(\text{H}\alpha), W_\lambda([\text{O III}]), \beta_{\text{UV}}, M_{1500}, E(B-V), z_{\text{spec}}\}$. Each column shows $\xi_{\text{ion},0}$ against a fixed covariate X_i , with each row having a designated covariate X_j as the z -axis variable colouring the markers (with the associated colour bar at the end of each row). On-diagonal panels display the Spearman rank correlation coefficient, ρ_s , and associated p -value between the X_i covariate in a given column with $\xi_{\text{ion},0}$. The upper and lower values of ρ_s represent the 95 per cent parametric confidence intervals, calculated using PINGOUIN (Vallat 2018). We denote $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) = 25.2$ as a horizontal grey line.

distribution $p(\theta | D_i) \propto p(\theta) \cdot p(D_i | \theta)$, where $p(D_i | \theta)$ written in Eq. 6, using the No-U-Turn Sampler (NUTS). The NUTS sampler is a gradient-based (Hamiltonian Monte Carlo) algorithm for sampling Monte Carlo Markov Chains, whereby the step-size and number of steps per sample are automatically tuned (e.g., see Hoffman & Gelman 2011, for further details). For the sampling, we use four independent sampling chains, with $N_{\text{tune}} = 1000$ tuning steps and $N_{\text{sample}} = 2000$ draws at a minimum.

We adopt a compound initialisation strategy for the NUTS sampler. First we perform a ‘warm-start’ by running a variational inference (VI) optimization with PyMC ADVI (automatic differentiation variation inference); ADVI produces computationally cheap and sensible approximations of the posteriors. Samples drawn from these variational approximations are then used as the initial values for NUTS. Secondly, we use the PyMC ‘jitter+adapt_diag’ initialisation method, which enables appropriate tuning for the NUTS sampler.

Together, these two steps reduce the possibility of early divergences in more complex parameter spaces and enable NUTS to sample more efficiently.

For the default fitting configuration used throughout this work, we adopt weakly informative priors for $p(\theta)$ on all regression parameters, setting the hyperparameters after performing prior predictive checks. This allows us to regularise the inference by selecting physically plausible ranges (aided by visually inspecting our data, e.g., Fig. 6), while enabling the data to drive the posterior sampling. Specifically, we set the intercept prior and gradient prior(s) as uniform:

$$\alpha \sim \text{Uniform}(23, 27), \quad \underbrace{\beta_k}_{k=1, \dots, K} \sim \text{Uniform}(-2, +2)$$

The intrinsic scatter prior is stated as a half-normal, with the scale set empirically from the data:

$$\sigma_{\text{int}} \sim \text{HalfNormal}(\sigma_{\text{hyp}}), \quad \sigma_{\text{hyp}} = 0.5 \times \sigma_{\text{MAD}}(\langle \log_{10}(\xi_{\text{ion},0}) \rangle)$$

which anchors the scatter to the typical per-galaxy spread (and disfavours extremely large values). The level of scatter in the heavy-tailed region of the Student- t likelihood is parametrised with ν degrees of freedom and a prior as:

$$\nu + 2 \sim \text{Exponential}(\lambda = \frac{1}{30})$$

which ensures finite variance ($\nu > 2$) with the possibility of heavy tails if driven for by the data.

Lastly, we note that we performed a number of prior sensitivity tests in which we replaced the uniform regression priors with Gaussian priors (e.g., $\sim \mathcal{N}(0, 1)$), and the intrinsic scatter prior with a number of alternatives (e.g., Uniform, Half-Cauchy, Half-Student- t). The posterior estimates of any of the key parameters described by θ remained stable to these changes, with the parameter medians shifting within their $\pm 1\sigma$ uncertainties.

4 RESULTS

4.1 Single property predictions of $\xi_{\text{ion},0}$

Recent imaging and spectroscopic surveys, such as those carried out with JWST, have enabled a rich set of physical properties to be measured or inferred on an individual galaxy basis across diverse populations; including the predictors used in this work (e.g., β_{UV} , M_{UV}). In Section 4.2 below, we explore combining various observed properties to obtain the most robust predictions of $\xi_{\text{ion},0}$. Here however, we present results from linear regressions between $\xi_{\text{ion},0}$ and single observables, including the intrinsic scatter, to facilitate more direct comparisons to the literature.

4.1.1 Relation between $\xi_{\text{ion},0}$ and $W_\lambda(H\alpha)$

Of the six predictors at the focus of this work, $W_\lambda(H\alpha)$ has perhaps the clearest physical connection to the ionizing photon production, tracing the ratio of recent star-formation and hence current production of ionizing photons to the underlying continuum. The link is illustrated plainly in Fig. 6 showing the joint distributions between $\xi_{\text{ion},0}$ and one of the predictors in each column (coloured by another covariate in each row), with $H\alpha$ equivalent width on the left side. In agreement with the previous literature (e.g., see Faist et al. 2019; Atek et al. 2022; Chen et al. 2024), we see a clear positive correlation between $\xi_{\text{ion},0}$ and $W_\lambda(H\alpha)$, reporting a Spearman rank correlation coefficient of $\rho_S = 0.36$ ($p \sim 10^{-4}$). This becomes more evident when mitigating the impact of outliers,

Table 1. Correlation analysis for the EXCELS sample between the ionizing photon production efficiency, $Y = \xi_{\text{ion},0}$, and the six key physical parameters, $\mathbf{X}_n = \{W_\lambda(H\alpha), W_\lambda([O III]), \beta_{\text{UV}}, M_{\text{UV}}, E(B-V), z_{\text{spec}}\}$, displayed in Fig. 6. The **top** table displays the pairwise Pearson r correlation coefficients $\rho(Y, X_i)$, whilst the **centre** table lists the Skipped-Spearman rank correlation coefficient $\rho_{\text{robust}}(Y, X_i)$, which is considered a "robust" measurement accounting for outlier removal. Lastly, the **bottom** table presents the *partial* (Spearman rank) correlation coefficient, $\rho(Y, X_i | \{\mathbf{X}_j\}_{j \neq i})$, which controls for the multicollinearity between the physical variables. The table reports the individual correlation coefficients (ρ), the 95% confidence intervals (95% CI) and significance levels (p - value) for each covariate, calculated using PINGOUIN Python package (Vallat 2018).

Covariate	ρ	95% CI	p - value
Pearson Correlation: $\rho(Y, X_i)$			
$W_\lambda(H\alpha)$	+0.172	[-0.02, +0.35]	0.072
$W_\lambda([O III])$	+0.102	[-0.09, +0.29]	0.304
β_{UV}	-0.397	[-0.54, -0.23]	$< 10^{-4}$
M_{UV}	+0.218	[-0.03, +0.39]	0.021
$E(B-V)$	+0.657	[+0.54, +0.75]	$< 10^{-14}$
z_{spec}	-0.015	[-0.20, +0.17]	0.879
Skipped Spearman Correlation: $\rho_{\text{robust}}(Y, X_i)$			
$W_\lambda(H\alpha)$	+0.538	[+0.38, +0.66]	$< 10^{-8}$
$W_\lambda([O III])$	+0.535	[+0.37, +0.67]	$< 10^{-7}$
β_{UV}	-0.383	[-0.54, -0.20]	$< 10^{-4}$
M_{UV}	+0.110	[-0.09, +0.30]	0.278
$E(B-V)$	+0.266	[+0.07, +0.44]	$\approx 10^{-2}$
z_{spec}	+0.268	[+0.08, +0.44]	$\approx 10^{-2}$
Partial Spearman Correlation: $\rho(Y, X_i \{\mathbf{X}_j\}_{j \neq i})$			
$W_\lambda(H\alpha)$	+0.315	[+0.12, +0.48]	$\approx 10^{-3}$
$W_\lambda([O III])$	+0.089	[-0.11, +0.28]	0.386
β_{UV}	-0.698	[-0.79, -0.58]	$< 10^{-14}$
M_{UV}	+0.327	[+0.14, +0.49]	$\approx 10^{-3}$
$E(B-V)$	+0.779	[+0.69, +0.85]	$< 10^{-20}$
z_{spec}	+0.149	[-0.05, +0.34]	0.146

as revealed by the increased "Skipped" correlation coefficient in Table 1 ($\rho_{\text{robust}} = 0.54$, $p \lesssim 10^{-8}$), and a similar coefficient when controlling for the remaining predictors ($\rho_{\text{partial}} = 0.32$, $p \sim 10^{-3}$).

Applying our Bayesian fitting methodology and fitting the the logarithm of $W_\lambda(H\alpha)$ (as standard in the literature), we find the relation: $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (0.49 \pm 0.09) \times \log_{10}(W_\lambda(H\alpha) / \text{\AA}) + (23.97 \pm 0.23)$, as shown in the left panel of Fig. 7. The values and associated uncertainties are quoted as the median and 68 per cent credible intervals (16th – 84th percentiles) from the parameter marginal posterior distributions - shown in the accompanying Fig. 8⁵. Assessing the convergence diagnostics for the sampling, we confirm $\hat{R} \approx 1.00$ (indicative of stable convergence and satisfactory mixing of the chains) for each of the parameters (in addition to large effective sample sizes, and small posterior mean Monte Carlo errors (MCSE); see also Fig. A5).

As expected given the moderate positive correlation between the ionizing photon production efficiency and $W_\lambda(H\alpha)$, the median posterior gradient of the relation implies a ~ 0.5 dex increase in $\xi_{\text{ion},0}$

⁵ We note that we achieve satisfactory MCMC sampling performance in all the Bayesian linear regression results presented throughout, unless otherwise stated.




Figure 7. Relation between the ionizing photon production efficiency $\xi_{\text{ion},0}$ and $\text{H}\alpha$ equivalent width ($W_\lambda(\text{H}\alpha)$). In the **left panel** we show the median posterior relation (blue solid line, with blue shaded bands representing $\pm 1\sigma$ and 2σ) from our Bayesian linear regression to the EXCELS data (black circles) described in Section 3.2, given by: $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (0.49 \pm 0.09) \times \log_{10}(W_\lambda(\text{H}\alpha) / \text{\AA}) + (23.97 \pm 0.23)$. The intrinsic scatter in this relation is best described as $\epsilon \sim \text{StudentT}(\sigma_{\text{intr}} = 0.20 \pm 0.03, \nu = 2.67 \pm 0.75)$, highlighting that the heavy-tailed residuals from outliers to the relation (see also Fig. 8). The combined uncertainty from the intrinsic scatter and regression parameter posteriors are indicated by the magenta dashed lines. The **lower right panel** compares our relation (now in the background as a black line and grey shaded areas) with recent literature measurements (Tang et al. 2019; Faisst et al. 2019; Emami et al. 2020; Atek et al. 2022; Prieto-Lyon et al. 2023; Chen et al. 2024; Laseter et al. 2025). In the **upper right panel**, we show the difference between the predicted $\xi_{\text{ion},0}$ using our inferred relation and those from the literature ($\Delta = \log_{10}(\xi_{\text{ion},0})_{\text{This work}} - \log_{10}(\xi_{\text{ion},0})_{\text{Literature}}$), as a function of $W_\lambda(\text{H}\alpha)$.

from $W_\lambda(\text{H}\alpha) = 100-1000\text{\AA}$. The intrinsic scatter, $\sigma_{\text{intr}} = 0.20 \pm 0.03$ (assuming a StudentT($\nu \sim 2.67$) model), is comparable to the typical measurement error of $\sigma_{\xi_{\text{ion},0}} \sim 0.15$ dex which suggests that there is systematic source-to-source variation in $\xi_{\text{ion},0}$ not accounted for by $W_\lambda(\text{H}\alpha)$. A number of separate pieces of evidence substantiate this conclusion including: the heavier-than-normal residual tails (particularly at $\epsilon \gtrsim 1.0$ dex) seen in Fig. 8 (lower panel) from large outliers, the relatively moderate correlation coefficients shown in Table 1, and the auxiliary predictor dependence seen in Fig. 6.

The right panel of Fig. 7 compares our work to other literature relations (Tang et al. 2019; Faisst et al. 2019; Emami et al. 2020; Atek et al. 2022; Prieto-Lyon et al. 2023; Chen et al. 2024). Broadly speaking, there is agreement that galaxies typically cross the canonical ‘ionising galaxy’ threshold of $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) \simeq 25.2$ when they display $W_\lambda(\text{H}\alpha) \simeq 250 - 350 \text{\AA}$ (corresponding to a specific star-formation rate of $\text{sSFR} \simeq 4 - 7 \text{ Gyr}^{-1}$; Mármol-Queraltó et al. 2016). The relation presented here more closely aligns with the most recent literature of Prieto-Lyon et al. (2023) and Chen et al. (2024), deviating by no more than $\sim 0.1 - 0.15$ dex at $W_\lambda(\text{H}\alpha) \gtrsim 100 \text{\AA}$ as shown in the top-right panel. Specifically comparing with the relation in Chen et al. (2024), we find a very close slope ($\beta_1 = 0.54 \pm 0.03$) but a slightly higher normalisation ($\Delta = +0.21$ dex but within $\sim 1\sigma$). This difference in part may be due to different assumed dust attenuation laws (e.g., see Reddy et al. 2016) - the relation we compare to used a Calzetti et al. (2000) curve, however steeper laws (like the ‘intermediate’ curve adopted here) would systematically increase $\xi_{\text{ion},0}$ by up to $0.1 - 0.2$ dex. Chen et al. (2024) find a relation almost identical to the one here ($\beta_1 = 0.51 \pm 0.04$, $\beta_0 = 23.96 \pm 0.11$) when using an SMC (Gordon et al. 2003) curve (although we note that curve is steeper than ours). Other mild deviations could arise from the differences in sample selection. The sample studied in Chen et al. (2024) comprises significantly of low-mass (photometrically measured) $\text{H}\alpha$ emitters (HAEs) at $z \sim 2.3$. A systematically lower-mass sample of HAEs may be younger or burstier, which would shift the relation to higher $W_\lambda(\text{H}\alpha)$ at a given $\xi_{\text{ion},0}$ (as seen in Fig. 7).

