

A super-linear ‘radio-AGN main sequence’ links mean radio-AGN power and galaxy stellar mass since $z \sim 3$

I. Delvecchio¹, E. Daddi², M. T. Sargent^{3,4}, J. Aird^{5,6}, J. R. Mullaney⁷, B. Magnelli², D. Elbaz², L. Bisigello^{8,9,10}, L. Ceraj¹¹, S. Jin^{12,13}, B. S. Kalita², D. Liu¹⁴, M. Novak¹⁵, I. Prandoni¹⁶, J. F. Radcliffe^{17,18,19}, C. Spingola¹⁶, G. Zamorani¹⁰, V. Allevato²⁰, G. Rodighiero^{8,9}, and V. Smolčić²¹

¹ INAF – Osservatorio Astronomico di Brera, Via Brera 28, 20121, Milano, Italy and Via Bianchi 46, 23807 Merate, Italy
e-mail: ivan.delvecchio@inaf.it

² Université Paris-Saclay, Université Paris-Cité, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France

³ International Space Science Institute (ISSI), Hallerstrasse 6, 3012 Bern, Switzerland

⁴ Astronomy Centre, Department of Physics & Astronomy, University of Sussex, Brighton BN1 9QH, UK

⁵ Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK

⁶ School of Physics and Astronomy, University of Leicester, University Road, Leicester LE1 7RJ, UK

⁷ Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH, UK

⁸ Department of Physics and Astronomy, Università degli Studi di Padova, Vicolo dell’Osservatorio 3, 35122 Padova, Italy

⁹ INAF – Osservatorio Astronomico di Padova, Vicolo dell’Osservatorio 5, 35122 Padova, Italy

¹⁰ INAF – Osservatorio Astronomico di Bologna, Via P. Gobetti 93/3, 40129 Bologna, Italy

¹¹ Ruđer Bošković Institute, Bijenička cesta 54, 10000 Zagreb, Croatia

¹² Cosmic Dawn Center (DAWN), Rådmandsgade 62, 2200 Copenhagen, Denmark

¹³ DTU-Space, Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark

¹⁴ Max-Planck-Institut für extraterrestrische Physik (MPE), Giessenbachstrasse 1, 85748 Garching, Germany

¹⁵ Independent researcher, formerly at Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

¹⁶ INAF – Istituto di Radioastronomia, Via P. Gobetti 101, 40129 Bologna, Italy

¹⁷ Department of Physics, University of Pretoria, Lynnwood Road, Hatfield, Pretoria 0083, South Africa

¹⁸ Jodrell Bank Centre for Astrophysics, University of Manchester, Oxford Road, Manchester M13 9PL, UK

¹⁹ National Institute for Theoretical and Computational Sciences (NITheCS), Matieland, Stellenbosch 7602, South Africa

²⁰ INAF – Osservatorio Astronomico di Capodimonte, Via Moiariello 16, 30131 Naples, Italy

²¹ Department of Physics, University of Zagreb, Bijenička cesta 32, 10002 Zagreb, Croatia

Received 29 July 2022 / Accepted 26 September 2022

ABSTRACT

Mapping the average active galactic nucleus (AGN) luminosity across galaxy populations and over time reveals important clues regarding the interplay between supermassive black hole and galaxy growth. This paper presents the demography, mean power, and cosmic evolution of radio AGN across star-forming galaxies (SFGs) of different stellar masses (M_*). We exploit deep VLA-COSMOS 3 GHz data to build the rest-frame 1.4 GHz AGN luminosity functions at $0.1 \leq z \leq 4.5$ hosted in SFGs. Splitting the AGN luminosity function into different M_* bins reveals that, at all redshifts, radio AGN are both more frequent and more luminous in higher M_* than in lower M_* galaxies. The cumulative kinetic luminosity density exerted by radio AGN in SFGs peaks at $z \sim 2$, and it is mostly driven by galaxies with $10.5 \leq \log(M_*/M_\odot) < 11$. Averaging the cumulative radio AGN activity across all SFGs at each (M_*, z) results in a ‘radio-AGN main sequence’ that links the time-averaged radio-AGN power $\langle L_{1.4}^{\text{AGN}} \rangle$ and galaxy stellar mass, in the form: $\log \langle [L_{1.4}^{\text{AGN}} / \text{W Hz}^{-1}] \rangle = (20.97 \pm 0.16) + (2.51 \pm 0.34) \cdot \log(1+z) + (1.41 \pm 0.09) \cdot (\log[M_*/M_\odot] - 10)$. The super-linear dependence on M_* , at fixed redshift, suggests enhanced radio-AGN activity in more massive SFGs as compared to star formation. We ascribe this enhancement to both a higher radio AGN duty cycle and a brighter radio-AGN phase in more massive SFGs. A remarkably consistent M_* dependence is seen for the evolving X-ray AGN population in SFGs. This similarity is interpreted as possibly driven by secular cold gas accretion fuelling both radio and X-ray AGN activity in a similar fashion over the galaxy’s lifetime.

Key words. galaxies: nuclei – radio continuum: galaxies – galaxies: luminosity function, mass function – galaxies: active – galaxies: evolution

1. Introduction

Understanding what drives the interplay between supermassive black holes (SMBHs) and their host galaxies is among the most debated topics in extragalactic astrophysics. Multi-wavelength surveys in the last decade have enabled us to reconstruct the cosmic history of SMBH accretion and star formation (SF; [Madau & Dickinson 2014](#)), finding a seem-

ingly consistent decline since ‘cosmic noon’ ($z \sim 1-3$, e.g. [Förster Schreiber & Wuyts 2020](#)). Fossil remnants of such interplay, the empirical black hole–galaxy scaling relations seen at $z \sim 0$, might suggest an intertwined evolution ([Kormendy & Ho 2013](#)), which is likely self-regulated by the galaxy’s baryon cycle of feeding and feedback mechanisms (e.g. [Harrison 2017](#)).

To explain black hole–galaxy scaling relations, cosmological simulations advocate a twofold phase of active galactic nucleus

(AGN) feedback, characterised by high radiative (‘quasar mode’) and high kinetic (‘jet mode’) power, whose cumulative effects are able to regulate SF in massive galaxies (stellar mass $M_* > 10^{10} M_\odot$), preventing the runaway galaxy mass growth (Di Matteo et al. 2005; Croton et al. 2006; Zubovas & King 2012). While the scenario of ‘expulsive’ AGN feedback is not fully backed up by observations (e.g. Sanders et al. 2022), alternative manifestations of AGN feedback have recently been gaining consensus. For instance, compact (<1 kpc), low-power radio jets have been widely observed in local (‘radio-quiet’) Seyferts (e.g. Jarvis et al. 2020; Venturi et al. 2021; Girdhar et al. 2022), powering ionised-gas outflows that inject heat and turbulence into the interstellar medium, thus possibly reducing the host’s SF efficiency. Nevertheless, the long-term impact (or lack thereof) of AGN feedback on galaxy growth is still highly controversial (e.g. Harrison 2017).

In the local Universe ($z < 0.3$), a number of studies have reported an increasing incidence of jet-driven AGN activity (often visible at radio frequencies) in more massive galaxies (Heckman & Best 2014), which reaches close to 100% at $M_* > 10^{11.5} M_\odot$ (Sabater et al. 2019), especially if above low-frequency radio luminosities of $\log[L_{150 \text{ MHz}}/W \text{ Hz}^{-1}] \geq 21.7$. Given the known black hole–galaxy mass scaling relations (Kormendy & Ho 2013), this trend might reflect the increasing ability of more massive black holes to launch more powerful jets (e.g. Best et al. 2005).

However, capturing the instantaneous effect of AGN feedback has proven to be difficult. The distribution of the black hole accretion rate (BHAR) normalised by galaxy M_* – or ‘specific BHAR’ (sBHAR $\propto L_X/M_*^1$; Aird et al. 2012) – is notably broad (>1 dex; e.g. Mendez et al. 2013; Azadi et al. 2015; Aird et al. 2018) and is subject to short-term variability relative to the stellar emission from the host (0.1–1 Myr vs. 100 Myr; Novak et al. 2011; Mullaney et al. 2012; Aird et al. 2013; Hickox et al. 2014; Schawinski et al. 2015). Moreover, the competing host-galaxy light from star-forming processes or the effect of circum-nuclear gas and/or dust obscuration may easily wash out AGN emission in individual objects. Thus, a growing practice is now to measure the average AGN power imprinted on large and homogeneous galaxy samples. These techniques often include image stacking (e.g. Mullaney et al. 2012; Chen et al. 2013; Rodighiero et al. 2015; Yang et al. 2017; Carraro et al. 2020; Ito et al. 2022), phenomenological modelling of the AGN luminosity function (Caplar et al. 2015, 2018; Weigel et al. 2017; Jones et al. 2019; Bernhard et al. 2019; Delvecchio et al. 2020), Bayesian modelling of sBHAR distributions with detections and non-detections (e.g. Aird et al. 2018; 2019; Grimm et al. 2019), or N-body simulations via continuity equations and ‘abundance matching’ (e.g. Behroozi et al. 2013; Grylls et al. 2019; Shankar et al. 2019; Allevalo et al. 2021).

The emerging consensus from deep X-ray observations at $0 \leq z < 3$ is that radiative (X-ray) AGN activity appears to be prevalent in more massive and distant star-forming galaxies (SFGs; e.g. Yang et al. 2018; Aird et al. 2019; Carraro et al. 2020; Delvecchio et al. 2020). Specifically, Aird et al. (2019) find that X-ray AGN at a given redshift are more frequently triggered in more massive galaxies, while – at fixed M_* – the typical sBHAR (or L_X) increases with redshift, possibly induced by the increasing cold gas fractions (Tacconi et al. 2018, 2020; Liu et al. 2019).