This effect could be exacerbated by their assumed fixed stellar-to-nebular dust reddening ratio (≈ 0.8), which we find (calculating on a individual galaxy-by-galaxy basis) to have significant scatter (and a likely mass-dependence; see Fig. A2). Specifically, we find a median $\langle E(B-V)_{\text{neb}} / E(B-V)_{\text{stellar}} \rangle \simeq 1.6$, but with a large galaxy-to-galaxy scatter (quantified by a 16th–84th percentile range of $\simeq 0.0$ to $\gtrsim 5.0$) and a heavily non-Gaussian distribution (with ≈ 25 per cent of the sample ($N = 40$) consistent with zero nebular attenuation; i.e., $E(B-V)_{\text{neb}} \leq 0.02$).

In contrast, Prieto-Lyon et al. (2023) report a steeper slope ($\beta_1 = 0.73 \pm 0.04$) from a similar redshift range ($z \simeq 3 - 7$). Several methodological differences could be responsible; they adopt an SMC-like dust curve for both nebular and stellar corrections (where a greyer law would increase the discrepancy), and their sample is substantially fainter ($\langle M_{\text{UV}} \rangle \sim -18$, $\Delta \sim 1.5$ dex fainter). In this regime they also find higher median $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) \simeq 25.33$ (which is ≈ 0.11 dex above EXCELS). Selection effects may also play a role, with their sample comprising LAEs and a high-purity photometric subsample, both of which preferentially select young, strongly star-forming systems with intense line emission and elevated ξ_{ion} (Erb et al. 2016; Cullen et al. 2020; Tang et al. 2021), which may therefore bias the faint end toward objects with stronger lines, inflating inferred relations (see Pahl et al. 2025b; Simmonds et al. 2024a; Begley et al. 2025)).

Possible variations in sample properties, as well as different methods and/or assumptions may also explain the contrast between this work and the other relations shown (e.g., Tang et al. 2019; Faisst et al. 2019; Emami et al. 2020; Atek et al. 2022). For example, Atek et al. (2022) is the lowest-redshift study we compare to here. A relatively mild evolution in $\xi_{\text{ion},0}$ with redshift (e.g., as seen in Pahl et al. 2025b; Simmonds et al. 2024a; Begley et al. 2025; Llerena et al. 2025), combined with a comparatively stronger redshift dependence with $W_\lambda(\text{H}\alpha)$ (Fumagalli et al. 2012; Mármol-Queraltó et al. 2016; Reddy et al. 2018b; Khostovan et al. 2024), could shift the relation left (i.e., higher $\xi_{\text{ion},0}$ at fixed $W_\lambda(\text{H}\alpha)$). Looking at




Figure 8. Supporting figures for the Bayesian linear regression of $\xi_{\text{ion},0}$ as a function of $W_\lambda(\text{H}\alpha)$ (e.g., see Section 4.1.1 and Fig. 8). **Top:** Corner plot showing the joint and marginal posterior probability distributions from our Bayesian fitting methodology (see Section 3.2) including the linear regression parameters (slope β_1 , intercept β_0) and those describing the StudentT intrinsic scatter (and likelihood parameterisation; σ_{int} and ν). **Bottom:** Distribution of the residuals ($\epsilon = \log_{10}(\xi_{\text{ion},0}) - \log_{10}(\hat{\xi}_{\text{ion},0})$) from the $\xi_{\text{ion},0} - W_\lambda(\text{H}\alpha)$ linear regression. A clear non-normality is observed with the residuals having heavier tails, thus justifying the use of a StudentT likelihood. The residuals as a function of the independent variable (i.e., $W_\lambda(\text{H}\alpha)$), revealing no systematic offsets, is shown in the inset panel.

the relation in Tang et al. (2019), the authors find a steeper trend ($\beta_1 = 0.86 \pm 0.06$). However, a direct comparison is difficult as their sample is very low mass, low-redshift, and generally focuses on emission line galaxies (specifically, $W_\lambda([\text{O III}]+\text{H}\beta) \gtrsim 200 \text{ \AA}$).

Given the shallower relation obtained here relative to the other literature shown in Fig. 7, we explore some variations in our fiducial fitting process. Firstly, we test the sensitivity of our regression to the functional form of the likelihood model. As discussed above,

the heavy-tailed residuals statistically prefer a StudentT scatter (with $\nu \sim 2.67$), nonetheless repeating the fitting with a more typical Normal likelihood we find an even flatter slope ($\beta_1 \approx 0.41$). This fit also implies a factor $\sim 2\times$ larger intrinsic scatter, $\sigma_{\text{int}} \approx 0.36$, but fails to accurately represent the observed heavy-tailed scatter. We also test an alternative fitting methodology using the SciPy ODR package, which is commonly adopted in the literature. This returns a slightly steeper (increasing by $\sim +1\sigma$) slope of $\beta_1 \approx 0.6$, which is closer to some of the recent literature ($\beta_1 \approx 0.6 - 0.85$). However this fitting method requires collapsing the posterior distributions of $\xi_{\text{ion},0}$ for each object to $y \pm \sigma_y$ assuming Gaussian uncertainties, which for some objects is an improper assumption and disregards the full information available (see Fig. A3).

Complementary to the aforementioned tests of simpler fitting methods, we also explore further modelling components. One such extension includes γ parameterising the predictor dependence of the intrinsic scatter (known as model heteroscedasticity, i.e., given as $\sigma_{\text{int}}(x) = \sigma_{\text{int},0} \cdot (x/x_0)^\gamma$). Here we find $\gamma \approx -0.3 \pm 0.5$ (at $\lesssim 1\sigma$ significance, with the remaining parameters consistent at the $\leq 1\sigma$ level), implying more population scatter at lower $\text{H}\alpha$ equivalent widths. Robust inferences of any possible evolution in σ_{int} with $W_\lambda(\text{H}\alpha)$ (or any other predictor) will require much larger sample sizes however.

As an alternative to the StudentT model, we also explore a Gaussian mixture model as another credible model configuration. Specifically, we parameterise the likelihood as two Normal distributions describing two distinct populations: $p(y | \theta, \dots) \sim (1 - a) \cdot \mathcal{N}(\mu(\theta), \sigma_{\text{int},0}^2) + a \cdot \mathcal{N}(\mu(\theta), \sigma_{\text{int},1}^2)$, where a is the mixing fraction and $\sigma_{\text{int},0} > \sigma_{\text{int},1}$. Median values (from the sampled posterior distributions) indicate $\sigma_{\text{int},0} = 0.06 \pm 0.05$, $\sigma_{\text{int},1} = 0.31 \pm 0.05$, and $a = 0.05 \pm 0.03$. We compare the performance of the mixture and StudentT models using leave-one-out cross validation (LOO), finding neither are statistically preferred⁶. As a result, the StudentT model is our fiducial setup given it is more interpretable and has lower complexity (with less free parameters).

4.1.2 Relation between $\xi_{\text{ion},0}$ and $W_\lambda([\text{O III}])$

Similar to $\text{H}\alpha$, the $[\text{O III}]\lambda 5007$ emission line is directly produced as a result of young, massive stars powering their surrounding ionized nebulae. This is clearly visualised in Fig. 6, and backed-up quantitatively by the correlation coefficient with $\xi_{\text{ion},0}$ ($\rho_{\text{robust}} = 0.54$, $p < 10^{-7}$). However, $W_\lambda([\text{O III}])$ is also secondarily sensitive to the ionization parameter, spectral hardness and metallicity, all of which may enhance or suppress $W_\lambda([\text{O III}])$ at fixed ξ_{ion} (e.g., see Reddy et al. 2018b; Laseter et al. 2025). These confounding factors may explain why the partial correlation is much lower ($\rho_{\text{partial}} = 0.09$, $p = 0.39$) than for $W_\lambda(\text{H}\alpha)$. Even when $\text{H}\alpha$ and $[\text{O III}]$ strongly correlate ($\rho_S = 0.93$, $p < 10^{-6}$) and thus provide overlapping information on $\xi_{\text{ion},0}$, $\text{H}\alpha$ is less impacted by secondary factors and therefore is the more dominant predictor.

Nonetheless, $[\text{O III}]$ remains an important avenue from which to empirically infer ξ_{ion} (e.g., see Tang et al. 2019; Begley et al. 2025), and applying the Bayesian regression methodology we determine:

⁶ Here, we compare models using the *expected log pointwise predictive density* (elpd) calculated with the PSIS-LOO algorithm (Pareto-smoothed importance sampling; Vehtari et al. 2015, implemented within ArViz; Kumar et al. 2019). In this scheme, one model is considered moderately *better* than another if $\Delta \text{elpd} \gtrsim 2 \cdot \text{SE}_\Delta$, where SE_Δ is the standard error on the elpd estimate.




Figure 9. Relation between $\xi_{\text{ion},0}$ and $W_\lambda([\text{O III}])$ (in the same format as the right panel of Fig. 10), given as: $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (0.30 \pm 0.08) \times W_\lambda([\text{O III}]) + (24.44 \pm 0.21)$, with intrinsic scatter StudentT($\sigma_{\text{int}} = 0.20 \pm 0.03$, $\nu = 2.56 \pm 0.63$) (shown in the **top panel** as black line with the $\pm 1\sigma$, $\pm 2\sigma$ regions as grey shaded and the σ_{int} denoted as thin dashed lines). We compare to literature results at cosmic noon (e.g., with samples of EELGS and HAEs / non-HAEs Tang et al. 2019; Chen et al. 2024, as black dash-dotted and dotted lines, respectively), as well as those spanning a wider high-redshift regime ($z \sim 4 - 10$ Simmonds et al. 2024a; Pahl et al. 2025b; Llerena et al. 2025; Papovich et al. 2025,; see legend). In the **lower panel**, we show the difference in the predicted ionizing photon production efficiency as function of $W_\lambda([\text{O III}])$ ($\Delta \log_{10}(\xi_{\text{ion},0})$).

$\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) = (0.30 \pm 0.08) \times \log_{10}(W_\lambda([\text{O III}]) / 15007) + (24.44 \pm 0.21)$. We show this relation in comparison to recent literature in Fig. 9. Our relation suggests only a $\Delta \sim 0.4$ dex evolution from $W_\lambda([\text{O III}]) = 100 \text{ \AA}$ to $W_\lambda([\text{O III}]) = 1000 \text{ \AA}$, with scatter at the $\sigma_{\text{int}} \approx 0.2$ dex level (and $\nu \sim 2.56$ in t_ν). We find the closest agreement with (Chen et al. 2024), differing by $\Delta \lesssim 0.05 - 0.1$ dex across the full dynamic range in our sample. This dependence is well matched by simulations such as FLARES, as shown in Seeay et al. (2023). The relation presented here and in Chen et al. (2024) (see also Emami et al. 2020; Atek et al. 2022, not shown, with similar slopes, but systematically offset as in Section 4.1.1) is shallower than other results shown in Fig. 9 ($\beta_1 = 0.30 \pm 0.08$ compared to $\beta_1 \sim 0.4 - 0.8$). This difference can likely be ascribed to dissimilarities in the samples used. Firstly, we highlight that the majority of the literature does not probe the low- $W_\lambda([\text{O III}])$ ($\lesssim 100 \text{ \AA}$) regime, whereas here this constitutes ~ 27 per cent of the sample subset with $[\text{O III}]$ measurements. Secondly, an underlying redshift dependence could be altering the inferred relations. In the higher redshift samples of Simmonds et al. (2024b); Pahl et al. (2025b); Llerena et al. (2025), lower metallicities or harder ionizing spectra may produce a steeper $\xi_{\text{ion},0} - W_\lambda([\text{O III}])$ dependence (Nakajima et al. 2016; Hayes et al. 2025, and see Section 4.2.4).

It is noteworthy that sample *selection* is unlikely to be a significant contributing factor here. The EXCELS parent sample is selected mostly from large multi-band photometric catalogues (i.e., not emission line selected, and tracing the star-forming main sequence) while the spectroscopy has sufficient depth to provide sensitivity down to very faint flux limits ($f_{\text{line}} \sim 10^{-18} / \text{erg s}^{-1} \text{cm}^{-2}$). Nonetheless, on the basis of our requirement to detect a minimum of two Balmer

emission lines at $\geq 3\sigma$, beyond a certain UV luminosity limit only the strongest emission lines will satisfy this criteria.

4.1.3 Relation between $\xi_{\text{ion},0}$ and β_{UV}

Based on Fig. 6 and the correlation coefficients in Table 1 (i.e., $\rho_{\text{partial}} \approx -0.7$, $p \ll 10^{-6}$), the UV continuum slope has one of the strongest connections with $\xi_{\text{ion},0}$. Although strongly sensitive to the degree of dust attenuation, β_{UV} is likewise indicative of the age and metallicity of the stellar populations present; with bluer galaxies expected to be younger, have lower-metallicities and present harder ionizing spectra - all properties conducive to elevated $\xi_{\text{ion},0}$ (Bouwens et al. 2014, 2016; Topping et al. 2024).