When averaging over all X-ray luminosities, the linear mean $\langle L_X \rangle$ is found to strongly correlate with galaxy M_* , featuring a super-linear ‘X-ray AGN main sequence’ (gradient ~ 1.5 ; e.g. Aird et al. 2019). This behaviour suggests enhanced radiative AGN activity in massive galaxies relative to SF, which by contrast sub-linearly increases with M_* along the star-forming main sequence (MS; e.g. Noeske et al. 2007; Elbaz et al. 2011; Schreiber et al. 2015; Lee et al. 2015; Rinaldi et al. 2022). Therefore, the shape and evolution of $\langle L_X \rangle$ encapsulates important clues regarding the interplay between SMBH and galaxy growth.

Building upon the above studies, our main goal is to investigate the relationship between radio-AGN activity and galaxy SF within SFGs, at various stellar masses and redshifts. Specifically, this study aims to elucidate the role of the host galaxy in triggering and sustaining long-term radio-AGN activity over cosmic time. Similarly to what was done for X-ray AGN, we thus explore the existence of a possible ‘radio-AGN main sequence’ (RAMS) that links mean rest-frame 1.4 GHz AGN luminosity ($\langle L_{1.4}^{\text{AGN}} \rangle$) and galaxy M_* at different redshifts ($0.1 \leq z \leq 4.5$). Measurements of the typical radio AGN power across M_* -selected SFGs are currently inferred only in the local Universe ($z < 0.3$; e.g. Sabater et al. 2019). At higher redshifts, a common roadblock for the calculation of $\langle L_{1.4}^{\text{AGN}} \rangle$ is how to quantify contamination from SF-driven synchrotron emission, which not only varies across the galaxy population, but is notoriously dominant over AGN emission at 1.4 GHz flux densities below 100 μJy (e.g. Prandoni & Seymour 2015; Smolčić et al. 2017b; Novak et al. 2018; Kono & Takeuchi 2021), where the bulk of radio SFGs and radio-faint AGN lie (Padovani 2016).

Radio-synchrotron emission at rest-frame 1–10 GHz offers an independent baseline to estimate dust-unbiased star formation rates (SFRs; Condon 1992; Murphy et al. 2011; Prandoni & Seymour 2015; but see Algera et al. 2021 for free-free emission at higher frequencies) and thereby decompose the radio emission into SF and AGN contributions. Local SFGs follow a strikingly tight ($\sigma < 0.2$ dex; Molnár et al. 2021) correlation between total IR luminosity (rest-frame 8–1000 μm , L_{IR}) and 1.4 GHz luminosity $L_{1.4}^{\text{SF}}$ arising from SF. This so-called infrared-radio correlation (IRRC) is often expressed in terms of $q_{\text{IR}} \approx \log(L_{\text{IR}}/L_{1.4}^{\text{SF}})$ (e.g. Helou et al. 1985; Condon 1992; Bell 2003 and references therein). At higher redshifts, flux-limited IR-radio samples yield a significant decline in q_{IR} with redshift (e.g. Magnelli et al. 2015; Delhaize et al. 2017). However, when accounting for various selection effects, Delvecchio et al. (2021, hereafter D21) find that the median SF-driven q_{IR} (or q_{IRRC}) shows a primary dependence on M_* and a much less significant redshift evolution. Their relation can be rewritten in log-space as

$$q_{\text{IRRC}}(M_*, z) = A + B \cdot \log(1 + z) + (C) \cdot \left(\log \frac{M_*}{M_\odot} - 10 \right), \quad (1)$$

where $A = (2.64 \pm 0.024)$, $B = (-0.137 \pm 0.048)$, and $C = (0.148 \pm 0.013)$. This was obtained by exploiting >400 000 *NUVrJ*-selected SFGs in the Cosmic Evolution Survey (COSMOS) field and stacking their ancillary IR images (*Herschel*, SCUBA, and AzTEC; Jin et al. 2018) and radio images (VLA-3 GHz; Smolčić et al. 2017b, as well as depth-matched MeerKAT-1.3 GHz; Jarvis et al. 2016; Heywood et al. 2022) across an unprecedented $M_* - z$ range, removing radio AGN contamination through a recursive approach. These findings suggest that more massive SFGs are radio brighter, at fixed L_{IR} , than lower- M_* analogues. A broadly similar M_* dependence has been independently confirmed by deep Low Frequency Array

¹ If not otherwise specified, L_X is the absorption-corrected rest-frame [2–10] keV luminosity from AGN.

(LOFAR) 150-MHz data at $z \lesssim 1$ over a ten times larger area than COSMOS (Smith et al. 2021; Bonato et al. 2021), as well as by local derivations (Molnár et al. 2021; Matthews et al. 2021; Heesen et al. 2022).

The implications of these new M_* -dependent recipes are twofold. Firstly, they enable us to reliably convert radio emission into SFRs, from L_{IR} (e.g. Kennicutt & Evans 2012) or from SFRs based on the spectral energy distribution (SED; e.g. Smith et al. 2021). For instance, the prescription from D21 has proved useful for reproducing the evolution of the star formation rate density (SFRD) at $z > 3$ (van der Vlugt et al. 2022), reaching the current best agreement with the dust-corrected UV estimate by Madau & Dickinson (2014) as compared to pure z -declining IRRCs. Secondly, these M_* -dependent prescriptions are instrumental for identifying ‘radio-excess’ AGN, that is, AGN that display excess radio emission (i.e. lower q_{IR} or a higher radio/SFR ratio) relative to that expected from SF alone (e.g. Donley et al. 2005; Del Moro et al. 2013; Delvecchio et al. 2017) across a wide range of M_* and redshifts.

In the present study we leverage the above-mentioned (M_*, z)-dependent IIRC to quantify $\langle L_{1.4}^{AGN} \rangle$ as a function of M_* and redshift across the global population of M_* -selected SFGs. From known radio-excess AGN placed at $>2\sigma$ from the IIRC (with $\sigma \sim 0.22$ dex, D21), we adopt a novel approach for factoring in the statistical contribution of radio AGN within the IIRC, which is critical for computing a representative sample-averaged AGN power. Firstly, we construct the 1.4 GHz luminosity function of radio-excess AGN in SFGs (AGN RLF hereafter) at different redshifts, following previous studies (e.g. Smolčić et al. 2017c; Novak et al. 2018; Ceraj et al. 2018; Butler et al. 2019; Kono & Takeuchi 2021). Secondly, to quantify the mean radio-AGN power at each M_* , we split the AGN RLF across different M_* and redshift bins, fitting and integrating each luminosity function down to the exact same $L_{1.4}^{SF}$ set by the IIRC at that (M_*, z). This self-consistent approach allows us to assess the cumulative power exerted by radio AGN in SFGs at various M_* , including the elusive contribution of radio-faint AGN within the IIRC.

The layout of this paper is as follows. Section 2 describes the sample selection. Section 3 presents the M_* -dependent radio source classification. The 1.4 GHz AGN RLF split into different M_* and redshift bins is discussed in Sect. 4. Section 5 illustrates the integrated and mean power of radio AGN across the entire SFG population, including our first derivation of the RAMS. In Sect. 6 its shape and evolution are compared to those of X-ray AGN and SF, with the broad implications for AGN feedback in SFGs over cosmic time discussed. We report our main conclusions in Sect. 7.

Throughout this paper, magnitudes are given in the AB system (Oke 1974). We assume a Chabrier (2003) initial mass function and a Λ cold dark matter cosmology with $\Omega_m = 0.30$, $\Omega_\Lambda = 0.70$, and $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Spergel et al. 2003).

2. Sample selection

2.1. Selection criteria and final sample

For this analysis we exploited deep radio-continuum data from the Very Large Array (VLA) COSMOS 3 GHz Large Project (Smolčić et al. 2017b), one of the most sensitive radio surveys ever conducted across a medium sky area like COSMOS (rms = $2.3 \mu\text{Jy beam}^{-1}$). With an angular resolution of $0.75''$, the total number of $S/N > 5$ detections reaches 10 830 over an area of 2.6 deg^2 . As detailed in Smolčić et al. (2017a), 8696 detec-




Fig. 1. Flux distribution of VLA-3 GHz sources detected at $S/N \geq 5$ (8696, black) over the COSMOS 1.77 deg^2 field (Smolčić et al. 2017b). The subset with an optical/NIR counterpart in the COSMOS2015 catalogue (Laigle et al. 2016) is also highlighted (7729 sources, red dashed).

tions (Fig. 1, black histogram) are contained within 1.77 deg^2 with optical and near-infrared (NIR) coverage. Of these, 7729 (Fig. 1, red histogram) were assigned an optical/NIR counterpart from the COSMOS2015 catalogue (Laigle et al. 2016) via a maximum likelihood approach. From these 7729 counterparts, we further restrict ourselves to the central 1.5 deg^2 area of the COSMOS field with deeper Ultra-VISTA coverage (6742 sources), in order to also exploit available de-blended far-IR/sub-millimetre photometry (Jin et al. 2018) extracted on K_s , Multi-Band Imaging Photometer for *Spitzer* (MIPS) $24 \mu\text{m}$ and VLA positional priors (see also Liu et al. 2018). This dataset was used in D21 to retrieve de-blended L_{IR} estimates for individual far-infrared (FIR)/sub-millimetre detections, as well as to obtain median-stacked L_{IR} from non-detections in different M_* and redshift bins.

Furthermore, within our 6742 VLA-3 GHz detections, we selected only radio sources within the ‘blue’ wedge of the [NUV-r]/[r-J] diagram (about 85%), using dust-corrected rest-frame magnitudes from Laigle et al. (2016). This criterion restricts our final sample to 5658 3 GHz detected SFGs (see e.g. Davidzon et al. 2017). Adopting the (M_*, z)-dependent IIRC from D21, calibrated on SFGs, enables us to self-consistently identify radio-excess AGN within our sample. We acknowledge that radio-detected galaxies are biased towards higher stellar masses, at fixed redshift, than a M_* -selected sample (as in D21). Interpreting the properties of radio-detected AGN from the viewpoint of M_* -selected SFGs will be addressed in Sect. 5.2.

We motivate our choice of not including *NUVrJ*-selected passive galaxies in our analysis in Sect. 4.4. Finally, for consistency with the $M_* - z$ space over which the IIRC was calibrated (D21), we consider sources within the same range: $0.1 \leq z \leq 4.5$ and $9 \leq \log(M_*/M_\odot) \leq 12$. These cuts yield our final sample of 5,658 radio-detected ($S/N > 5$ at 3 GHz) SFGs across 1.5 deg^2 . Table 1 outlines the various steps towards the final sample.

2.2. Parameter estimates

For consistency with the publicly available VLA catalogue (Smolčić et al. 2017a), we adopt the same 3 GHz radio flux

Table 1. Main numbers that lead to our final sample of 5658 radio-detected galaxies.

Definition	#
VLA-3 GHz (1.77 deg ²)	8696
VLA-3 GHz + Opt/NIR counterparts (1.77 deg ²)	7729
– UltraVISTA area (1.5 deg ²)	6742
– <i>NUVrJ</i> star-forming galaxies	5734
– final sample:	
$0.1 \leq z \leq 4.5$ and $9 \leq \log(M_*/M_\odot) \leq 12$	5658

Notes. The subset of radio AGN used throughout this work will be introduced in Sect. 3.

densities. These are scaled to 1.4 GHz assuming $S_\nu \propto \nu^\gamma$, with the 1.4–3 GHz spectral index (γ) being inferred from the observed flux densities at both frequencies whenever available ($\approx 25\%$ of the sample). Otherwise, we assume $\gamma = -0.75 \pm 0.1$ (Condon 1992; Ibar et al. 2009, 2010; Magnelli et al. 2015), which is close to the median spectral index of the 3 GHz population using survival analysis (see Sect. 4 in Smolčić et al. 2017b). We verify that taking $\gamma = -0.75 \pm 0.1$ for all galaxies would not affect the conclusions of this work (see Appendix C).

Estimates of photometric redshifts and M_* are taken from the COSMOS2015 catalogue (Laigle et al. 2016) via SED fitting of the optical–mid-infrared (MIR) photometry, reading the median value of the likelihood distribution for each source. The typical photometric redshift accuracy is $\langle |\Delta z|/(1+z) \rangle = 0.007$ at $z < 3$, and 0.021 at $3 < z < 6$ (Laigle et al. 2016), which increases up to only 0.057 for the faintest galaxies ($25 < i^+ < 26$). The ‘super-deblended’ catalogue presented by Jin et al. (2018) also contains publicly available spectroscopic redshifts ($\approx 38\%$ of the sample, courtesy of M. Salvato), which were prioritised over photometric measurements if deemed reliable (z_s quality flag > 3), or allowing for redshift variations within $\pm 10\%$ for sources with only a photometric redshift (see Jin et al. 2018). The same redshifts were used to compute rest-frame 1.4 GHz spectral luminosities ($L_{1.4}$). We note that M_* estimates were all based on photometric redshifts from COSMOS2015, after verifying a good consistency also for spec- z sources (Jin et al. 2018).

Uncertainties on $L_{1.4}$ are obtained by propagating the errors on flux density and spectral index. As mentioned in Sect. 2.1, for each radio source with FIR detection (i.e. with combined $S/N > 3$ over all FIR/sub-millimetre bands) from Jin et al. (2018), a single L_{IR} measurement is taken from their catalogue. For radio sources with no FIR detection, instead, stacked fluxes in each (M_*, z) bin were obtained from FIR/sub-millimetre stacking in D21 (see their Sect. 3). These were then converted to L_{IR} via SED fitting and finally re-scaled to the M_* and redshift of each source via the MS (Sect. 4.2.1 in D21) to mitigate underlying sample variance within each bin. Potential mid-IR AGN contamination is accounted for through empirical AGN templates (Mullaney et al. 2011) in the SED fitting (Liu et al. 2018; 2021). These L_{IR} measurements are preferred to 3σ L_{IR} upper limits inferred from FIR/sub-millimetre SED fitting (Jin et al. 2018) as they provide more stringent constraints for non-detections.

3. Radio source classification

In this section we describe the methodology adopted in this work to identify AGN through the radio-excess criterion. This method aims to distribute galaxies into two categories depending on the physical process likely producing the radio emission, regardless

of other multi-band AGN diagnostics. This can be carried out by decomposing the observed q_{IR} (or equivalent formalism) distribution of radio detections into a dominant (at flux densities below $100 \mu\text{Jy beam}^{-1}$; e.g. Prandoni & Seymour 2015) population of SFGs and a smaller, skewed low- q_{IR} tail, ascribed to radio AGN. The dividing threshold is somewhat empirical, but reflects a trade-off between AGN purity and completeness. In the literature, such radio-excess thresholds have been taken as redshift-invariant (e.g. Del Moro et al. 2013), or mildly redshift-dependent (e.g. Delvecchio et al. 2017 and used in Smolčić et al. 2017b). Contrary to these studies, we base our radio-excess criterion on the (M_*, z) -dependent IRRC derived in D21 (Eq. (1)), as detailed below. A discussion on the impact of using a different radio-excess criterion is given in Appendix C.

3.1. A M_* -dependent radio source classification

To categorise radio detections into SFGs or radio-excess AGN, we rely upon the recent (M_*, z) -dependent IRRC from D21. For each source we calculate its q_{IRRC} (from Eq. (1)), and we measure the offset from the observed q_{IR} , namely $\Delta q_{IRRC} = q_{IR} - q_{IRRC}$. As detailed in D21, this decomposition analysis in the space of Δq_{IRRC} can be used to separate the AGN and SFG populations on statistical basis. We recall that D21 restricted this procedure to a subset of radio detections for which the observed q_{IR} range is accessible by a *NUVrJ*-selected SFG, at each (M_*, z) . Their requirement translated into a M_* cut (i.e. $M_* > 10^{10.5} M_\odot$). This procedure was then extended to lower M_* using stacked q_{IR} measurements.

On the other hand, in the present analysis we aim at constructing the AGN luminosity function of 3 GHz detections, and thus we used their full observed q_{IR} (and M_*) range. Hence, we study the Δq_{IRRC} distribution of all radio detections in four different M_* bins (see Fig. 2): $9 < \log(M_*/M_\odot) < 10$; $10 < \log(M_*/M_\odot) < 10.5$; $10.5 < \log(M_*/M_\odot) < 11$ and $11 < \log(M_*/M_\odot) < 12$. The internal redshift variations are already factored in the q_{IRRC} term.

Figure 2 (top panels) displays the observed Δq_{IRRC} distribution of our 5658 radio-detections in all four M_* bins (increasing from left to right). The total distribution (black) is split between SFGs (blue dot-dashed) and AGN (red) following these steps: (i) the mode of the distribution is identified and assumed to trace pure SF; (ii) the right-hand side of the Δq_{IRRC} distribution is mirrored to the left, building a symmetric (log-normal like) function interpreted as purely driven by SF (blue dot-dashed); (iii) the residual distribution in excess to SFGs is statistically ascribed to AGN (red). We note that the SF histogram peaks at $\Delta q_{IRRC} \approx 0$ in all M_* bins, which highlights the validity of our IRRC prescription to the entire sample. For simplicity, a fixed threshold of $\Delta q_{IRRC} = -0.44$ dex (green vertical line) is set to separate SFGs (above) from radio-excess AGN (below). This threshold indicates a 2σ offset from the IRRC, with σ being the dispersion of the SF Gaussian found at $M_* > 10^{10.5} M_\odot$ (D21).

D21 estimated that roughly 80% of the AGN distribution lies below this dividing threshold, assuming a log-normal function. Here we refrain from fitting the AGN distribution at each M_* , since a proper characterisation of the shape of the AGN population is beyond the scope of this paper. However, we acknowledge that radio AGN within the IRRC might be under-represented from our method, as they are defined from the residual of the SFG distribution. Hence, we do not attempt at quantifying radio AGN within 2σ from the IRRC directly from the observed distribution. Instead, the contribution of such radio-faint AGN will be estimated indirectly from the AGN RLF and is discussed in Sect. 4.




Fig. 2. Radio source classification as a function of M_* . *Top panels:* distribution of 3 GHz detections (black) as a function of $\Delta q_{\text{IRRC}} = q_{\text{IR}} - q_{\text{IRRC}}$, split into different M_* bins (increasing from left to right). The subsample of radio SFGs (blue dot-dashed) and AGN (red) are separated on a statistical basis at a threshold of $\Delta q_{\text{IRRC}} = -0.44$ dex (vertical green line), that is, a 2σ offset from the IRRC. *Bottom panels:* distribution of AGN classification purity (or f_N^{AGN} , solid red line) and AGN 1.4 GHz luminosity purity ($f_{1.4}^{\text{AGN}}$, dashed grey line), both shown as a function of Δq_{IRRC} . At the threshold $\Delta q_{\text{IRRC}} = -0.44$ dex, we get $f_N^{\text{AGN}} \approx f_{1.4}^{\text{AGN}} \approx 64\%$. For details, see Sect. 3.2.