From our regression, we infer the relation: $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (-0.22 \pm 0.06) \times \beta_{\text{UV}} + (23.97 \pm 0.23)$, shown in Fig. 10. Broadly speaking, galaxies in the general population with UV slopes bluer than $\beta_{\text{UV}} \lesssim -1.6$ have ionizing photon production efficiencies around or above $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) \approx 25.2$, in approximate agreement with stacking-based analyses in the literature (e.g., Bouwens et al. 2016; Chen et al. 2024, as shown in the right panel of Fig. 10).

This relation agrees remarkably well with the photometric analysis of star-forming galaxies at $z \approx 6.9 - 7.6$, presented in our previous work (Begley et al. 2025), supporting a redshift-invariant $\xi_{\text{ion},0} - \beta_{\text{UV}}$ relation. This is at odds with the lack of any correlation recovered by Pahl et al. (2025b). Interestingly, we also find excellent agreement between the gradient ($\beta_1 = \partial \xi_{\text{ion},0} / \partial \beta_{\text{UV}}$) of the relation here and that presented in Prieto-Lyon et al. (2023); but with a ≈ 0.2 dex offset in $\xi_{\text{ion},0}$ normalisation. We remind the reader of the ≈ 0.1 dex higher median $\xi_{\text{ion},0}$ (and ≈ 1.5 mag fainter) sample in Prieto-Lyon et al. (2023), and suggest the arguments made in Section 4.1.1 could also explain this offset seen in $\xi_{\text{ion},0} - \beta_{\text{UV}}$.

As with the previous predictors, we find a moderate and heavy tailed scatter given as $\sigma_{\text{int}} = 0.23 \pm 0.04$ and StudentT($\nu \approx 3.22$). A key observation that we highlight here is that the outliers are discernible from the typical star-forming galaxy population that makes up the majority of the EXCELS sample - namely they show high $E(B - V)_{\text{neb}}$ (e.g., see Fig. 6). High $\xi_{\text{ion},0}$ in combination with high $E(B - V)_{\text{neb}}$, which describes a large fraction of the heavy-tailed outliers, may be indicative of sources that contain young obscured star-forming regions. However, it is also probable that these types of galaxies do not yet contribute significantly to the ionizing photon budget, as the high $E(B - V)_{\text{neb}}$ would indicate very little of the ionizing photons produced could escape due to the dust attenuation. Another important consideration is that we observe a moderate correlation between $E(B - V)_{\text{neb}}$ and β_{UV} ($\rho_S \approx 0.39$, $p \lesssim 10^{-4}$), but again these sources are typically outliers in the $E(B - V)_{\text{neb}} - \beta_{\text{UV}}$ as well. This aligns with the large scatter seen between stellar and nebular attenuation, particularly at high redshifts (Reddy et al. 2018b, see also Fig. A2), and that these galaxies lie in the extreme high $E(B - V)_{\text{neb}} / E(B - V)_{\text{stellar}}$ tail. Complicating the interpretation of these objects, is the underlying assumption of a foreground dust geometry when using the Balmer decrement (and Case B) to obtain the intrinsic Balmer line luminosity. It may be the case that a more mixed, clumpy dust geometry is more appropriate (producing greyer attenuation curves; e.g., see Vijayan et al. 2024, for a more in-depth discussion), in which case the Balmer decrement method leads to over-estimated $\xi_{\text{ion},0}$.




Figure 10. Inferred $\xi_{\text{ion},0} - \beta_{\text{UV}}$ relation, shown in the same two-panel layout (and formatting) as Fig. 7. On the **left panel**, we show the median posterior relation constrained by the data: $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (-0.22 \pm 0.06) \times \beta_{\text{UV}} + (23.97 \pm 0.23)$, with intrinsic scatter $\text{StudentT}(\sigma_{\text{intr}} = 0.23 \pm 0.04, \nu = 3.22 \pm 1.46)$. We compare to a number of different literature constraints in the **right panel**, including relations from Prieto-Lyon et al. (2023) and Begley et al. (2025); results from stacking (Bouwens et al. 2016; Shivaiei et al. 2018; Chen et al. 2024); and the running mean across the spectroscopic sample analysed in Pahl et al. (2025b).

4.1.4 Relation between $\xi_{\text{ion},0}$ and M_{UV}

The rest-frame UV luminosity dependence of the ionizing photon production is vital for accurate empirical modelling of the reionization epoch (e.g., see Donnan et al. 2025). Here we establish a $\approx 2\sigma$ relation given by: $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (0.08 \pm 0.04) \cdot (M_{\text{UV}} + 20) + (25.13 \pm 0.04)$, implying $\Delta \sim 0.3$ dex evolution in $\xi_{\text{ion},0}$ across a ~ 4 mag dynamic range ($M_{\text{UV}} \sim -21$ to $M_{\text{UV}} \sim -17$) (with $\epsilon \sim \text{StudentT}(\sigma_{\text{int}} = 0.24 \pm 0.04, \nu = 2.98 \pm 1.09)$). We show this, alongside comparisons to recent observational results in the top two panels of Fig. 11. What is fully transparent considering the landscape of recent literature on the $\xi_{\text{ion},0} - M_{\text{UV}}$ relation (i.e., middle panel of Fig. 11), is the lack of consensus. A number of studies, like the result here, find that elevated $\xi_{\text{ion},0}$ are found in more M_{UV} -faint populations (Prieto-Lyon et al. 2023; Llerena et al. 2025; Papovich et al. 2025). This trend would not be surprising as M_{UV} -faint galaxies will typically have younger ages, lower metallicities, and lower dust content (Reddy et al. 2018a; Cullen et al. 2019, 2020; Curti et al. 2024; Jain et al. 2025, and Stanton et al. 2025b in prep.), all of which are expected to drive higher $\xi_{\text{ion},0}$ (Reddy et al. 2018b; Shivaiei et al. 2018). Nevertheless, we highlight that the $\xi_{\text{ion},0} - M_{\text{UV}}$ relations are mostly still very shallow and only marginally significant (e.g., Prieto-Lyon et al. 2023 with $\beta_1 \approx 0.03 \pm 0.02$; Papovich et al. 2025 with $\beta_1 \approx 0.06 \pm 0.03$ at $4.5 \leq z \leq 5.5$). This picture is made further ambiguous by the remaining confounding variables that are not taken into account when considering only M_{UV} .

Caution should inherently be taken when considering inferred trends in $\xi_{\text{ion},0}$ with M_{UV} , as there is mounting evidence that sample selection can have a nuanced but critical impact. Begley et al. (2025) from a photometrically-selected ($\gtrsim 90$ per cent complete at $M_{\text{UV}} \leq -18$) sample at $z \sim 6.9 - 7.6$, find that brighter M_{UV} SFGs have more intense $\xi_{\text{ion},0}$ contrary to the literature previously discussed. Recent spectroscopic work by Pahl et al. (2025b), and a carefully selected mass-complete sample from Simmonds et al. (2024a) (comparing with Simmonds et al. 2024b, a clear example of the importance of the impact of completeness), also mirror this result. This trend can be understood from the perspective of star-formation burstiness, which increases both towards higher redshifts and lower stellar masses (Atek et al. 2022; Endsley et al. 2022, 2024; Begley

et al. 2025). Specifically, in this scenario M_{UV} -bright galaxies have undergone a recent strong up-turn in star-formation providing a boost to $\xi_{\text{ion},0}$ (Endsley et al. 2025). On the other hand, M_{UV} -faint sources are in a current ‘lulling’ phase (e.g., see Looser et al. 2023, but also Carvajal-Bohorquez et al. 2025; Cole et al. 2025; Simmonds et al. 2025; Trussler et al. 2025 for recent literature discussions on bursty star-formation). Crucially however, like in the case for relations where the fainter galaxies are seen to have higher $\xi_{\text{ion},0}$, the $\xi_{\text{ion},0} - M_{\text{UV}}$ trends in Simmonds et al. (2024a); Begley et al. (2025) and Pahl et al. (2025b) are mild, close to flat, and relatively insignificant ($\lesssim 2\sigma$). Here we note that when only restricting our sample to galaxies with $M_{\text{UV}} \leq -19$, effectively making our sample more comparable to the literature discussed here (broadly removing the extreme low-redshift tail at $z < 2$, and where the M_{UV} distribution begins to fall off towards fainter- M_{UV}), the relation is fully consistent with flat at the $\lesssim 1\sigma$ level ($\beta_1 \approx 0.06^{+0.07}_{-0.08}$).

It is also possible that the importance (or lack thereof) of UV luminosity varies with redshift as well. The running median on the $\xi_{\text{ion},0} - M_{\text{UV}}$ plane from the combined (VANDELS + JWST) sample in Begley et al. (2025) is entirely consistent with being flat (as shown in Fig. 11). Moreover, current evidence suggests the $\xi_{\text{ion},0} - M_{\text{UV}}$ relation slope varies only mildly, if at all, with redshift (Pahl et al. 2025b; Simmonds et al. 2024a; Papovich et al. 2025).

The lower panel of Fig. 11 considers results from some of the latest cosmological hydrodynamical simulations; SPHINX (e.g., see Rosdahl et al. 2018; Katz et al. 2023) and FLARES (e.g., see Lovell et al. 2021; Seeyave et al. 2023). These two simulation suites together show an extremely mild-to-flat connection between UV luminosity and the ionizing photon production efficiency. There is also a minor modulation based on redshift, but this is due to the underlying physical property evolution. Fig. 11 also includes the latest results from the DELPHI semi-analytic model (Mauerhofer et al. 2025) at $z \sim 5$, including their ‘fiducial’ model (incorporating ISM physics derived from the SPHINX high-resolution radiative-hydrodynamical simulations) and their ‘eIMF’ model (folding in a redshift and metallicity dependent IMF). The fiducial model shown, which reliably reproduces observations of UV continuum slopes and Balmer jump strengths down to the end of the EOR, predicts an incredibly flat




Figure 11. Results for the relation between $\xi_{\text{ion},0}$ and the absolute UV luminosity (M_{UV}), discussed in Section 4.1.4. The **top** panel shows the $\log_{10}(\xi_{\text{ion},0} / \text{Hz erg}^{-1}) = (0.08 \pm 0.04) \times (M_{\text{UV}} + 20) + (25.13 \pm 0.04)$ relation from the Bayesian linear regression (median posterior as a blue line with $\pm 1\sigma$, $\pm 2\sigma$ shaded regions), with the intrinsic scatter indicated by the purple dashed lines; $\epsilon \sim \text{StudentT}(\sigma_{\text{int}} = 0.24 \pm 0.04, \nu = 2.98 \pm 1.09)$. The **centre** panel compares our work to observational analyses, indicating when M_{UV} -faint galaxies are predicted to have elevated $\xi_{\text{ion},0}$ (Prieto-Lyon et al. 2023; Llerena et al. 2025; Papovich et al. 2025, in yellow-to-red colour dashed lines) and when M_{UV} -bright galaxies have higher $\xi_{\text{ion},0}$ (Simmonds et al. 2024a; Begley et al. 2025; Pahl et al. 2025b, in blue-to-purple colour solid lines); also including the running median of the full $z \sim 3 - 8$ VANDELS+JWST sample studied in Begley et al. (2025). In parallel, we show results from leading high-resolution hydrodynamical simulations in the **bottom** panel, as a function of redshift (Seeyave et al. 2023; Katz et al. 2023, for FLARES and SPHINX, respectively), and for the latest iterations of the DELPHI semi-analytic model (e.g., see Mauerhofer et al. 2025, for details).

dependence of $\xi_{\text{ion},0}$ on M_{UV} . In mild contrast, the DELPHI ‘eIMF’ model appears to agree well with the results found here, which lends support to the possibility that any M_{UV} dependence is driven by an underlying IMF evolution (although we highlight as discussed previously, our relation is not statistically significant and may be impacted by sample bias in the M_{UV} -faint regime). An important aspect to consider in relation to the simulations is that they slightly under-represent the diversity seen in $\xi_{\text{ion},0}$ across observed populations (here; ~ 0.4 dex and spanning $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) \approx 24 - 27$, but see also Prieto-Lyon et al. 2023; Simmonds et al. 2024a; Begley et al. 2025; Pahl et al. 2025b; Laseter et al. 2025), particularly at the low- and high- $\xi_{\text{ion},0}$ tails. In contrast, for example, the SPHINX simulation has $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) \approx 24.8 - 25.8$, with $\sigma_{\text{SAD}} \approx 0.2$ dex. This can be explained by the adopted SED assumptions (e.g., see Katz et al. 2023, their Section 3.1.12), with possible impact from physics below the resolution scale of the simulations. Moreover, these values are intrinsic $\xi_{\text{ion},0}$ with sightline dependence and considerations of the complex stochastic SFHs not accounted for, as well as potential observational effects.

4.1.5 Relation between $\xi_{\text{ion},0}$ and $E(B - V)_{\text{neb}}$

From Fig. 6, a clear linear and monotonic $\xi_{\text{ion},0} - E(B - V)_{\text{neb}}$ relation is not expected, given the relatively shallow trend at low $E(B - V)_{\text{neb}}$ (< 0.5) followed by a tail of both high $\xi_{\text{ion},0}$ and high $E(B - V)_{\text{neb}}$. Quantitatively, things are no more transparent given the moderate Spearman rank correlation coefficient of $\rho_S \approx 0.38$ ($p < 10^{-4}$; and $\rho_{\text{partial}} \approx 0.78$), that significantly drops to $\rho_{\text{robust}} \approx 0.27$ when using a robust coefficient ($p \approx 0.01$; see Table. 1).