3.2. Statistical corrections to the classification method

A caveat of our approach concerns the purity of the radio-excess AGN sample. While galaxies within $\pm 2\sigma$ from the IRRC are likely to have radio emission mainly powered by SF, SFGs can partially contaminate the AGN population at the threshold $\Delta q_{\text{IRRC}} \lesssim -0.44$ dex (D21). This is because D21 prioritised a clean identification of SFGs rather than that of radio-excess AGN. Here we quantify this effect, while the corresponding correction for these misclassified galaxies is described in Appendix B.1.

We thus define $f_N^{\text{AGN}} = N_{\text{AGN}} / (N_{\text{AGN}} + N_{\text{SF}})$ as the number of radio AGN (N_{AGN} , red histogram) divided by the number of all radio-detections ($N_{\text{AGN}} + N_{\text{SF}}$, black histogram), in each Δq_{IRRC} bin. This is a proxy for AGN purity (in number), that is, the probability of a radio source being classified as a radio AGN in our sample at a given Δq_{IRRC} and M_* . This is marked in Fig. 2 (bottom panels). Error bars on f_N^{AGN} are propagated from the number ratio assuming Poissonian statistics if $N_{\text{AGN}} \geq 5$, otherwise we used the tabulated 1σ confidence intervals from Gehrels (1986). It is evident that $f_N^{\text{AGN}} \sim 60\text{--}70\%$ at the threshold of $\Delta q_{\text{IRRC}} = -0.44$ dex, and hence such a correction is non-negligible in any M_* bin. We also note that, although the f_N^{AGN} functions appear self-similar at all M_* , the relative fraction of radio-excess AGN apparently increases at $\log(M_*/M_\odot) < 10$. This is mainly a selection effect due to our sample being progressively less complete at lower M_* , and thus biased towards the brightest radio detections (i.e. AGN).

A further caveat caused by our simple thresholding method is that 100% of radio light is implicitly assumed to originate from the process of the corresponding class (AGN or SF). Instead, each object likely hides a composite nature that is important to

assess, particularly in radio-faint AGN close to the IRRC, where both emission processes can substantially contribute to the total (e.g. Maini et al. 2016; Herrera Ruiz et al. 2017; Radcliffe et al. 2018). Following previous studies (e.g. Ceraj et al. 2018, 2020), the offset from the IRRC can statistically trace the AGN fractional contribution at 1.4 GHz, which is defined as: $f_{1.4}^{\text{AGN}} = 1 - 10^{\Delta q_{\text{IRRC}}}$. Similar to f_N^{AGN} , $f_{1.4}^{\text{AGN}}$ also ranges from 0 to 1, as shown in the bottom panels of Fig. 2 (dashed grey line). Moreover, at the threshold $\Delta q_{\text{IRRC}} = -0.44$ dex we obtain $f_{1.4}^{\text{AGN}} \approx 64\%$, implying that roughly a third of the 1.4 GHz luminosity can be contaminated by SF. Instead, the median fractions obtained over all Δq_{IRRC} bins below -0.44 dex are 92% (for f_N^{AGN}) and 99% (for $f_{1.4}^{\text{AGN}}$), suggesting overall a highly reliable AGN selection. As discussed in more detail in Appendix B.2, each $L_{1.4}$ estimate can be scaled by the corresponding AGN fraction at 1.4 GHz to isolate the statistical AGN-related contribution in our sample, thus purifying the AGN RLF.

It is worth noting that f_N^{AGN} (solid red line) drops with Δq_{IRRC} more steeply than $f_{1.4}^{\text{AGN}}$ (dashed grey line). This suggests that our AGN identification method does not perform evenly at all Δq_{IRRC} , as it picks AGN more easily in the radio-excess regime than within the IRRC, after controlling for the varying AGN fraction at 1.4 GHz. However, we also acknowledge that both f_N^{AGN} and $f_{1.4}^{\text{AGN}}$ rely upon the measured Δq_{IRRC} , which becomes quite uncertain close to the IRRC, since the typical error bars on Δq_{IRRC} become larger than the actual Δq_{IRRC} . This argument, together with a poor sampling of the radio AGN distribution at $\Delta q_{\text{IRRC}} > -0.44$ dex (by construction), motivates our choice of not trusting the corrections for AGN classification and luminosity purity close to the IRRC. As described in Sect. 5.1, we only include radio AGN at $\Delta q_{\text{IRRC}} < -0.44$ dex, while the

contribution of radio-fainter AGN above that threshold will be statistically factored in the integrated luminosity function.

4. Luminosity function of radio-excess AGN in SFGs

We describe the methods used to construct the 1.4 GHz AGN RLF (Sect. 4.1) for our sample of radio-excess AGN. Details on the AGN classification purity and AGN $L_{1.4}$ purity are given in Appendix B. The computation of error bars in our data is detailed in Sect. 4.2. Later, we fit our data with z -evolving models, discussing the model parameters estimates and corresponding uncertainties (Sect. 4.3). Finally, in Sect. 4.5 we split the AGN RLF into various M_* bins, fitting the corresponding data with luminosity or density-evolving models at each M_* .

4.1. Building the radio-AGN luminosity function

To calculate the AGN RLF, we follow the procedure outlined in Novak et al. (2017, see their Sect. 3.1.). To summarise briefly, we employ the $1/V_{\max}$ method (Schmidt 1968), with V_{\max} being the maximum comoving volume over which a source is detectable, within the survey area and across a given redshift bin. Following Novak et al. (2017), we further corrected for a set of incompletenesses, including radio detection, heterogeneous noise, and resolution biases, as well as for missing optical/NIR counterparts (the latter being $\sim 10\%$; see Table 1).

We construct the 1.4 GHz AGN luminosity function in four redshift bins: $0.1 \leq z < 0.7$; $0.7 \leq z < 1.4$; $1.4 \leq z < 2.5$ and $2.5 \leq z \leq 4.5$. This redshift grid was chosen to be large enough to mitigate possible photometric uncertainties (i.e. sources falling into the wrong bin), while being centred on reference values of $z \approx 0.5, 1, 2,$ and 3 , respectively. We report AGN RLFs using the median $L_{1.4}$ in each luminosity bin (0.4 dex wide). The Poissonian uncertainty on the number density $\Phi(L)$ in each $(L_{1.4}, z)$ bin is calculated as in Marshall (1985) by weighting each galaxy by its contribution to the total $(1/V_{\max})^2$. However, if there are fewer than five sources in a luminosity bin, we used the tabulated 1σ values for small number statistics Gehrels (1986). We stress that the above steps and corrections to derive $\Phi(L)$ are identical to those used in previous luminosity functions based on VLA-COSMOS 3 GHz data (e.g. Smolčić et al. 2017c; Novak et al. 2018; Ceraj et al. 2018; van der Vlugt et al. 2022).

We refer the reader to Appendix B for a detailed explanation of how the AGN RLF was corrected for AGN classification purity and $L_{1.4}$ purity, following the reasoning in Sect. 3.2.

4.2. Error propagation via bootstrapping

The AGN RLF shown in Fig. 3 (red circles) represents our final dataset, whose error bars already incorporate several sources of uncertainty. Here we explain how these uncertainties propagate to the calculation of the final AGN RLF. Specifically, for each data point we account for errors on: (i) the 1.4 GHz luminosity $L_{1.4}$ of each object from the corresponding flux density and spectral index errors (see Sect. 2.2); (ii) the IR luminosity L_{IR} of each object (from SED-fitting), which translates into an error on q_{IR} ; (iii) the fraction f_N^{AGN} of AGN identified from our radio-excess criterion in a given Δq_{IRRC} bin, as displayed in Fig. 2; and (iv) the fraction $f_{1.4}^{\text{AGN}}$ of AGN-related luminosity at 1.4 GHz, whose uncertainty scales directly from (i) and from the scatter of the IRRC (0.22 dex).

We bootstrapped over all these uncertainties 1000 times, assuming at each step a Gaussian shape centred on the nominal parameter value, and a standard deviation given by the 1σ error bar on each parameter. For every realisation we re-calculate the position of all data points (in $L_{1.4}^{\text{AGN}}$ and $\Phi(L)$). Thus, we end up with 1000 bootstrapped AGN RLFs, which we interpolate at the 16th and 84th percentiles to delimit the $\pm 1\sigma$ confidence interval reported in Fig. 3. We note that the nominal sample size might slightly fluctuate among all AGN RLF realisations, due to some faint radio AGN crossing the threshold at $\Delta q_{\text{IRRC}} = -0.44$ dex.

The full baseline of AGN RLF data points and uncertainties is listed in Appendix D.

4.3. Fitting the redshift evolution of the AGN RLF

The redshift evolution of a galaxy-AGN population as a function of (L, z) is generally expressed by a joint density and luminosity evolution of its local luminosity function:

$$\Phi(L, z) = (1+z)^{\alpha_D} \cdot \Phi_0 \left[\frac{L}{(1+z)^{\alpha_L}} \right], \quad (2)$$

where α_D and α_L are the characteristic density and luminosity evolution parameters, respectively, $\Phi(L, z)$ is the observed AGN RLF, and Φ_0 is the local luminosity function. For consistency with the literature, we adopted the local AGN RLF from Mauch & Sadler (2007), which is parametrised as a double power law,

$$\Phi_0(L) = \frac{\Phi^*}{(L^*/L)^{\delta_1} + (L^*/L)^{\delta_2}}, \quad (3)$$

where the parameters are the normalisation $\Phi^* = \frac{1}{0.4} \times 10^{-5.5} \text{ Mpc}^{-3} \text{ dex}^{-1}$ calculated at the knee position $L^* = 10^{24.59} \text{ W Hz}^{-1}$, and the bright and faint end slopes $\delta_1 = -1.27$ and $\delta_2 = -0.49$ (with uncertainties of ± 0.18 and ± 0.04 , respectively; see Mauch & Sadler 2007). These parameters were obtained from 2661 radio detections at $z \sim 0$ spanning six decades in $L_{1.4}$ and are thus ideal for constraining both the faint and bright end of the local RLF.