Fitting a linear relationship, we find median posterior parameters of $\beta_1 = 1.24 \pm 0.17$, $\beta_0 = 24.90 \pm 0.05$ and $\epsilon \sim \text{StudentT}(\sigma_{\text{int}} = 0.24 \pm 0.04, \nu = 12.3^{+27.1}_{-7.6})$. This fit however produces a relatively broad residual distribution; see Fig. 8 for an example) $\sigma_{\text{res}} \approx 0.38$ dex, compared with $\sigma_{\text{res}} \approx 0.30$ dex for the $\xi_{\text{ion},0} - W_{\lambda}(\text{H}\alpha)$ relation. Alternatively, to better fit the trend observed in Fig. 6, we fit a power-law of the form $\xi_{\text{ion},0} = \beta_1 \cdot E(B - V)_{\text{neb}}^{\alpha} + \beta_0$ finding: $\alpha = 1.23 \pm 0.16$, $\beta_1 = 1.99 \pm 0.41$, and $\beta_0 = 25.06 \pm 0.05$ with $\epsilon \sim \text{StudentT}(\sigma_{\text{int}} = 0.24 \pm 0.04, \nu = 10.3^{+24.2}_{-5.9})$. The residuals from this functional fit remain broader than for $\text{H}\alpha$ or $[\text{O III}]\lambda 5007$, with $\sigma_{\text{res}} \approx 0.34$ dex.

Rather than being a unique single-variate predictor of the ionizing output of a galaxy, Fig. 6 qualitatively shows that $E(B - V)_{\text{neb}}$ may be crucial in explaining strong outliers in the relatively tight connections between $\xi_{\text{ion},0}$ and other observables, namely $W_{\lambda}(\text{H}\alpha)$, $W_{\lambda}(\text{O III})$ and β_{UV} (as well as the scatter in M_{UV} to a minor extent). Quantitatively, this is seen by the $\nu > 10$ posterior values preferred in the scatter of the $\xi_{\text{ion},0} - E(B - V)_{\text{neb}}$ relations fit, indicating the tail heaviness has been suppressed significantly (with ≈ 98.7 per cent of the probability within $\pm 3\sigma$, compared with ≈ 99.7 per cent under normality). As discussed in Section 4.1.3, these outliers are likely sources with obscured, but intense star-formation. These systems will have high differential nebular-to-stellar dust attenuation (and possibly more complex geometry) leading to high inferred $\xi_{\text{ion},0}$.

4.1.6 Relation between $\xi_{\text{ion},0}$ and z

From the EXCELS sample spanning a wide $z \sim 1 - 8$ redshift range, we recover a very mild (non-significant) redshift evolution given as: $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) = (0.03 \pm 0.02) \cdot z + (25.03 \pm 0.10)$. Such a shallow redshift evolution supports recent work by Simmonds et al.

(2024a), Begley et al. (2025) and Pahl et al. (2025b) in which the average $\xi_{\text{ion},0}$ only increases by $\Delta \sim 0.1 - 0.2$ dex between cosmic noon and the epoch reionization. The mild evolution is primarily driven by the large diversity in SFG populations unearthed by JWST (i.e., not being restricted to the brightest galaxy population subsets) which suppresses the redshift evolution that was widely anticipated by literature in the pre-JWST era (Finkelstein et al. 2019; Stefanon et al. 2022).

Given the shallow redshift dependence, the relation intrinsic scatter significantly outweighs the underlying evolution, here given by $\epsilon \sim \text{StudentT}(\sigma_{\text{int}} = 0.25 \pm 0.04, \nu = 3.19 \pm 1.27)$. This is in excellent agreement with estimates of the $\xi_{\text{ion},0}$ population scatter in Begley et al. (2025) ($\approx 0.2 - 0.3$ dex). We note that we do not observe the increase in intrinsic scatter with redshift (~ 0.2 dex at $z \sim 3 - 4$ to ~ 0.3 dex at $z \sim 7.0 - 7.5$; Begley et al. 2025). This is likely due to differences in sample properties, as the $z \sim 3 - 4$ sample in Begley et al. (2025) (from VANDELS) is relatively bright ($M_{\text{UV}} \in [-19.5, -22.0]$), whereas the range in M_{UV} for the EXCELS sample at similar redshifts is shifted $\Delta \sim 1.5$ mag fainter. In parallel, only a narrow magnitude regime ($M_{\text{UV}} \sim -19.5 \pm 0.5$) is covered by the highest redshift subset of the EXCELS sample (compared to $M_{\text{UV}} \in [-17.75, -21.0]$ at $z \sim 7.0 - 7.5$ in Begley et al. 2025).

4.2 Multivariate fitting

As shown in Section 4.1, a number of observed properties trace the ionizing photon production relatively well, including $W_{\lambda}(\text{H}\alpha)$, $W_{\lambda}([\text{O III}])$ and β_{UV} . For example, in the $\xi_{\text{ion},0} - W_{\lambda}(\text{H}\alpha)$ regression, we find an intrinsic scatter of $\sigma_{\text{int}} \approx 0.2$ dex, with the model predicting $\sigma(\mu_{\theta}) \equiv \sqrt{\text{Var}(\mu_{\theta})} \approx 0.18$ dex galaxy-to-galaxy spread across the sample population. We can also assess the predictive power of the model using $R_{\text{intr}}^2 = \text{Var}(\mu_{\theta}) / (\text{Var}(\mu_{\theta}) + \sigma_{\text{int}}^2)$, which measures the intrinsic galaxy-to-galaxy variance accounted for by the model variables (excluding measurement noise), as well as $R_{\text{Bayes}}^2 = \text{Var}(\mu_{\theta}) / (\text{Var}(\mu_{\theta}) + \text{Var}(\epsilon_{\text{residuals}}))$ ⁷, which includes all sources of scatter (i.e., measurement noise and intrinsic scatter; see Gelman et al. 2019), and evaluates how well the *full* variance is matched.

From the regression between $\xi_{\text{ion},0}$ and $W_{\lambda}([\text{O III}])$, we find $R_{\text{intr}}^2 = 0.34$ and $R_{\text{Bayes}}^2 = 0.10$. Together these metrics indicate that $W_{\lambda}([\text{O III}])$ show a statistically significant connection with $\xi_{\text{ion},0}$ and reasonable predictive power (with a $\approx +0.3$ dex increment in $\xi_{\text{ion},0}$ from a ≈ 1 dex increase in $W_{\lambda}([\text{O III}])$ in addition to $\text{RMAE} \approx 0.33$ dex⁸). Nonetheless, $W_{\lambda}([\text{O III}])$ (or $W_{\lambda}(\text{H}\alpha)$ ⁹) as lone predictors only accounts for ~ 10 per cent of the total *observed* variance. This is due to the fact that the intrinsic scatter ($\sigma_{\text{int}} \approx 0.2$ dex), combined with the measurement uncertainty across the dataset, and specifically in $\xi_{\text{ion},0}$ (e.g., see Fig. A3), is relatively large compared

with the predicted model spread $\sigma(\mu_{\theta})$. Moreover, a single property-based prediction of $\xi_{\text{ion},0}$ neglects the impact of other physical properties on $\xi_{\text{ion},0}$ (as discussed throughout this work, e.g., see Section 4.1 and Fig. 6).

Motivated by the necessity to improve models of the EOR, and in particular by providing more robust predictions of the ionizing photon production in the high-redshift SFG population, below we explore multivariate trends with $\xi_{\text{ion},0}$.

4.2.1 Joint modelling of $\xi_{\text{ion},0}$ and $\mathbf{X}$

Several of the measured properties explored here tie closely with the physical processes relating to the ionizing photon production efficiency; for example, the nebular emission line strengths ($W_{\lambda}(\text{H}\alpha)$ and $W_{\lambda}([\text{O III}])$) trace the presence of young massive stars. With the goal of improving the predictive metrics for $\xi_{\text{ion},0}$, we apply the Bayesian linear regression framework described in Section 3.2 to our full set of six key observables, $\mathbf{X}_n = \{W_{\lambda}(\text{H}\alpha), W_{\lambda}([\text{O III}]), \beta_{\text{UV}}, M_{1500}, E(B - V), z_{\text{spec}}\}$. The predictive power gained when moving from a single measurement prediction (e.g., $W_{\lambda}(\text{H}\alpha)$), to folding in the full set of observables is clearly illustrated in Fig. 12 and Fig. 13. Here we show a comparison between the observed and model predicted $\xi_{\text{ion},0}$ (with the full posterior corner plot shown in Fig. A4), finding practically all the modelled variance can be explained by the model ($R_{\text{intr}}^2 \approx 0.99$). This indicates that together, the chosen variables thoroughly trace the underlying physical processes that drive the range in $\xi_{\text{ion},0}$. This is qualitatively supported by the large population diversity spanned by the model ($\sigma(\mu_{\theta}) \approx 0.47$ dex), in combination with the little remaining intrinsic scatter inferred from the Bayesian model ($\sigma_{\text{int}} \approx 0.035 \pm 0.02$). Crucially, we find that multivariate median posterior model has significantly more accurate predictive power, with $\text{RMAE} \approx 0.16$ dex (a factor of two lower than the best-performing single property regression with $W_{\lambda}(\text{H}\alpha)$).

4.2.2 Identifying a minimal set of predictive observables

From the multivariate regression, we recover (conditional) relations between $\xi_{\text{ion},0}$ and UV luminosity and redshift of $\beta_{M_{\text{UV}}} = 0.04 \pm 0.03$ and $\beta_z = 0.01 \pm 0.02$, respectively. These parameters already had weak, marginal trends with $\xi_{\text{ion},0}$ from their individual regressions (see Sections 4.1.4 and 4.1.6), and now appear consistent with no relation at the $< 2\sigma$ level when fitting to the full observed data. This strengthens the argument that these weak trends (or otherwise as found in the literature; e.g., see Pahl et al. 2025b; Llerena et al. 2025 and Fig. 11) are due to more fundamental evolution in the properties of star-forming galaxies.

Another important consideration is the strong collinearity between $W_{\lambda}(\text{H}\alpha)$ and $W_{\lambda}([\text{O III}])$, as shown in Fig. 3. Such tight correlations can potentially lead to unstable Monte Carlo sampling, and inflate individual posterior slope distributions. Here we find that the uncertainties in $\beta_{W_{\lambda}(\text{H}\alpha)}$ and $\beta_{W_{\lambda}([\text{O III}])}$ are a factor of $\approx 1.5 - 1.8$ larger than in the single predictor fits. Moreover, the individual conditional gradients become shallower ($\beta_{W_{\lambda}(\text{H}\alpha)} = 0.43^{+0.17}_{-0.16}$ and $\beta_{W_{\lambda}([\text{O III}]}) = 0.19^{+0.12}_{-0.13}$) compared to the fits in Sections 4.1.1 and 4.1.2; $\beta_{W_{\lambda}(\text{H}\alpha)} = 0.49 \pm 0.09$ and $\beta_{W_{\lambda}([\text{O III}]}) = 0.30 \pm 0.08$.

We therefore perform two complementary fits, using one of $W_{\lambda}(\text{H}\alpha)$ or $W_{\lambda}([\text{O III}])$, combined with the remaining properties. As expected, we find steeper gradients with $\xi_{\text{ion},0}$ for both emission lines, and tighter posteriors compared with the full fit above:

⁷ Notably, R_{Bayes}^2 accounts for the added uncertainty from the intrinsic and measurement scatter. Specifically, μ_{θ} is the latent model prediction (i.e., drawing from the model posteriors), and $\epsilon_{\text{residuals}} = \tilde{\mu}_{\theta} - y_{\text{obs}}$, where y_{obs} is drawn from the $\xi_{\text{ion},0}$ posterior distributions (e.g., see Fig. A3) and $\tilde{\mu}_{\theta}$ represents posterior-predictive draws from the model: $\tilde{\mu}_{\theta} = \mu_{\theta} + \epsilon_{\text{intr}}$. The intrinsic scatter is drawn from $\epsilon^{\text{intr}} \sim \text{StudentT}(\sigma_{\text{int}}, \nu)$.

⁸ We define the *robust* median absolute deviation, $\text{RMAE} = \{|\epsilon_{\text{residuals}}|\}$, where $\epsilon_{\text{residuals}}$ is as above, and we take the median to mitigate against the impact of strong outliers in the metric.

⁹ The $\xi_{\text{ion},0} - W_{\lambda}(\text{H}\alpha)$ relation gives $R_{\text{intr}}^2 = 0.43$ and $R_{\text{Bayes}}^2 = 0.11$. According to the relation a ≈ 1 dex increase in $W_{\lambda}(\text{H}\alpha)$ produces a ≈ 0.5 dex increase in $\xi_{\text{ion},0}$, with $\text{RMAE} = 0.33$ (or $\text{RMAE} = 0.20$ from the model alone).




Figure 12. Left panel: Observed $\xi_{\text{ion},0}$ versus the median posterior predicted $\xi_{\text{ion},0}$ from the full multivariate Bayesian linear regression to $W_{\lambda}(\text{H}\alpha)$, $W_{\lambda}([\text{O III}])$, β_{UV} , M_{1500} , $E(B - V)_{\text{neb}}$ and z_{spec} . The 1 : 1 line is shown as black dashed diagonal, and the orange shaded region represents ± 0.2 dex. The zoom inset shows the ± 1.0 dex region around the canonical value required for reionization, $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) = 25.2$. The text within the figure shows a number of key metric and model statistics: $R_{\text{intr}}^2 = 0.99$, $R_{\text{Bayes}}^2 = 0.60$, $\sigma_{\text{int}} = 0.035$ and $\sigma(\bar{\mu}) = 0.472$ (see Section 4.2 for details). Right panels: The regression residuals ($\epsilon = \xi_{\text{ion},0} - \hat{\xi}_{\text{ion},0}$) as a function of each parameter in the model (one panel per property). The figures demonstrate that the model recovers the observed dynamic range in $\xi_{\text{ion},0}$, without any systematic offsets as a function of the physical properties.




Figure 13. Relation between ionizing photon production efficiency $\xi_{\text{ion},0}$ and $W_{\lambda}(\text{H}\alpha)$ from the full multivariate linear regression (see also Fig. A4 for the parameter posterior distributions). In the left panel we show lines of constant β_{UV} with $\beta_{\text{UV}} = -3.0$, -1.5 , 0.0 , and in the right panel we show lines of fixed $E(B - V)_{\text{neb}}$, with $E(B - V)_{\text{neb}} = 0.0$, 0.25 , 1.0 . In each panel, we evaluate the remaining parameters at the sample median values ($\langle W_{\lambda}([\text{O III}]) \rangle \approx 250 \text{\AA}$, $\langle M_{1500} \rangle \approx -19.5$, $\langle z_{\text{spec}} \rangle \approx 3.25$, and where applicable $\langle \beta_{\text{UV}} \rangle \approx -1.6$ and $\langle E(B - V)_{\text{neb}} \rangle \approx 0.18$), while marginalising over their associated gradient posterior distributions. The inset panels show the joint relation between the primary and auxiliary independent variables.

$\beta_{W_{\lambda}(\text{H}\alpha)} = 0.65 \pm 0.06$ and $\beta_{W_{\lambda}([\text{O III}]}) = 0.51^{+0.05}_{-0.05}$, with the remaining parameter gradients consistent to within $\pm 1\sigma$.

Crucially, both these fits perform equally as well, mimicking the tight residuals relation ($R_{\text{intr}}^2 \gtrsim 0.95$ and $R_{\text{Bayes}}^2 \approx 0.42$ for each) while explaining the large diversity seen in observed $\xi_{\text{ion},0}$ (both with $\sigma(\bar{\mu}) \approx 0.47$ dex and $\sigma_{\text{int}} \approx 0.04$ dex and ≈ 0.06 dex for $W_{\lambda}(\text{H}\alpha)$ or $W_{\lambda}([\text{O III}])$, respectively).

Notably, these slopes are higher than in both their individual $\xi_{\text{ion},0} - W_{\lambda}(X)$ relations (e.g., see Figures 7 and 9). The elevated conditional slopes measure the change in $\xi_{\text{ion},0}$ with $W_{\lambda}(X)$, while holding the

other properties fixed, which is particularly impactful for properties with moderate to strong covariance such as β_{UV} and $E(B - V)_{\text{neb}}$. As such the steeper relations found in other literature (e.g., Tang et al. 2019 in $\xi_{\text{ion},0} - W_{\lambda}([\text{O III}])$ or Prieto-Lyon et al. 2023 in $\xi_{\text{ion},0} - W_{\lambda}(\text{H}\alpha)$), can be understood in terms of the level of diversity in the physical properties across the samples used. For example, the sample in Prieto-Lyon et al. (2023) (see also Papovich et al. 2025) is much bluer on average ($\langle \beta_{\text{UV}} \rangle = -2.1$ compared with $\langle \beta_{\text{UV}} \rangle = -1.6$ here, see also their Fig. 3), and given the correlation between $W_{\lambda}(\text{H}\alpha)$ and β_{UV} , this will drive a change in the inferred $\xi_{\text{ion},0} - W_{\lambda}(\text{H}\alpha)$ slope.

Similarly, the sources in the [Tang et al. \(2019\)](#) sample are selected as EELGs, which are an extreme subset of the SFG population, likely with more extreme ionization conditions and elevated specific star-formation rates (as well as differences in other physical properties not considered, such as metallicity).

Taking this into account, from the original set of key properties considered in this work, we conclude that the most important in predicting $\xi_{\text{ion},0}$ while striving to minimise the degree of systematic scatter are; either $W_\lambda(\text{H}\alpha)$ or $W_\lambda([\text{O III}])$, the UV continuum slope β_{UV} and the nebular attenuation, $E(B - V)_{\text{neb}}$. These results support the expected physical connections between the properties and the ionizing photon output. The nebular emission lines ($W_\lambda(\text{H}\alpha)$ / $W_\lambda([\text{O III}])$) trace the presence of recently formed massive stars (i.e., recent bursts and high sSFR, as well as responding to metallicity and ionization state), while bluer β_{UV} are indicative of lower metallicity, younger, and less dust attenuated systems. Lastly, $E(B - V)_{\text{neb}}$ dictates the Balmer line attenuation correction, required to accurately infer the intrinsic $\xi_{\text{ion},0}$.

The nebular dust attenuation in and of itself is not a strong predictor of $\xi_{\text{ion},0}$ (the residual scatter is large, $\sigma_{\text{res}} \approx 0.38$ dex, and qualitatively seen in Fig. 6), but it is crucial to explain the outliers (and the heavy-tailed residuals). From the single variable regressions (see Section 4.1) we see only weak-to-moderate behaviour at $E(B - V) \lesssim 0.5$, with much larger excursions coming from the high- $E(B - V)$ tail, suggesting $E(B - V)_{\text{neb}}$ carries the bulk of the $\approx 0.20 - 0.25$ dex of intrinsic scatter that remains when it is omitted from the fitting. This is demonstrable by contrasting the regression results in a two-parameter fit with β_{UV} and $W_\lambda([\text{O III}])$ against a regression additionally including $E(B - V)_{\text{neb}}$. The former gives median posterior parameters of $\beta_{W_\lambda([\text{O III}])} = 0.26 \pm 0.08$, $\beta_{\beta_{\text{UV}}} = -0.13 \pm 0.06$, $\beta_0 = 24.58 \pm 0.21$ (with $\sigma_{\text{int}} = 0.19 \pm 0.04$, $\nu = 2.65^{+0.98}_{-0.46}$), while additionally including $E(B - V)_{\text{neb}}$ produces: $\beta_{W_\lambda([\text{O III}])} = 0.55 \pm 0.05$, $\beta_{\beta_{\text{UV}}} = -0.25 \pm 0.03$, $\beta_{E(B - V)_{\text{neb}}} = 1.73 \pm 0.10$, $\beta_0 = 24.58 \pm 0.21$ (with $\sigma_{\text{int}} = 0.05 \pm 0.02$, $\nu = 25^{+33}_{-16}$). These results clearly suggest that accounting for nebular dust attenuation is crucial for breaking degeneracies as a result of the covariance between observed properties (e.g., $W_\lambda(\text{O III})$ and β_{UV}).

Moreover, the model including only $W_\lambda(\text{O III})$ and β_{UV} gives $R^2_{\text{intr}} \approx 0.47$, $R^2_{\text{Bayes}} \approx 0.13$, $\sigma_{\text{int}} \approx 0.19$ and $\sigma(\bar{\mu}) \approx 0.19$ dex, whereas adding $E(B - V)_{\text{neb}}$ gives $R^2_{\text{intrinsic}} \approx 0.99$, $R^2_{\text{Bayes}} \approx 0.57$, $\sigma_{\text{int}} \approx 0.05$ and $\sigma(\bar{\mu}) \approx 0.46$ dex. Significantly, including $E(B - V)_{\text{neb}}$ removes the heavy-tailed residuals (the StudentT tail parameter rises from $\nu \approx 3$ as shown in Fig. 8 to $\nu \gg 10$), indicating that the sources that appeared to be outliers are driven by dust extremes that the ($W_\lambda(\text{O III})$, β_{UV} , $E(B - V)_{\text{neb}}$) model correctly accounts for.

As a useful exercise, we assess the predictive power when using the most commonly measured properties in samples at the highest redshifts. Specifically, M_{UV} , β_{UV} and redshift can be measured from photometric samples of SFGs selected in studies estimating the UV luminosity function for example (e.g., see [McLeod et al. 2024](#); [Cullen et al. 2024](#)). Such a model (with posterior parameters; $\beta_{\beta_{\text{UV}}} = -0.24 \pm 0.06$, $\beta_{M_{\text{UV}}} = 0.09 \pm 0.04$, $\beta_0(\beta_{\text{UV}} = -2, M_{\text{UV}} = -20) = 25.22 \pm 0.04$, $\sigma_{\text{int}} = 0.23^{+0.05}_{-0.04}$, and $\nu = 3.4^{+2.7}_{-1.0}$) is only able to predict $\xi_{\text{ion},0}$ at the ≈ 0.25 dex level (from the latent model parameters, or worse at RMAE ≈ 0.35 dex when folding in intrinsic and measurement scatter).

A number of past studies have sought to infer the ionizing photon production efficiencies of SFGs from photometric samples (e.g., see [Endsley et al. 2024](#); [Simmonds et al. 2024a](#); [Begley et al. 2025](#)), either applying empirical relations (i.e., like those presented in

Section 4.1.1) or from SED modelling codes. Such methods rely on the flux excess emission line signature in a photometric band, and are therefore not able to directly recover the individual line measurements that may be present (e.g., the F410M band is sensitive to $[\text{O III}] + \text{H}\beta$ at $z \sim 7$ [Endsley et al. 2024](#); [Begley et al. 2025](#)). To enable future comparisons with photometric-based samples, or to easier facilitate predictions of $\xi_{\text{ion},0}$ from photometric flux excess measurements, we perform our regression to the combined nebular $[\text{O III}]\lambda\lambda 4960, 5007 + \text{H}\beta$ equivalent widths. The median posterior parameters are $\beta_1 \approx 0.33^{+0.08}_{-0.09}$, $\beta_0 = 24.30^{+0.24}_{-0.21}$, $\sigma_{\text{int}} = 0.20^{+0.04}_{-0.03}$ and $\nu = 2.5^{+0.8}_{-0.4}$. Predictions based purely on the $[\text{O III}]\lambda\lambda 4960, 5007 + \text{H}\beta$ equivalent width perform marginally better than the above model (as a function of β_{UV} and M_{UV}), achieving RMAE ≈ 0.21 dex from the model parameters (and ≈ 0.33 when folding in the additional scatter). Similarly to other commonly used empirical relations for $\xi_{\text{ion},0}$ (e.g., [Tang et al. 2019](#)), galaxies with high equivalent widths ($W_\lambda([\text{O III}]\lambda\lambda 4960, 5007 + \text{H}\beta) \gtrsim 500 \text{ \AA}$) display $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{ Hz}) \gtrsim 25.2$, suggesting that they are significant contributors to the ionizing photon budget ([Robertson et al. 2015](#)).

For a more direct comparison with the commonly used [Tang et al. \(2019\)](#) relation, we repeat our regression isolating the “extreme emission line” galaxies in the EXCELS sample with $W_\lambda([\text{O III}]\lambda\lambda 4960, 5007 + \text{H}\beta) \gtrsim 500 \text{ \AA}$. Here we find a steeper relation than found across the full EXCELS sample but still shallower than found by [Tang et al. \(2019\)](#), with $\beta_1 = 0.48 \pm 0.13$ (plus $\beta_0 = 23.88 \pm 0.36$, $\sigma_{\text{int}} = 0.17 \pm 0.03$, and $\nu \approx 20^{+31}_{-14}$). The model on the EELG subset returns RMAE ≈ 0.15 dex (and ≈ 0.24 including the measurement scatter), suggesting these extreme galaxies represent a more homogeneous subset of the wider population. The steeper trend with $\xi_{\text{ion},0}$ is also likely driven in part by the fact these sources are typically less dust attenuated, brighter, and have bluer UV continuum slopes on average compared with the full EXCELS sample (e.g., see Fig. 3).

4.2.3 The impact of outliers

From inspecting Fig. 6, the panels using $E(B - V)_{\text{neb}}$ as an auxiliary variable reveal that the majority of outliers in the relations between $\xi_{\text{ion},0}$ and $W_\lambda(\text{H}\alpha)$, $W_\lambda([\text{O III}])$ and β_{UV} have elevated nebular attenuation, $E(B - V)_{\text{neb}} \gtrsim 0.5$. To test the impact of these outliers on the regression, we repeat the $\xi_{\text{ion},0} - W_\lambda(\text{H}\alpha)$ fit after removing the upper quartile of nebular attenuation from the sample $E(B - V)_{\text{neb}} \geq 0.4$. The $\xi_{\text{ion},0} - W_\lambda(\text{H}\alpha)$ slope is marginally steeper $\beta_1 \approx 0.55 \pm 0.09$, but consistent within the uncertainties. However, removing the high $E(B - V)_{\text{neb}}$ tail does in-fact lower the residual scatter, with $\sigma_{\text{int}} \approx 0.18 \pm 0.03$ and $\nu \sim 21$ indicating a return to a more typical Gaussian residuals distribution.