Recent AGN RLF studies have also performed a global fit of all redshift bins to retrieve an evolution of $\Phi(L)$ with redshift (e.g. Novak et al. 2018; Kono & Takeuchi 2021). In that case, an additional parameter β_L (or β_D) for the redshift evolution is introduced, thus re-formulating the luminosity function of Eq. (2) as

$$\Phi(L, z) = (1+z)^{\alpha_D+z\beta_D} \times \Phi_0 \left[\frac{L}{(1+z)^{\alpha_L+z\beta_L}} \right]. \quad (4)$$

While this formalism is preferable to increase the statistics of input data, it implicitly assumes a simple linear trend of the total evolution parameter $(\alpha + z \cdot \beta)$ with redshift. This has been proved to be successful in describing the redshift evolution of the Mauch & Sadler (2007) AGN RLF (e.g. Novak et al. 2018; Kono & Takeuchi 2021).

We attempt a similar approach by using the Markov chain Monte Carlo (MCMC) algorithm, available in the Python package EMCEE (Foreman-Mackey et al. 2013), to perform a multivariate fit to these data. Nevertheless, we find that a combined four-parameter fit is loosely constrained by our data. Thus, we investigate two extreme cases of evolution: pure luminosity evolution (PLE; i.e. $\alpha_D = 0$) and pure density evolution (PDE; i.e. $\alpha_L = 0$). In both cases, we fitted each redshift slice independently (i.e. with $\beta_L = \beta_D = 0$), inferring a best-fit α_L (or α_D) in each bin (e.g. McAlpine et al. 2013; Smolčić et al. 2017c;




Fig. 3. Observed AGN RLF (red circles) fitted in each redshift bin with two parametrisations: PLE (black lines) and PDE (grey lines). The corresponding best-fitting function obtained through MCMC using the analytical form of [Mauch & Sadler \(2007\)](#) is marked with a solid line. The 1000 bootstrapped RLFs (dotted lines) delimit the $\pm 1\sigma$ confidence interval on the fit in each redshift bin. Further details are given in Sect. 4.2, and best-fit parameters and uncertainties are listed in Table 2.

Table 2. Results of modelling the evolution of the radio-excess AGN population in SFGs.

z-bin	$\langle z \rangle$	PLE			PDE		
		$\log L^*(z)$ [log(W Hz ⁻¹)]	α_L	χ^2_ν	$\Phi^*(z) \times 10^{-6}$ [Mpc ⁻³ dex ⁻¹]	α_D	χ^2_ν
0.1 ≤ z < 0.7	0.59	24.16 ± 0.07	-2.15 ± 0.51	0.34	4.10 ± 0.46	-1.42 ± 0.35	0.22
0.7 ≤ z < 1.4	1.12	24.64 ± 0.04	0.16 ± 0.19	0.49	8.46 ± 0.66	0.09 ± 0.13	0.44
1.4 ≤ z < 2.5	1.97	24.95 ± 0.04	0.77 ± 0.11	1.26	15.52 ± 1.56	0.62 ± 0.08	0.83
2.5 ≤ z ≤ 4.5	3.11	24.77 ± 0.07	0.29 ± 0.10	1.79	11.26 ± 1.89	0.25 ± 0.09	1.78

Notes. We assume either pure luminosity evolution (PLE) or pure density evolution (PDE) models, relative to the local relation from [Mauch & Sadler \(2007\)](#), as detailed in Sect. 4.3. The best-fitting RLF for each model is shown in Fig. 3.

[Ceraj et al. 2018, 2020; Kondapally et al. 2022](#)). We acknowledge that such a simplified method is conceptually similar to fitting L^* (for PLE) or Φ^* for PDE) at each redshift, assuming a fixed luminosity function shape. However, this formalism enables us to compare our evolving AGN RLF with existing radio-based literature. This approach also allows us to demonstrate that the evolution of radio AGN in SFGs cannot be described by a monotonic trend of the alpha parameters (see below).

Figure 3 shows the best-fit AGN RLF obtained in each redshift bin through a PLE (black curves) or PDE (grey curves) formalism. The dashed lines around each best-fit RLF are derived by bootstrapping 1000 times over the uncertainties of the evolution parameter α and delimit the uncertainty on the fit in each redshift bin. We list all output parameters and related 1σ uncertainties in Table 2. In particular, we convert the best-fit evolution parameters into the knee luminosity L^* for PLE and knee normalisation Φ^* for PDE, in each redshift slice. We note that only L^* is free to vary for PLE, whereas only Φ^* is free to vary for PDE (the other parameter is fixed to its local value). The reduced χ^2 (or χ^2_ν) highlights that both fitting forms successfully reproduce the observed data points (red circles). In agreement with previous studies, we find that either the L^* or Φ^* of radio AGN do increase with redshift, peaking at $z \sim 2$ and declining towards $z \sim 3$ (e.g. [Kondapally et al. 2022](#)).

In Fig. 4 we show our best-fit α_L (upper panel, PLE) α_D (bottom panel, PDE) as a function of redshift. We also display

the same parameters obtained from [Smolčić et al. \(2017c, grey squares\)](#), for the VLA-COSMOS 3 GHz sample, but including both SF and passive hosts. Grey lines by [Smolčić et al. \(2017c\)](#) are obtained by fitting all redshift bins at ($0 < z < 4$) and denote a declining evolution of $(\alpha + z \cdot \beta)$ with redshift. On the contrary, our data points show a reversal below $z \sim 2$.

This difference is likely induced by the lack of passive galaxies in our sample. As mentioned in Sect. B.2, about 30% of radio-excess AGN in the parent VLA-COSMOS 3 GHz sample are hosted in passive galaxies ([Smolčić et al. 2017c](#)). However, this fraction strongly drops with increasing redshift, namely: 57% at $0.1 \leq z < 0.7$; 46% at $0.7 \leq z < 1.4$; 14% at $1.4 \leq z < 2.5$; and 7% at $2.5 \leq z \leq 4.5$. Because the relative fraction of passive galaxies strongly varies over redshift, excluding them induces a notable drop in the normalisation of the AGN RLF (hence α_L or α_D) at $z < 2$ and thus a deviation from the monotonic decline of $(\alpha + z \cdot \beta)$ with redshift. Therefore, it is no surprise that the best agreement with previous AGN RLFs including all (SF+passive) AGN hosts (e.g. [Smolčić et al. 2017c; Novak et al. 2018; Kono & Takeuchi 2021; Kondapally et al. 2022](#)) is found at the highest redshifts (see Appendix C). We motivate our choice of removing passive galaxies in Sect. 4.4. As a reminder, we are mainly interested in exploring the relationship between radio AGN activity and galaxy growth in SFGs, but disentangling the two populations can elucidate their differential AGN demography and cosmic evolution.

If we were to perform a fitting of RLF data points in all redshift bins, we would obtain $\alpha_L = 0.39 \pm 0.22$ and $\beta_L = 0.02 \pm 0.09$ for PLE ($\chi^2_\nu = 3.64$), while $\alpha_D = 0.11 \pm 0.17$ and $\beta_D = 0.11 \pm 0.07$ for PDE ($\chi^2_\nu = 3.55$). With either formalism, the total evolution parameter ($\alpha + z \cdot \beta$) would return a slightly decreasing trend with redshift that systematically overestimates our AGN RLF at $z \lesssim 2$, yielding a not very good fit ($\chi^2_\nu > 3$; see above). We conclude that a linear evolutionary form is not applicable to fitting subsets of radio AGN (e.g. hosted by SFGs), if these follow a different redshift distribution relative to the global radio AGN sample. For this reason, we keep the evolution parameters derived in individual redshift bins (i.e. with $\beta_L = \beta_D = 0$), as listed in Table 2.

We further clarify that our non-monotonic trend of the evolution parameters with redshift (Fig. 4) only impacts the assumed evolution of the RLF, not its functional shape from Mauch & Sadler (2007). Nevertheless, we emphasise that our study does not imply that the functional form by Mauch & Sadler (2007) provides the best possible fit to our data. We simply argue that, at least in the observed $L_{1.4}$ range, the faint- and bright-end slopes seem to fit our data points quite well. As long as the faint-end (bright-end) slope does not become steeper (flatter) than unity, the integral will converge (see Sect. 5.1) and the results will broadly remain unchanged.

4.4. Removal of NUVrJ-passive galaxies

As mentioned in the previous sections, we refrain from including NUVrJ-passive galaxies in our sample, although they dominate the number density of radio-excess AGN at $z \lesssim 1$ (Sect. 4.3). Our motivations are the following.

4.4.1. Same IRRC?

Interpreting excess radio emission in passive galaxies implicitly assumes that the same IRRC calibrated for SFGs also applies to passive systems. Nonetheless, this is highly uncertain, since passive galaxies might be affected by contamination in their L_{IR} and $L_{1.4}$ (hence q_{IR}) estimates, which are instead needed in order to disentangle radio emission from star formation or AGN activity. Specifically, ‘cirrus’ emission associated with cold dust heated by old (>A-type) stellar populations might lower the intrinsic SFR at fixed L_{IR} , especially at low s-SFR (e.g. Yun et al. 2001). Moreover, radio emission can suffer from (largely unconstrained) millisecond pulsar contamination, which is negligible for SFGs but it can severely contaminate radio emission from SF in massive quiescent galaxies (e.g. Sudoh et al. 2021), thus complicating a proper IRRC calibration for this class.

4.4.2. Different cosmic evolutions

As highlighted in Sects. 4.3 and 5.1, passive galaxies host the bulk of radio AGN feedback at $z \lesssim 1$, and they strongly decline in number density at higher redshifts. As a consequence, mixing passive and SFGs in the same RLF washes out their differential cosmic evolution, as clearly demonstrated by the non-monotonic redshift trend of $(\alpha + z \cdot \beta)$ in SFGs (see Fig. 4).

4.4.3. Different AGN triggering?

As discussed in Sect. 6.4.1, radio-excess AGN show systematically higher BHARs in SF than in passive galaxies (e.g. Delvecchio et al. 2018), supporting a broad link between




Fig. 4. Best-fit evolution parameters α_L (upper panel) and α_D (bottom panel) obtained from PLE and PDE fitting forms, respectively. Our estimates (black circles) at each redshift are compared against those by Smolčić et al. (2017c, grey squares), which include both SF and passive hosts of radio-excess AGN. Grey lines by Smolčić et al. (2017c) indicate a declining evolution of $(\alpha + z \cdot \beta)$ with redshift, while our data points display a reversal at $z \sim 2$ likely ascribed to the lack of passive galaxies that would outnumber SFGs at lower redshifts. See Sect. 4.3 for details.

available cold gas and SMBH growth in radio AGN. Furthermore, as discussed in Sect. 6.4.1, recent studies find evidence that radio AGN in SF versus passive hosts have different duty cycle and cosmic evolution (e.g. Kondapally et al. 2022). These findings seem to suggest that radio AGN activity in SF versus passive galaxies is intrinsically driven by different mechanisms, and hence they should be analysed separately.

4.4.4. Comparison with literature

We aim to quantify the incidence, power and evolution of radio AGN activity specifically within SF hosts. This is to compare radio AGN activity, SF (set by the MS; e.g. Speagle et al. 2014) and X-ray AGN activity (from X-ray stacking; e.g. Carraro et al. 2020), all measured across the SFG population, via a self-consistent framework. Moreover, a recent paper by Ito et al. (2022) already inferred the mean radio-AGN luminosity in a sample of M_* -selected passive galaxies from COSMOS2020 (Weaver et al. 2022), via stacking of 3 GHz images (Smolčić et al. 2017b). This is discussed in Sect. 6.4.2. Unlike for SFGs, the stacking of quiescent hosts enables the underlying radio-AGN emission to be captured thanks to the high contrast of AGN-over-SF emission. On the contrary, in SFGs radio stacking is insensitive to radio-AGN emission (see D21), and hence we follow a novel approach based on the RLF.

4.5. Splitting the AGN RLF into M_* bins

The main goal of this work is measuring the mean and integrated radio AGN power across SFGs of different M_* and redshifts. To capture the dependence of radio AGN power on M_* , we decompose the AGN RLF obtained in the previous sections in various M_* bins, at fixed redshift. To the best of our knowledge,




Fig. 5. AGN RLF split into redshift and M_* . Coloured circles are the observed data points at each M_* , which add up to make the total RLF (black circles) at that redshift. Vertical dotted lines indicate the 1.4 GHz luminosity at $+2\sigma$ above the IRRC at a given (M_*, z) , which sets our $L_{1.4}$ threshold in the same bin to mitigate AGN incompleteness. Solid lines mark the best-fitting RLF obtained with PLE form. See Sect. 4.5 for details.




Fig. 6. Same as in Fig. 5, but showing the best-fit AGN RLF obtained with PDE form.

this approach is new in literature, probably since such a distinction requires large statistical samples of radio AGN up to high redshifts, and more importantly a proper treatment of AGN completeness as a function of M_* . The M_* -dependent IRRC prescription from D21 sets an ideal ground for this method (Sect. B.1 and Fig. 2). Therefore, we split our AGN RLF data points (coloured circles) into four different M_* bins, as displayed in Fig. 5. Empty black circles indicate the global RLF at that redshift bin across the full range $9 < \log(M_*/M_\odot) < 12$. As done for the RLF in each redshift bin, we again define the 1.4 GHz luminosity corresponding to $+2\sigma$ from the IRRC, this time calculated from the median M_* and redshift of the corresponding bin (vertical dotted lines). Because of the nearly linear relation of $L_{1.4}^{\text{SF}}$ with M_* set by the IRRC (Delvecchio et al. 2021), this radio luminosity threshold tends to increase with increasing mass, though at low M_* (e.g. middle panels of Fig. 6) it can reach below the reference $L_{1.4}$ threshold taken at that red-

shift, since this latter was computed at the median M_* of the full underlying sample (dotted black lines in the bottom panel of Fig. B.1). All RLF data points used in this analysis (i.e. above the corresponding $L_{1.4}$ threshold) are listed in Table D.1, for each M_* and redshift bin.

In order to track the macroscopic evolution of the radio-excess AGN population across different M_* , here we explore the same PLE and PDE fitting forms, as used in Sect. 4.3. Hence, for PLE we fix the knee normalisation to $\Phi^* = \frac{1}{0.4} \times 10^{-5.5} \text{ Mpc}^{-3} \text{ dex}^{-1}$, and for PDE we fix the knee luminosity $L^* = 10^{24.59} \text{ W Hz}^{-1}$ (i.e. their local values; Mauch & Sadler 2007). Figures 5 and 6 show the best-fit RLF split into M_* bins in the case of PLE and PDE fitting, respectively. We stress that the sum of best-fit RLFs over all M_* bins (coloured lines) is fully consistent with the best-fit RLF obtained at that redshift (black solid line), despite having been derived independently from one another. As a sanity check, we further verified that in

Table 3. AGN RLF fitting parameters derived in various M_* bins, at each redshift.

z -bin	M_* -bin $\log(M_*/M_\odot)$	PLE		PDE	
		$\log L^*(M_*, z)$ [$\log(\text{W Hz}^{-1})$]	χ^2_ν	$\Phi^*(M_*, z) \times 10^{-6}$ [$\text{Mpc}^{-3} \text{ dex}^{-1}$]	χ^2_ν
$0.1 \leq z < 0.7$	9–10	23.35 ± 0.29	0.15	1.32 ± 0.60	0.37
	10–10.5	23.27 ± 0.19	0.07	1.10 ± 0.41	<0.01
	10.5–11	23.58 ± 0.15	0.26	1.70 ± 0.44	0.20
	11–12	23.43 ± 0.35	0.43	1.16 ± 0.40	0.16
$0.7 \leq z < 1.4$	9–10	23.25 ± 0.15	1.34	0.44 ± 0.17	1.91
	10–10.5	23.81 ± 0.14	0.86	1.98 ± 0.48	1.13
	10.5–11	24.14 ± 0.12	0.48	3.48 ± 0.69	0.49
	11–12	24.00 ± 0.13	1.81	2.73 ± 0.54	1.08
$1.4 \leq z < 2.5$	9–10	23.66 ± 0.21	1.01	1.14 ± 0.34	1.52
	10–10.5	24.07 ± 0.10	0.89	2.65 ± 0.51	1.02
	10.5–11	24.46 ± 0.09	1.07	6.13 ± 1.30	1.30
	11–12	24.35 ± 0.12	0.71	4.77 ± 1.05	0.59
$2.5 \leq z \leq 4.5$	9–10	23.56 ± 0.21	0.13	0.62 ± 0.19	0.16
	10–10.5	23.85 ± 0.10	0.25	1.25 ± 0.30	0.41
	10.5–11	24.38 ± 0.11	0.23	4.86 ± 1.25	0.19
	11–12	23.94 ± 0.14	0.73	1.69 ± 0.57	0.66

Notes. Output parameter are listed for both PLE and PDE forms, as detailed in Sect. 4.5. The best-fitting RLF for each model is shown in Figs. 5 and 6.

each (M_*, z) bin, the integrated number density of radio AGN (from either the PLE or PDE form) never exceeds the number density of all M_* -selected SFGs.

We retrieve the best-fit L^* or Φ^* in each bin, and estimate their 1σ uncertainties via bootstrapping over the error bars of all data points in the same bin, as described in Sect. 4.2. Output parameters for each (M_*, z) bin are listed in Table 3. As expected, binning also with M_* adds-on noise in the evolution of the AGN RLF. Nevertheless, we observe a clear stratification in M_* for both PLE and PDE fitting forms. AGN at lower M_* are typically both less common and less luminous compared to AGN at higher M_* , with a peak being reached at intermediate-to-high stellar masses, that is, $10.5 < \log(M_*/M_\odot) < 11$ (cyan lines).

For sake of clarity, hereafter we only show the results obtained from the PLE form. Although the results are fully consistent with one another, we acknowledge that a PLE fitting performs slightly better than PDE in reproducing the RLF at low M_* , which is a critical domain to establish the global M_* dependence of the integrated radio AGN power. This choice is also in line with several previous studies of the AGN RLF (e.g. Sadler et al. 2007; Smolčić et al. 2017c, 2009; McAlpine et al. 2013; Padovani et al. 2015).

5. Results: Radio AGN activity across the SFG population

In this section we explore the integrated power emitted by radio AGN across the M_* -selected SFG population. First we calculate the cumulative (kinetic) AGN luminosity exerted as a function of redshift and M_* , highlighting its evolution compared to other works (Sect. 5.1). Then we follow a statistical approach to average the integrated radio AGN power at fixed (M_*, z) across the entire M_* -selected population of SFGs, deriving the RAMS that links mean radio AGN power and SFG stellar mass over cosmic time (Sect. 5.2).