4.2.4 Redshift evolution in regression gradients

A further aspect to consider is the possibility of an underlying evolutionary factor as a function of redshift when using nebular emission lines as a tracer of $\xi_{\text{ion},0}$. Such a change could be driven by the global trend of higher ionization state, lower metallicity or increased star-formation burstiness at higher redshifts. To examine this possibility, we construct three sub-samples (of approximately equal size; $N \approx 35 - 40$ per bin) in redshift having $z < 3$, $3 \leq z < 4$ and $z \geq 4$. Fitting for $\xi_{\text{ion},0} - W_\lambda(\text{H}\alpha)$, we find no clear trend across the three redshift samples. In contrast,

we find a positive (albeit not statistically significant) trend in redshift for the slope strength with the $\xi_{\text{ion},0} - W_{\lambda}([\text{O III}])$ regressions: $\beta_{W_{\lambda}([\text{O III}])} = 0.14 \pm 0.21$ ($\beta_0 = 24.87 \pm 0.46$) at $z < 3$, $\beta_{W_{\lambda}([\text{O III}])} = 0.27 \pm 0.14$ ($\beta_0 = 24.47 \pm 0.35$) at $3 \leq z < 4$, and $\beta_{W_{\lambda}([\text{O III}])} = 0.42 \pm 0.16$ ($\beta_0 = 24.14 \pm 0.43$) at $z \geq 4$. This is physically plausible as $[\text{O III}]$ is more sensitive to the ionization state (i.e., $\log(U)$ and hardness) and metallicity than $\text{H}\alpha$, which more directly tracks the ionizing photon production rate. It therefore follows, given the modest evolution in redshift with metallicity (e.g., Sanders et al. 2016; Cullen et al. 2020; Reddy et al. 2021), that a redshift dependent trend in $\xi_{\text{ion},0} - W_{\lambda}([\text{O III}])$ exists without an associated $\text{H}\alpha$ evolution.

Lastly, we also detect a systemic decline in the intrinsic scatter (for $W_{\lambda}([\text{O III}])$) with redshift; $\sigma_{\text{int}} = 0.41^{+0.14}_{-0.12}$, 0.29 ± 0.06 and $0.20^{+0.06}_{-0.05}$ in our progressively higher redshift bins. At earlier epochs, star-forming galaxies have lower metallicities and have more uniformly young stellar populations, potentially leading to a tighter mapping between $[\text{O III}]$ and ionizing photon production. Observational sample effects could also be amplifying this trend to some extent. Specifically, it is likely that we probe a narrower dynamic range in SFG properties in our higher redshift sample compared to those at the lowest redshifts in our sample (e.g., see also Fig. 3). We stress that larger sample sizes across a large dynamic range in redshift will be required to solidify these findings, as well as robustly determine any evolutionary trends in $\xi_{\text{ion},0} - W_{\lambda}([\text{O III}])$.

4.2.5 The influence of metallicity

A full investigation into the role of metallicity in the ionizing photon production efficiency of galaxies is beyond the scope of this paper. However, for a small subset ($N = 56$, ≈ 35 per cent of our EXCELS sample), we have $[\text{O II}]\lambda\lambda 3727, 3730$ flux measurements which can be combined with $[\text{O III}]\lambda\lambda 4960, 5008$ to measure $\text{O}_{32} (= [\text{O III}]\lambda\lambda 4960, 5008 / [\text{O II}]\lambda\lambda 3727, 3730)$; after dust corrections following the method outlined in Section 2.5.1) as a proxy of metallicity (e.g., see Curti et al. 2023; Hirschmann et al. 2023; Scholte et al. 2025; Shapley et al. 2025).

In Fig. 14, we show $\xi_{\text{ion},0}$ as a function of O_{32} , with the median posterior relation from the Bayesian linear regression: $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) = (0.19 \pm 0.11) \times \log_{10}(\text{O}_{32}) + (25.05 \pm 0.07)$. The gradient is only marginally significant (at the $\approx 2\sigma$ level), as expected given the relatively small sample size and the moderate correlation coefficient of $\rho_S \approx 0.38$ ($p \approx 3 \times 10^{-3}$). This is shallower than the relation found by Llerena et al. (2025) (see also Papovich et al. 2025), but shows excellent consistency with Pahl et al. (2025b).

Broadly speaking, there is a clear qualitative agreement across the literature that higher O_{32} ratios, which trace lower metallicities and higher ionization parameters, are indicative of higher ionizing photon production efficiencies. We do note that there is moderately large intrinsic scatter $\sigma_{\text{int}} = 0.23 \pm 0.03$ (and $\nu \approx 27^{+33}_{-16}$) with low predictive power ($R^2_{\text{intr}} \approx 0.1$, $\sigma(\mu_{\theta}) \approx 0.08$ dex, $\text{RMAE} \approx 0.28$), and that larger deep spectroscopic samples will be required to obtain a more robust relation between metallicity and $\xi_{\text{ion},0}$ while controlling for other physical properties.

5 CONCLUSIONS

In this study we have used ultra-deep JWST/NIRSpec spectroscopy taken as part of the EXCELS survey (e.g., see Fig. 1) to investigate the ionizing photon production efficiencies of $N = 159$ star-forming galaxies across a wide $1 \lesssim z \lesssim 8$ redshift range. With the aim of




Figure 14. The ionizing-photon production efficiency $\xi_{\text{ion},0}$ as a function of $\text{O}_{32} = [\text{O III}]\lambda\lambda 4960, 5008 / [\text{O II}]\lambda\lambda 3727, 3730$, a tracer of ionization parameter and gas-phase metallicity. From a Bayesian linear regression to the $N = 56$ galaxy subsample with O_{32} measurements we obtain $\log_{10}(\xi_{\text{ion},0} / \text{erg}^{-1} \text{Hz}) = (0.19 \pm 0.11) \log_{10}(\text{O}_{32}) + (25.05 \pm 0.07)$ (the red solid line and shaded region indicating the $\pm 1\sigma$ posterior uncertainty). Dashed lines show the inferred intrinsic scatter $\sigma_{\text{int}} = 0.23 \pm 0.03$ dex. This trend is in good qualitative agreement with recent literature (e.g. Llerena et al. 2023; Pahl et al. 2025b; Papovich et al. 2025), and is consistent with higher O_{32} (i.e., higher ionization parameter / lower metallicity) being associated with elevated $\xi_{\text{ion},0}$.

statistically characterising how $\xi_{\text{ion},0}$ connects to key physical properties; including the $[\text{O III}]$ and $\text{H}\alpha$ nebular emission line equivalent widths, the UV luminosity (M_{UV}) and UV spectral slope (β_{UV}), as well as the dust attenuation ($E(B - V)_{\text{neb}}$) and redshift; we adopt a Bayesian methodology to perform single and multi-variate linear regression between these and $\xi_{\text{ion},0}$ (as summarised in Table A1). The key results of our analysis can be summarised as follows:

(i) The ionizing photon production efficiency is most tightly linked with signatures of recent massive-star formation, specifically the equivalent widths of strong optical nebular emission lines $\text{H}\alpha$ and $[\text{O III}]$ are the best predictors. From our Bayesian linear regression fitting $\log_{10}(\xi_{\text{ion},0}) = \beta_1 \times \log_{10}(W_{\lambda}(X)) + \beta_0 + \epsilon(\sigma_{\text{int}}, \nu)$, we find: $\beta_1 = 0.49 \pm 0.09$, $\beta_0 = 23.97 \pm 0.23$, $\sigma_{\text{int}} = 0.20 \pm 0.03$ and $\nu = 2.67 \pm 0.75$, for $\text{H}\alpha$ as shown in Fig. 7 and Fig. 8. Similarly, for $[\text{O III}]$ we infer $\beta_1 = 0.30 \pm 0.08$, $\beta_0 = 24.44 \pm 0.21$, $\sigma_{\text{int}} = 0.20 \pm 0.03$ and $\nu = 2.56 \pm 0.63$ (see Section 4.1.2 and Fig. 4.1.2).

(ii) A multivariate model combining either nebular W_{λ} with the UV continuum slope (β_{UV}) and nebular attenuation level ($E(B - V)_{\text{neb}}$) recovers essentially all the of the intrinsic galaxy-to-galaxy variation in $\xi_{\text{ion},0}$, showing that these three physical properties (from the full set studied; e.g., see Fig. 3 and Fig. 6) capture the dominant physics setting the ionizing photon production efficiency.

(iii) Trends with redshift and absolute UV luminosity (M_{UV}) are shallow and marginal across the EXCELS sample, and are best interpreted as due to the underlying evolution of the primary drivers above. Broadly speaking, the literature reports a wide variety of $\xi_{\text{ion},0} - M_{\text{UV}}$ relations with no clear consensus, reflecting that other physical properties (e.g., $W_{\lambda}(\text{H}\alpha)$ or β_{UV}) are indeed more influential on $\xi_{\text{ion},0}$.

(iv) Although $E(B - V)_{\text{neb}}$ is not a strong single predictor of $\xi_{\text{ion},0}$, it is essential for explaining large outliers and removes the heavy tailed residuals in relations with $W_{\lambda}(\text{H}\alpha)$, $W_{\lambda}([\text{O III}])$ or β_{UV} .

that are consistent with intensely star-forming but obscured regions producing high $\xi_{\text{ion},0}$.

(v) The $\xi_{\text{ion},0} - W_{\lambda}([\text{O III}])$ relation shows a marginal steepening and a systematic decrease in intrinsic scatter at higher redshift, while no clear trend is seen for $H\alpha$. This is physically plausible as a result of evolving nebular conditions (lower metallicity, higher / harder ionization states) impacting $W_{\lambda}([\text{O III}])$ more than than $W_{\lambda}(H\alpha)$.

(vi) Sample selection remains an important consideration when predicting $\xi_{\text{ion},0}$ - our diverse sample yields generally flatter $\xi_{\text{ion},0} - W_{\lambda}(X)$ relations than other literature studies with more focused samples, and differences between these samples (e.g., alternative selections, or $\xi_{\text{ion},0}$ measurement methods) can shift normalisations and slopes.

(vii) From a practical perspective, with only a small set of observables, namely $W_{\lambda}(H\alpha)$ or $W_{\lambda}([\text{O III}])$, β_{UV} , and $E(B - V)_{\text{neb}}$ to reduce the impact of dusty outliers, $\xi_{\text{ion},0}$ is predictable across the SFG population at the ≈ 0.15 dex level.

ACKNOWLEDGEMENTS

R. Begley, R. J. McLure, J. S. Dunlop and D.J. McLeod acknowledge the support of the Science and Technology Facilities Council. F. Cullen and T. M. Stanton acknowledge the support from a UKRI Frontier Research Guarantee Grant [grant reference EP/X021025/1]. A. C. Carnall acknowledges support from a UKRI Frontier Research Guarantee Grant [grant reference EP/Y037065/1]. JSD acknowledges the support of the Royal Society via the award of a Royal Society Research Professorship. RSE acknowledges generous financial support from the Peter and Patricia Gruber Foundation.

This work is based in part on observations made with the NASA/ESA/CSA James Webb Space Telescope. The data were obtained from the Mikulski Archive for Space Telescopes 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.

This research made use of Astropy, a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013, 2018), NumPy (Harris et al. 2020) and SciPy (Virtanen et al. 2020b), Matplotlib (Hunter 2007), IPython (Pérez & Granger 2007), PyMC (Abril-Pla et al. 2023, DOI: <https://doi.org/10.5281/zenodo.4603970>), PINGOUIN (Vallat 2018) and NASA's Astrophysics Data System Bibliographic Services.

DATA AVAILABILITY

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission. All raw JWST and HST data products are available via the Mikulski Archive for Space Telescopes <https://mast.stsci.edu/>, and additional data underlying the main figures and tables are available from the corresponding author upon reasonable request.

REFERENCES

Abril-Pla O., et al., 2023, *PeerJ Comput Sci.*
 Adams N. J., et al., 2025, *MNRAS*, **542**, 1705
 Aird J., Coil A. L., Georgakakis A., Nandra K., Barro G., Pérez-González P. G., 2015, *MNRAS*, **451**, 1892
 Arellano-Córdova K. Z., et al., 2024, *arXiv e-prints*, p. [arXiv:2412.10557](https://arxiv.org/abs/2412.10557)

Astropy Collaboration et al., 2013, *A&A*, **558**, A33
 Astropy Collaboration et al., 2018, *AJ*, **156**, 123
 Atek H., Furtak L. J., Oesch P., van Dokkum P., Reddy N., Contini T., Illingworth G., Wilkins S., 2022, *MNRAS*, **511**, 4464
 Atek H., et al., 2023, *arXiv e-prints*, p. [arXiv:2308.08540](https://arxiv.org/abs/2308.08540)
 Begley R., et al., 2022, *MNRAS*, **513**, 3510
 Begley R., et al., 2023, *arXiv e-prints*, p. [arXiv:2306.03916](https://arxiv.org/abs/2306.03916)
 Begley R., et al., 2025, *MNRAS*, **537**, 3245
 Bertin E., Arnouts S., 1996, *A&AS*, **117**, 393
 Bosman S. E. I., et al., 2021, *arXiv e-prints*, p. [arXiv:2108.03699](https://arxiv.org/abs/2108.03699)
 Bouwens R. J., et al., 2014, *ApJ*, **793**, 115
 Bouwens R. J., et al., 2015, *ApJ*, **803**, 34
 Bouwens R. J., Smit R., Labbé I., Franx M., Caruana J., Oesch P., Stefanon M., Rasappu N., 2016, *ApJ*, **831**, 176
 Bowler R. A. A., et al., 2015, *MNRAS*, **452**, 1817
 Branch M. A., Coleman T. F., Li Y., 1999, *SIAM Journal on Scientific Computing*, **21**, 1
 Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *ApJ*, **533**, 682
 Cardelli J. A., Clayton G. C., Mathis J. S., 1989, *ApJ*, **345**, 245
 Carnall A. C., et al., 2024, *arXiv e-prints*, p. [arXiv:2405.02242](https://arxiv.org/abs/2405.02242)
 Carvajal-Bohorquez C., et al., 2025, *arXiv e-prints*, p. [arXiv:2507.13160](https://arxiv.org/abs/2507.13160)
 Charlot S., Fall S. M., 2000, *ApJ*, **539**, 718
 Chary R., Petitjean P., Robertson B., Trenti M., Vangioni E., 2016, *Space Sci. Rev.*, **202**, 181
 Chen N., Motohara K., Spitler L., Nakajima K., Terao Y., 2024, *ApJ*, **968**, 32
 Chisholm J., et al., 2022, *MNRAS*, **517**, 5104
 Choustikov N., et al., 2024, *MNRAS*, **529**, 3751
 Cole J. W., et al., 2025, *ApJ*, **979**, 193
 Conroy C., 2013, *ARA&A*, **51**, 393
 Cullen F., et al., 2019, *MNRAS*, **487**, 2038
 Cullen F., et al., 2020, *MNRAS*, **495**, 1501
 Cullen F., et al., 2024, *MNRAS*, **531**, 997
 Curti M., et al., 2023, *MNRAS*, **518**, 425
 Curti M., et al., 2024, *A&A*, **684**, A75
 Donnan C. T., et al., 2022, *MNRAS*,
 Donnan C. T., Dunlop J. S., McLure R. J., McLeod D. J., Cullen F., 2025, *MNRAS*, **539**, 2409
 Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, *Publ. Astron. Soc. Australia*, **34**, e058
 Emami N., Siana B., Alavi A., Gburek T., Freeman W. R., Richard J., Weisz D. R., Stark D. P., 2020, *ApJ*, **895**, 116
 Endsley R., Stark D. P., Whitler L., Topping M. W., Chen Z., Plat A., Chisholm J., Charlot S., 2022, *arXiv e-prints*, p. [arXiv:2208.14999](https://arxiv.org/abs/2208.14999)
 Endsley R., et al., 2024, *MNRAS*, **533**, 1111
 Endsley R., Chisholm J., Stark D. P., Topping M. W., Whitler L., 2025, *ApJ*, **987**, 189
 Erb D. K., Pettini M., Steidel C. C., Strom A. L., Rudie G. C., Trainor R. F., Shapley A. E., Reddy N. A., 2016, *ApJ*, **830**, 52
 Faist A. L., Capak P. L., Emami N., Tacchella S., Larson K. L., 2019, *ApJ*, **884**, 133
 Finkelstein S. L., et al., 2015, *ApJ*, **810**, 71
 Finkelstein S. L., et al., 2019, *ApJ*, **879**, 36
 Fletcher T. J., Tang M., Robertson B. E., Nakajima K., Ellis R. S., Stark D. P., Inoue A., 2019, *ApJ*, **878**, 87
 Flury S. R., et al., 2022, *ApJ*, **930**, 126
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, **125**, 306
 Fumagalli M., et al., 2012, *ApJ*, **757**, L22
 Garilli B., et al., 2021, *A&A*, **647**, A150
 Gelman A., Goodrich B., Gabry J., Vehtari A., 2019, *The American Statistician*, **73**, 307
 Gordon K. D., Clayton G. C., Misselt K. A., Landolt A. U., Wolff M. J., 2003, *ApJ*, **594**, 279
 Goto H., et al., 2021, *ApJ*, **923**, 229
 Grogin N. A., et al., 2011, *ApJS*, **197**, 35
 Harris C. R., et al., 2020, *Nature*, **585**, 357

- Hayes M. J., Saldana-Lopez A., Citro A., James B. L., Mingozi M., Scarlata C., Martinez Z., Berg D. A., 2025, *ApJ*, **982**, 14
- Hirschmann M., Charlot S., Somerville R. S., 2023, *MNRAS*, **526**, 3504
- Hoffman M. D., Gelman A., 2011, *arXiv e-prints*, p. [arXiv:1111.4246](#)
- Horne K., 1986, *PASP*, **98**, 609
- Hunter J. D., 2007, *Computing in Science & Engineering*, **9**, 90
- Inoue A. K., Iwata I., 2008, *MNRAS*, **387**, 1681
- Inoue A. K., Shimizu I., Iwata I., Tanaka M., 2014, *MNRAS*, **442**, 1805
- Izotov Y. I., Worseck G., Schaerer D., Guseva N. G., Thuan T. X., Fricke Verhamme A., Orlitová I., 2018, *MNRAS*, **478**, 4851
- Jain S., et al., 2025, *arXiv e-prints*, p. [arXiv:2508.18369](#)
- Jaskot A. E., 2025, *ARA&A*, **63**, 45
- Jaskot A. E., Dowd T., Oey M. S., Scarlata C., McKinney J., 2019, *ApJ*, **885**, 96
- Jaskot A. E., et al., 2024, *ApJ*, **973**, 111
- Jung I., et al., 2023, *arXiv e-prints*, p. [arXiv:2304.05385](#)
- Katz H., et al., 2023, *The Open Journal of Astrophysics*, **6**, 44
- Khostovan A. A., Malhotra S., Rhoads J. E., Sobral D., Harish S., Tilvi V., Coughlin A., Rezaee S., 2024, *MNRAS*, **535**, 2903
- Kimm T., Bieri R., Geen S., Rosdahl J., Blaizot J., Michel-Dansac L., Garel T., 2022, *ApJS*, **259**, 21
- Koekemoer A. M., et al., 2011, *ApJS*, **197**, 36
- Kriek M., Conroy C., 2013, *ApJ*, **775**, L16
- Kulkarni G., Worseck G., Hennawi J. F., 2019, *MNRAS*, **488**, 1035
- Kumar R., Carroll C., Hartikainen A., Martin O., 2019, *Journal of Open Source Software*, **4**, 1143
- Laseter I. H., et al., 2025, *ApJ*, **988**, 73
- Le Fèvre O., et al., 2005, *A&A*, **439**, 845
- Llerena M., et al., 2023, *A&A*, **676**, A53
- Llerena M., et al., 2025, *A&A*, **698**, A302
- Looser T. J., et al., 2023, *arXiv e-prints*, p. [arXiv:2302.14155](#)
- Lovell C. C., Vijayan A. P., Thomas P. A., Wilkins S. M., Barnes D. J., Irodoutou D., Roper W., 2021, *MNRAS*, **500**, 2127
- Luridiana V., Morisset C., Shaw R. A., 2015, *A&A*, **573**, A42
- Ma X., Quataert E., Wetzel A., Hopkins P. F., Faucher-Giguère C.-A., Kereš D., 2020, *MNRAS*, **498**, 2001
- Maji M., et al., 2022, *A&A*, **663**, A66
- Marchi F., et al., 2018, *A&A*, **614**, A11
- Mármol-Queraltó E., McLure R. J., Cullen F., Dunlop J. S., Fontana A., McLeod D. J., 2016, *MNRAS*, **460**, 3587
- Mascia S., et al., 2023, *A&A*, **672**, A155
- Mascia S., et al., 2024, *A&A*, **685**, A3
- Maseda M. V., et al., 2023, *ApJ*, **956**, 11
- Mason C. A., Treu T., Dijkstra M., Mesinger A., Trenti M., Pentericci L., de Barros S., Vanzella E., 2018, *ApJ*, **856**, 2
- Matsuoka Y., et al., 2023, *ApJ*, **949**, L42
- Mauerhofer V., Dayal P., Haehnelt M. G., Kimm T., Rosdahl J., Teyssier R., 2025, *A&A*, **696**, A157
- Maulick S., Saha K., Rutkowski M. J., 2025, *ApJ*, **984**, 40
- McClymont W., et al., 2025, *MNRAS*, **540**, 190
- McLeod D. J., McLure R. J., Dunlop J. S., Cullen F., Carnall A. C., Duncan K., 2021, *MNRAS*, **503**, 4413
- McLeod D. J., et al., 2024, *MNRAS*, **527**, 5004
- McLure R. J., et al., 2013, *MNRAS*, **432**, 2696
- McLure R. J., et al., 2018a, *MNRAS*, **476**, 3991
- McLure R. J., et al., 2018b, *MNRAS*, **479**, 25
- Mingozi M., et al., 2022, *ApJ*, **939**, 110
- Morisset C., Luridiana V., García-Rojas J., Gómez-Llanos V., Bautista M., Mendoza C., 2020, *Atoms*, **8**, 66
- Muñoz J. B., Mirocha J., Chisholm J., Furlanetto S. R., Mason C., 2024, *MNRAS*, **535**, 2903
- Naidu R. P., Tacchella S., Mason C. A., Bose S., Oesch P. A., Conroy C., 2020, *ApJ*, **892**, 109
- Nakajima K., Ellis R. S., Iwata I., Inoue A. K., Kusakabe H., Ouchi M., Robertson B. E., 2016, *ApJ*, **831**, L9
- Oke J. B., Gunn J. E., 1983, *ApJ*, **266**, 713
- Osterbrock D. E., Ferland G. J., 2006, *Astrophysics of gaseous nebulae and active galactic nuclei*
- Pahl A. J., Shapley A., Steidel C. C., Chen Y., Reddy N. A., 2021, *MNRAS*, **505**, 2447
- Pahl A. J., Shapley A., Steidel C. C., Reddy N. A., Chen Y., Rudie G. C., Strom A. L., 2023, *MNRAS*, **521**, 3247
- Pahl A. J., et al., 2025a, *arXiv e-prints*, p. [arXiv:2510.12177](#)
- Pahl A. J., et al., 2025b, *ApJ*, **981**, 134
- Papovich C., et al., 2025, *arXiv e-prints*, p. [arXiv:2505.08870](#)
- Pentericci L., et al., 2018, *A&A*, **616**, A174
- Pérez F., Granger B. E., 2007, *Computing in Science and Engineering*, **9**, 21
- Planck Collaboration et al., 2020, *A&A*, **641**, A6
- Prieto-Lyon G., et al., 2023, *A&A*, **672**, A186
- Reddy N. A., Steidel C. C., Pettini M., Bogosavljević M., Shapley A. E., 2016, *ApJ*, **828**, 108
- Reddy N. A., et al., 2018a, *ApJ*, **853**, 56
- Reddy N. A., et al., 2018b, *ApJ*, **869**, 92
- Reddy N. A., et al., 2020, *ApJ*, **902**, 123
- Reddy N. A., et al., 2021, *arXiv e-prints*, p. [arXiv:2108.05363](#)
- Reddy N. A., et al., 2025, *arXiv e-prints*, p. [arXiv:2506.17396](#)
- Roberts-Borsani G., et al., 2024, *arXiv e-prints*, p. [arXiv:2403.07103](#)
- Robertson B. E., 2021, *arXiv e-prints*, p. [arXiv:2110.13160](#)
- Robertson B. E., et al., 2013, *ApJ*, **768**, 71
- Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, *ApJ*, **802**, L19
- Robertson B. E., et al., 2023, *Nature Astronomy*, **7**, 611
- Rogers N. S. J., Strom A. L., Rudie G. C., Trainor R. F., Raptis M., von Raesfeld C., 2024, *ApJ*, **964**, L12
- Rosdahl J., et al., 2018, *MNRAS*, **479**, 994
- Saldana-Lopez A., et al., 2023, *MNRAS*, **522**, 6295
- Salim S., Narayanan D., 2020, *ARA&A*, **58**, 529
- Salim S., Boquien M., Lee J. C., 2018, *ApJ*, **859**, 11
- Salvatier J., Wiecki T., Fonnesbeck C., 2015, *arXiv e-prints*, p. [arXiv:1507.08050](#)
- Sanders R. L., et al., 2016, *ApJ*, **816**, 23
- Sanders R. L., et al., 2024, *arXiv e-prints*, p. [arXiv:2408.05273](#)
- Scarlata C., Hayes M., Panagia N., Mehta V., Haardt F., Bagley M., 2024, *arXiv e-prints*, p. [arXiv:2404.09015](#)
- Scholte D., et al., 2025, *MNRAS*, **535**, 2903
- Seeyave L. T. C., et al., 2023, *MNRAS*, **525**, 2422
- Shapley A. E., et al., 2025, *ApJ*, **980**, 242
- Shivaei I., et al., 2018, *ApJ*, **855**, 42
- Simmonds C., et al., 2024a, *arXiv e-prints*, p. [arXiv:2409.01286](#)
- Simmonds C., et al., 2024b, *MNRAS*, **527**, 6139
- Simmonds C., et al., 2025, *arXiv e-prints*, p. [arXiv:2508.04410](#)
- Stanton T. M., et al., 2025, *MNRAS*, **537**, 1735
- Stefanon M., Bouwens R. J., Illingworth G. D., Labbé I., Oesch P. A., Gonzalez V., 2022, *ApJ*, **935**, 94
- Steidel C. C., Bogosavljević M., Shapley A. E., Reddy N. A., Rudie G. C., Pettini M., Trainor R. F., Strom A. L., 2018, *ApJ*, **869**, 123
- Storey P. J., Hummer D. G., 1995, *MNRAS*, **272**, 41
- Stoughton C., et al., 2002, *AJ*, **123**, 485
- Tang M., Stark D. P., Chevallard J., Charlot S., 2019, *MNRAS*, **489**, 2572
- Tang M., Stark D. P., Chevallard J., Charlot S., Endsley R., Congiu E., 2021, *MNRAS*, **503**, 4105
- Tang M., Stark D. P., Mason C. A., Gelli V., Chen Z., Topping M. W., 2025, *arXiv e-prints*, p. [arXiv:2507.08245](#)
- Topping M. W., et al., 2024, *MNRAS*, **529**, 4087
- Trebtsch M., Hutter A., Dayal P., Gottlöber S., Legrand L., Yepes G., 2023, *MNRAS*, **518**, 3576
- Trussler J. A. A., et al., 2025, *MNRAS*, **537**, 3662
- Vallat R., 2018, *Journal of Open Source Software*, **3**, 1026
- Vehtari A., Gelman A., Gabry J., 2015, *arXiv e-prints*, p. [arXiv:1507.04544](#)
- Vijayan A. P., Thomas P. A., Lovell C. C., Wilkins S. M., Greve T. R., Irodoutou D., Roper W. J., Seeyave L. T. C., 2024, *MNRAS*, **527**, 7337
- Virtanen P., et al., 2020a, *Nature Methods*, **17**, 261
- Virtanen P., et al., 2020b, *Nature Methods*, **17**, 261
- Whitler L., et al., 2025, *arXiv e-prints*, p. [arXiv:2501.00984](#)
- Yanagisawa H., et al., 2024, *ApJ*, **974**, 180