5.1. Kinetic AGN luminosity density

The integral of the AGN RLF allows us to assess the cumulative luminosity released through radio AGN activity. This originates from mass accretion onto the central SMBH, and is channelled in kinetic form via collimated jets propagating throughout the galaxy (but see Panessa et al. 2019 for alternative origins of radio emission in radio quiet AGN). Though these jet signatures are widespread in powerful radio AGN at high- z (e.g. Nesvadba et al. 2017; Collet et al. 2016; Spingola et al. 2020; but see Radcliffe et al. 2018, 2021a for high- z radio-faint AGN) or in nearby radio AGN (e.g. Jarvis et al. 2019; Brienza et al. 2021; Venturi et al. 2021; Girdhar et al. 2022), only a small fraction of the kinetic energy carried by the jet and deposited in the interstellar medium is observable with monochromatic observations, the rest being dissipated in the environment (see reviews by e.g. McNamara & Nulsen 2007; Gitti et al. 2012). Several scaling relations have been proposed in the literature to convert monochromatic radio to kinetic luminosity (e.g. Willott et al. 1999; Bîrzan et al. 2004, 2008; Merloni & Heinz 2007; Cavagnolo et al. 2010; O’Sullivan et al. 2011; Daly et al. 2012; Godfrey & Shabala 2016). A comprehensive overview of these relations and their unknowns is given in Smolčić et al. (2017c, see their Appendix A).

For consistency with Smolčić et al. (2017c), we compute the kinetic AGN luminosity L_{kin} from the following relation formulated by Willott et al. (1999, scaled to 1.4 GHz),

$$\log(L_{\text{kin}}) = 0.86 \cdot \log(L_{1.4}^{\text{AGN}}) + 14.08 + 1.5 \cdot \log(f_W), \quad (5)$$

where L_{kin} is given in W, $L_{1.4}^{\text{AGN}}$ is the AGN-related 1.4 GHz luminosity in units of W Hz^{-1} . The parameter f_W encapsulates all uncertainties on the energetics and geometry of the jet, ranging from $f_W \approx 1$ to $f_W \approx 20$ (we assume $f_W = 4$; see below).

The kinetic AGN luminosity density at a given redshift, $\Omega_{\text{kin}}(z)$, is computed by multiplying L_{kin} by the AGN RLF,




Fig. 7. Kinetic AGN luminosity density, Ω_{kin} , as a function of redshift and dissected in M_* bins (coloured squares). The sum across all M_* at each redshift is marked with black open squares. For comparison, we show the integrated values from Smolčić et al. (2017c), purple lines), both for PLE (dot-dashed) and for PDE (triple dot-dashed) fitting forms. Model predictions from SAGE (Croton et al. 2016), including radio-mode AGN feedback, are displayed as a function of redshift (dashed black line). We assume $f_W = 4$ as in Smolčić et al. (2017c) for consistency. See Sect. 5.1 for further details.

$\Phi(L_{1.4}^{\text{AGN}})$, and integrating in 1.4 GHz luminosity,

$$\Omega_{\text{kin}}(z) = \int_{L_{1.4}^{\text{SF}}}^{L_{1.4}^{\text{max}}} \Phi(L_{1.4}^{\text{AGN}}) \cdot L_{\text{kin}} \cdot d(\log L_{1.4}^{\text{AGN}}). \quad (6)$$

We extend this calculation also in different M_* bins, computing $\Omega_{\text{kin}}(M_*, z)$ simply from the corresponding best-fit AGN RLF (Sect. 4.5). Unlike previous studies that usually assume an arbitrary minimum luminosity, we set the minimum to match $L_{1.4}^{\text{SF}}(M_*, z)$, which is the 1.4 GHz luminosity corresponding to the IRRC at a given (M_*, z) . This value is motivated by the need to account for radio-faint AGN that display no excess radio emission (i.e. lying within the scatter of the IRRC). We remind the reader that the RLF data points cover a $L_{1.4}^{\text{AGN}}$ range down to $+2\sigma$ above the IRRC, in which we are able to correct for AGN purity, and in which our sample is $\approx 90\%$ complete in $L_{1.4}^{\text{AGN}}$ (see Appendix B.3). Instead, we now extrapolate the best-fitting RLF down to the corresponding $L_{1.4}^{\text{SF}}$ to factor in extra radio-AGN emission not yet accounted for. The entity of such a correction is inherently tied to the assumed shape of the AGN RLF (i.e. Mauch & Sadler 2007) at the faint end regime, and will be discussed in Sect. 5.2.

Figure 7 displays the kinetic AGN luminosity density Ω_{kin} dissected in M_* (coloured squares), which add up to make the total Ω_{kin} at a given redshift (black open squares). Error bars at 1σ level are obtained by bootstrapping 1000 times over the uncertainties on the AGN RLF, listed in Tables 2 and 3.

The global $\Omega_{\text{kin}}(M_*, z)$ displays a clear stratification in M_* . Specifically, radio AGN in galaxies at $9 < \log(M_*/M_\odot) < 10$ display the smallest contribution, while the population at $10.5 < \log(M_*/M_\odot) < 11$ dominates the kinetic AGN luminosity density at all redshifts.

Irrespective of the M_* bin, the integrated kinetic AGN luminosity density in SFGs peaks at $z \sim 2$ and declines towards $z \sim 0$, though this drop is gentler in low- M_* galaxies. We emphasise

that the RLF in each M_* bin is fitted independently from the others, without forcing any internal monotonic M_* trend. Hence, the strong reported M_* stratification is genuine. We further note that the total $\Omega_{\text{kin}}(M_*, z)$ summed over all M_* bins (black open squares) is in good agreement with the $\Omega_{\text{kin}}(z)$ inferred from the integrated RLF at every redshift.

For comparison, we also show the global $\Omega_{\text{kin}}(z)$ inferred by Smolčić et al. (2017c) for PLE (dot-dashed line) and PDE (triple dot-dashed line) forms. For consistency with their work and for illustrative purposes, we also assume $f_W = 4$, though we note that any constant value in the range $f_W = 1-20$ would rigidly scale the data without affecting our main conclusions. However, we acknowledge that a M_* and/or z -dependent f_W could alter Ω_{kin} , but this possibility has been so far unexplored. Following this naive assumption ($f_W = 4$), we find fully consistent $\Omega_{\text{kin}}(z)$ measurements at $z \gtrsim 2$ with Smolčić et al. (2017c), while our data are a factor of 2–3 lower at lower redshifts. This is, again, most likely caused by the lack of passive galaxies in our sample. This apparent discrepancy closely resembles the offset seen in the evolution parameters α_L and α_D in Fig. 4. This further strengthens that radio-AGN activity taking place in passive galaxies dominates the kinetic luminosity density at $z \lesssim 1$.

Figure 7 shows the prediction from the semi-analytical galaxy evolution (SAGE; Croton et al. 2016) model. Compared to the former version by Croton et al. (2006), this updated model incorporates a realistic coupling between gas cooling and radio-mode (or jet-mode) AGN heating, which is a desirable refinement to predict the long-term impact of radio AGN feedback on the surrounding gas. This cooling-heating cycle is modulated through the so-called radio mode efficiency parameter ($k_R = 0.08$; see Eq. (16) and Sect. 9.1 in Croton et al. 2016). In SAGE, this parameter is used to modulate the BHAR (or \dot{m}_{BH}), and hence the AGN accretion luminosity $L_{\text{AGN}} = \eta \dot{m}_{\text{BH}} c^2$, where $\eta = 0.1$ is the standard radiative efficiency and c is the speed of light. Assuming that a given fraction of the accretion energy is channelled in kinetic – rather than radiative – form, we can re-scale L_{AGN} to our derived L_{kin} , and compare their volume-averaged luminosity density across cosmic time (Fig. 7). We find a good agreement at $z \gtrsim 2$ with the shape and normalisation of the SAGE model, while at lower redshifts our estimated $\Omega_{\text{kin}}(z)$ lies 3–6 times lower. This is not surprising, since the missing population of passive (i.e. quenched) galaxies in our study is the one undergoing most likely radio mode AGN feedback in the model, as a means to permanently turn off gas cooling and SF. Taking $f_W = 15$, as suggested by observations of radio lobes inflating X-ray cavities in local galaxy clusters (e.g. Birzan et al. 2004; 2008; Merloni & Heinz 2007; Cavagnolo et al. 2010; O’Sullivan et al. 2011), would match SAGE predictions at $z \sim 0.5$, albeit over-boosting our data at $z \gtrsim 1$ to 3–4 \times above the model. It is evident how the large uncertainties on f_W can easily accommodate an agreement, although we stress that alternative scaling relations to Eq. (5) would yield $\Omega_{\text{kin}}(z)$ systematically above SAGE predictions (see Appendix A in Smolčić et al. 2017c).

5.2. The radio-AGN main sequence (RAMS)

Each AGN RLF fit obtained in Sect. 4.5 is calibrated on a carefully selected sample of radio-excess AGN (at $>2\sigma$ from the IRRC) hosted in SFGs, taking into account flux incompleteness, classification purity and M_* -dependent radio emission from SF at each redshift. Therefore, we are well-placed to explore the intrinsic relationship between mean radio AGN power and galaxy M_* in SFGs, factoring

in the contribution of radio-faint AGN located within the IRRC.

We thus proceeded in two steps: firstly, we computed the cumulative power produced by the radio AGN population in SFGs at a given (M_*, z) ; secondly, we divided this integrated value by the number of all (*NUVRJ*-selected) SFGs contained in each bin to infer a representative sample-averaged radio-AGN luminosity.