Figure A1. The $H\alpha$ luminosity versus the $H\beta$ luminosity, scaled assuming Case B recombination with $T_e = 12,000$ K and $n_e = 100 \text{ cm}^{-3}$, for the $N = 105$ EXCELS galaxies in our sample with both $H\alpha$ and $H\beta$ measured at $S/N \geq 3$. Both measurements are corrected for dust attenuation using the observed Balmer decrement, as outlined in Section 2.5.1. The inset shows the ratio distribution, with $\bar{\mu} = 0.969$ and $\sigma_{\bar{\mu}} = 0.014$, highlighting the excellent agreement after accounting for the differential dust attenuation under the adopted Case B assumption. This is also illustrated by the tight clustering around the (black dashed) 1 : 1 line. We also indicate the typical $x - y$ error with the marker plotted in the top right.

APPENDIX A: SUPPLEMENTARY FIGURES

A1 Summary of Bayesian regression models

In Table A1, we summarise the regression coefficients $\theta = \{\beta_X, \dots, \alpha, \sigma_{\text{int}}, \nu\}$ from the fits throughout Section 4, that may be used for predictions of $\xi_{\text{ion},0}$ from a variety of different measurement types (both photometric and spectroscopic).

A2 Nebular dust attenuation corrections

As verification for the fiducial choices of Case B recombination parameters, as well as the adopted nebular dust attenuation law (e.g., see Section 2.5.1), in Fig. A1 we compare the intrinsic $H\alpha$ luminosity to the scaled intrinsic $H\beta$ luminosity. The $H\beta$ luminosity has been scaled by a factor = 2.826 according to case B recombination ($T_e = 12,000$ K and $n_e = 100 \text{ cm}^{-3}$). The comparison demonstrates an excellent agreement (see inset in Fig. A1) in the intrinsic Balmer line luminosities after corrections, with a median ratio of $\bar{\mu} = 0.969$ and distribution standard deviation of $\sigma_{\bar{\mu}} = 0.014$. Any residual difference remaining is likely due to galaxy-to-galaxy differences in the nebular dust attenuation law (which are not directly measurable in our data, but see Sanders et al. 2024; Reddy et al. 2025, for recent literature on the subject), or in a very small minority of cases, a deviation from Case B recombination conditions (e.g., see Scarlata et al. 2024; McClymont et al. 2025).

Here we also demonstrate the importance of direct spectroscopic measurements of the nebular dust attenuation. A common assumption in the literature, based on observations of local starburst galaxies (e.g., Calzetti et al. 2000, see also; Kriek & Conroy 2013 and Reddy et al. 2020), is that the stellar and nebular dust attenuation levels are related as $E(B - V)_{\text{stellar}}/E(B - V)_{\text{nebular}} = 0.44$, as nebular emission originates from dustier star-forming regions, or “birth-clouds” (Charlot & Fall 2000). Fig. A2 contrasts the stellar and nebular




Figure A2. Stellar versus nebular dust attenuation (parameterised by colour excess, $E(B - V)$), with typical error indicated in the top-right corner for our EXCELS galaxy sample. We show the 1 : 1 line (black dashed) as well as $E(B - V)_{\text{stellar}} = 0.44 \times E(B - V)_{\text{nebular}}$ relation typically adopted in the literature to account for the additional attenuation present in birth clouds (e.g., see Calzetti et al. 2000; Charlot & Fall 2000). Significant galaxy-to-galaxy scatter is seen in the EXCELS, justifying the requirement for spectroscopically measured $E(B - V)_{\text{neb}}$ to robustly infer $\xi_{\text{ion},0}$. In the inset we also show the stellar-mass dependence, whereby higher mass galaxies have larger $E(B - V)_{\text{neb}}/E(B - V)_{\text{stellar}}$ ratios (see also Section A2 for details).

dust attenuation (parameterised via the colour excess; see Section 2.5.1 and Sections 2.5.2 and 2.5.3 for details on $E(B - V)_{\text{neb}}$ and $E(B - V)_{\text{stellar}}$, respectively), revealing extremely large galaxy-to-galaxy scatter with $E(B - V)_{\text{stellar}}/E(B - V)_{\text{nebular}}$ ratios ranging from $\lesssim 0.1$ to $\gtrsim 5$. For reference, we also report a median ratio of $\langle E(B - V)_{\text{stellar}}/E(B - V)_{\text{nebular}} \rangle \approx 0.45$, and a $\pm 1\sigma$ (16th – 84th percentile) range of $\approx 0.1 - 1.1$.

From Fig. A2 it is clear that assuming a constant nebular-to-stellar dust attenuation ratio is a largely oversimplifying assumption, and may lead to large uncertainties in $\xi_{\text{ion},0}$. This is exacerbated by the underlying stellar mass dependence (with mass inferred from SED fits to emission line corrected photometry; see Stanton et al. 2025b in preparation for method details) observed in the EXCELS sample, whereby higher mass galaxies have higher $E(B - V)_{\text{neb}}/E(B - V)_{\text{stellar}}$ (see also Mingozi et al. 2022, for similar findings at $z \sim 0$).

A3 Individual galaxy constraints on $\xi_{\text{ion},0}$

To accurately propagate the full measurement uncertainties through to our $\xi_{\text{ion},0}$ constraints (e.g., see Section 2.5.3) and ultimately to our Bayesian linear regression (see Section 3.2), we have adopted a Monte Carlo approach as detailed throughout Section 2.5. This method allows us to maintain more accurate constraints than offered by the standard Normally distribution measurement errors (i.e., $x \pm \sigma_x$) as shown in Fig. A3. From the figure, we can see that not only do we have a large range in $\xi_{\text{ion},0}$ across the EXCELS sample, but there is also significant diversity in the $\xi_{\text{ion},0}$ constraints as evident from their distributions. Failing to consider this diversity would inevitably lead to underestimations of the model parameters at best, or strong systematic biases at worst.

Table A1. Coefficients from the Bayesian linear regression (e.g., see Section 3.2), for the various models described throughout Section 4.

Model Parameters	Model Coefficients							
$W_\lambda(\text{H}\alpha)$	0.49 ± 0.09	—	—	—	—	—	0.65 ± 0.06	—
$W_\lambda([\text{O III}])$	—	0.30 ± 0.08	—	0.26 ± 0.08	0.55 ± 0.05	—	—	0.51 ± 0.05
$\beta_{\text{UV}} + 2$	—	—	-0.22 ± 0.06	-0.13 ± 0.06	-0.25 ± 0.03	-0.24 ± 0.06	-0.33 ± 0.03	-0.30 ± 0.04
$M_{\text{UV}} + 20$	—	—	—	—	—	0.09 ± 0.04	0.05 ± 0.03	0.04 ± 0.03
$E(B - V)_{\text{neb}}$	—	—	—	—	1.73 ± 0.10	—	1.58 ± 0.09	1.74 ± 0.10
z_{spec}	—	—	—	—	—	—	0.01 ± 0.02	0.01 ± 0.02
α	23.97 ± 0.23	24.44 ± 0.21	25.25 ± 0.04	24.58 ± 0.21	24.58 ± 0.21	25.22 ± 0.04	$23.39^{+0.15}_{-0.14}$	23.71 ± 0.13
σ_{int}	0.20 ± 0.03	0.20 ± 0.03	0.23 ± 0.04	0.19 ± 0.04	0.05 ± 0.02	$0.23^{+0.05}_{-0.04}$	0.04 ± 0.02	0.06 ± 0.02
ν	2.67 ± 0.75	2.56 ± 0.63	3.22 ± 1.46	$2.65^{+0.98}_{-0.46}$	25^{+33}_{-16}	$3.4^{+2.7}_{-1.0}$	$24.2^{+32.7}_{-16.0}$	25^{+32}_{-16}




Figure A3. Example $\xi_{\text{ion},0}$ constraints obtained from the EXCELS spectroscopy for six sources. We show the kernel density estimates of the distributions generated from the Monte Carlo error uncertainty method used in Section 2.5.3, along with their median and 16th – 84th percentiles as dashed lines. A key observation is that the galaxy-to-galaxy distributions are diverse, and as such the typical usage of a measurement and error as $x \pm \sigma_x$ (i.e., assuming Normally distributed errors, here as $\log_{10}(\xi_{\text{ion},0}) \sim \mathcal{N}(x, \sigma_x^2)$) would generally fail to accurately represent the constraints. This then motivates the strategy taken in our Bayesian linear regression described in Section 3.2, whereby we marginalise over the $\xi_{\text{ion},0}$ distributions directly.

A4 Posterior distributions from multivariate regression

In Fig. A4, we show the corner plot of the multivariate regression that includes the six principal observed properties, as discussed in Section 4.2.1.

A5 Bayesian Linear Regression Diagnostics

We assess the convergence of the PyMC sampling using diagnostic trace plots for each model fit, for example shown in Fig. A5 for the single-variate $\xi_{\text{ion},0} - W_\lambda(\text{H}\alpha)$ fit (see also summary statistics discussed in Section 3.2; i.e., $\hat{R}$ and effective sample size, ESS). The figure shows the four independent NUTS chain samples as well as the corresponding posterior density for each fitted parameter (in this case $\theta = \{\beta_1, \alpha, \sigma_{\text{int}}, \nu\}$). Well-behaved sampling is typically represented convergence of the different chains, in addition to the same overall shape in each panel's posterior density, which indicate the sampler has explored the posterior reliably. If a parameter shows poor mixing or unusual excursions we treat it as a warning and re-run with the necessary adjustments (more tuning, a different parameterisation for that block of the model, or slightly stronger regularising priors) until the traces appear well mixed and stable. In the figure panels, we also report key diagnostic statistics including; $\hat{R}$, the Monte Carlo standard errors $\text{MCSE}_{\text{mean}}$ and MCSE_{SD} , and the bulk and tail effective sample sizes, $N_{\text{ESS}}(\text{Bulk})$ and $N_{\text{ESS}}(\text{Tail})$ (see PyMC documentation for details). These statistics further aide interpreting the quality of the model sampling; $\hat{R}$ near unity ($\lesssim 1.01$) indicates that the chains have converged to the same distribution; $\text{MCSE}_{\text{mean}}$ and MCSE_{SD} are the Monte-Carlo uncertainties on the posterior mean and standard deviation (they should be very small compared with the posterior width); and the N_{ESS} (bulk/tail) values report the effective number of independent samples for the central region and the extremes of the posterior (larger is better, with $\gtrsim 500$ typically being sufficient for accurate posteriors).

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




Figure A4. Corner plot for the multivariate Bayesian linear regression between $\xi_{\text{ion},0}$ and the six key observable properties $X_n = \{W_\lambda(H\alpha), W_\lambda([OIII]), \beta_{UV}, M_{UV}, E(B-V)_{\text{neb}}, z_{\text{spec}}\}$. The model has 10 parameters, including 6 gradients (β_{X_i}), the intercept β_0 , the power-law index for the assumed $g(E(B-V)_{\text{neb}})$ functional form, and the intrinsic scatter parameters (σ_{int} and ν).




Figure A5. Diagnostic trace plot for the PyMC Bayesian linear regression methodology introduced in Section 3.2. Each panel on the right shows the chain traces, with the marginal posterior distributions per chain on the left, for the four key parameters (in a single-predictor model, here for $W_\lambda(\text{H}\alpha)$); $\theta = \{\beta_1, \alpha, \sigma_{\text{int}}, \nu\}$, that is the regression slope and intercept, as well as the intrinsic relation scatter and “tail” parameter for a StudentT distribution. The initialisation samples ($N = 1000$) are not shown, and we mark the median posterior parameter value across all chains as a grey line (vertically on the right panels, horizontally on the left panels). We also report a number of key diagnostics, as discussed in the text of Section A5.