Similarly to the calculation of the kinetic AGN luminosity density (Sect. 5.1), we integrate the best-fit AGN RLF above the radio luminosity set by the IRRC, at each (M_*, z) . However, a key difference now is that we want to assess the integrated power released by radio AGN across the full SFG population. The sample of SFGs is selected by M_* (Jin et al. 2018) and counts 236 763 galaxies identified via *NUVRJ*-colours at $9 \leq \log(M_*/M_\odot) \leq 12$ (from D21). These galaxies lie on the star-forming MS, while the subset of radio-detections stands slightly above MS (i.e. at higher L_{IR} at fixed M_*), especially at low M_* , since the radio flux limit sets a roughly horizontal cut in SFR. Based on the IRRC, lower L_{IR} (or SFR) imply also lower $L_{1.4}$ in M_* -selected SFGs than in radio detections. Thus, we adjust the new radio luminosity to match the $L_{IR}/L_{1.4}$ ratio of a typical MS galaxy at a given (M_*, z) . Specifically, we read L_{IR} from the MS fitting form presented in Daddi et al. (2022a, see their Table 1). They used median L_{IR} measurements obtained in D21 via IR-mm stacking for the same sample of M_* -selected SFGs. From L_{IR} we compute the mean 1.4 GHz luminosity of the IRRC for MS galaxies ($\langle L_{1.4}^{SF} \rangle$), at each M_* and redshift. These values set the minimum of the integral in Eq. (7) and are reported in Table 4. The integrated radio-AGN luminosity can be therefore expressed as

$$\sum L_{1.4}^{AGN} = \int_{\langle L_{1.4}^{SF} \rangle}^{L_{1.4}^{\max}} \Phi(L_{1.4}^{AGN}) \cdot \langle V_{\max} \rangle \cdot L_{1.4}^{AGN} \cdot d(\log L_{1.4}^{AGN}). \quad (7)$$

Differently from the Ω_{kin} calculation, our $\sum L_{1.4}^{AGN}$ is dimensionally a luminosity, not a luminosity density. Hence, in Eq. (7) we multiply the best-fit $\Phi(L_{1.4}^{AGN})$ by a characteristic $\langle V_{\max} \rangle$, taken as the median value across the underlying population in the same (M_*, z) bin. We propagate the dispersion around $\langle V_{\max} \rangle$ when assessing the uncertainty on $\sum L_{1.4}^{AGN}$. By integrating above $\langle L_{1.4}^{SF} \rangle$, we are implicitly assuming that no radio AGN is ‘active’ below the value set by the IRRC. This limit is empirically motivated by a (M_*, z) -dependent IRRC prescription (D21)². This approach ensures a fully self-consistent treatment of radio emission from SF and AGN activity.

As a sanity check, we quantified the effect of changing integration limits and extrapolating the luminosity function in the faint end. We find that the extra portion of the integral counted in the faint-end extrapolation (i.e. between $L_{1.4}^{SF}$ and the faintest observed $L_{1.4}$ -bin) is only about 20%, while the total number of AGN included in the extrapolation increases by a factor of 3–4. Therefore, this extrapolation does not substantially alter the integrated AGN luminosity density, while it is necessary to account for the global incidence of radio AGN.

We stress that setting the lower integration bound to a fixed canonical value (e.g. 10^{22} W Hz⁻¹, Ceraj et al. 2018) overestimates $\sum L_{1.4}^{AGN}$ by up to 50% at the highest M_* . Thus, we argue that accounting for the evolving $\langle L_{1.4}^{SF} \rangle$ with (M_*, z) is necessary for minimising AGN versus SF cross-contamination. Instead, increasing the upper integration limit $L_{1.4}^{\max}$ from 10^{28} (as

² We note that the contribution of ‘quiescent SMBHs’ (here assumed to have $L_{1.4}^{AGN}=0$) will be factored in by applying Eq. (8).

Table 4. Parameters used to compute the radio-AGN main sequence (RAMS) in Sect. 5.2.

RAMS parameters (Sect. 5.2)				
$\langle z \rangle$	$\langle \log(M_*) \rangle$ $\log(M_\odot)$	$\log(\langle L_{1.4}^{SF} \rangle)$ [log(W Hz ⁻¹)]	N_{SFG}	$\log(\langle L_{1.4}^{AGN} \rangle)$ [log(W Hz ⁻¹)]
0.52	9.42	21.67	15 463	21.01 ± 0.34
0.50	10.23	22.32	3610	21.54 ± 0.38
0.50	10.70	22.54	1708	22.19 ± 0.25
0.48	11.15	22.65	251	22.83 ± 0.43
1.05	9.42	22.02	49 264	20.85 ± 0.19
1.06	10.23	22.74	10 593	22.09 ± 0.26
1.06	10.70	23.02	5166	22.79 ± 0.15
1.08	11.16	23.19	813	23.44 ± 0.17
1.91	9.43	22.23	63 448	21.41 ± 0.30
1.86	10.23	23.05	13 265	22.55 ± 0.16
1.87	10.71	23.45	6677	23.30 ± 0.13
1.90	11.15	23.71	1315	23.85 ± 0.19
3.23	9.46	22.51	52 301	21.56 ± 0.30
3.19	10.21	23.33	9365	22.72 ± 0.14
3.09	10.69	23.78	2902	23.76 ± 0.14
3.21	11.20	24.17	622	23.87 ± 0.19

Notes. From left to right: the median redshift ($\langle z \rangle$), stellar mass ($\langle M_* \rangle$), and 1.4 GHz luminosity from SF ($\langle L_{1.4}^{SF} \rangle$) averaged over all M_* -selected SFGs in each bin; the number of all *NUVRJ*-selected SFGs (N_{SFG}); the mean radio AGN power averaged across N_{SFG} ($\langle L_{1.4}^{AGN} \rangle$).

in Novak et al. 2018; Ceraj et al. 2020) to infinity would boost $\sum L_{1.4}^{AGN}$ by only $\lesssim 10\%$. Over 50% of the cumulative AGN luminosity density is produced by AGN within ± 1 dex from the corresponding L^* (M_*, z), although these sources are a relatively small fraction of the galaxy population.

By dividing the cumulative radio-AGN luminosity $\sum L_{1.4}^{AGN}$ by the number of all M_* -selected SFGs in the same bin (N_{SFG}), we can compute the mean radio-AGN luminosity:

$$\langle L_{1.4}^{AGN} \rangle = \frac{\sum L_{1.4}^{AGN}}{N_{\text{SFG}}}. \quad (8)$$

This method follows the same logic of stacking, in which a representative sample-averaged measurement is inferred by combining detections and non-detections. The main difference is that stacked luminosities are corrected for SF contamination a posteriori; instead, our prescription set by the IRRC allows us to quantify and remove galaxy contamination a priori, by integrating down to $L_{1.4}^{SF}$. Moreover, stacking is broadly sensitive to the average signal from the dominant underlying population. Since radio emission from 3 GHz-undetected SFGs in COSMOS is primarily originated from SF (D21), radio stacking reveals a notable radio-excess only for passive galaxies (Ito et al. 2022). Instead, our statistical approach, backed-up with a detailed assessment of the evolving AGN RLF, can also account for hidden radio AGN in SFGs. Some caveats and limitations inherent to our approach are discussed in Sect. 6.1.

The RAMS between $\langle L_{1.4}^{AGN} \rangle$ and galaxy M_* (in log-space) is presented in Fig. 8 (left). Mean $\langle L_{1.4}^{AGN} \rangle$ measurements (squares) are coloured by redshift to highlight the evolution of this relationship. We used the IDL routine MPFIT2DFUN.PRO to perform a three-dimensional fitting in the $\langle L_{1.4}^{AGN} \rangle$ - M_* -($1+z$) log-space,




Fig. 8. RAMS versus M_* and redshift. *Left:* RAMS relating mean radio AGN power $\langle L_{1.4}^{\text{AGN}} \rangle$ averaged across all ($NUVrJ$ -based) SFGs and galaxy M_* , at different redshifts. Individual points are inferred from the integral of the AGN RLF at each (M_*, z) , following Eq. (8) and coloured by M_* . We fit all points with a three-dimensional function (Eq. (9)). Dashed areas encompass the $\pm 1\sigma$ scatter around the best-fit values obtained by bootstrapping over the uncertainties. *Right:* redshift evolution of the same $\langle L_{1.4}^{\text{AGN}} \rangle$ measurements, coloured by M_* and fitted as a function of $(1+z)^{P_2}$ (see Eq. (9)), with $P_2 = 2.51 \pm 0.34$.

assuming the following analytical expression,

$$\log\langle L_{1.4}^{\text{AGN}} / \text{W Hz}^{-1} \rangle = P_1 + P_2 \cdot \log(1+z) + (P_3) \cdot \left(\log \frac{M_*}{M_\odot} - 10 \right). \quad (9)$$

The three best-fit parameters are: the intercept $P_1 = (20.97 \pm 0.16)$, the $\log(1+z)$ slope $P_2 = (2.51 \pm 0.34)$, and the $\log(M_*)$ slope $P_3 = (1.41 \pm 0.09)$. Our three-dimensional fitting yields $\chi^2_v = 1.08$. Error bars on the best-fit parameters are given at 1σ level. By bootstrapping over the uncertainties on (P_1, P_2, P_3) at the mean redshift of each bin, the error on $\log\langle L_{1.4}^{\text{AGN}} \rangle$ is about 0.08 dex, at fixed (M_*, z) . For visual purposes, we show this error as a function of M_* at the mean redshift of each bin (coloured dashed areas) in Fig. 8 (left). All measurements are listed in Table 4.

With a significantly super-linear slope of 1.41 ± 0.09 (i.e. roughly 4σ steeper than unity) between $\langle L_{1.4}^{\text{AGN}} \rangle$ and galaxy M_* , this trend suggests that more massive SFGs contain. Moreover, as shown in Fig. 8 (right), the RAMS normalisation at fixed M_* clearly evolves with redshift ($\propto (1+z)^{2.51 \pm 0.34}$), mimicking the evolution of the MS relation (e.g. Daddi et al. 2022a) and the evolution of the molecular gas fraction in galaxies (roughly $\propto (1+z)^{2.5}$; e.g. Saintonge et al. 2017; Tacconi et al. 2018, 2020; Liu et al. 2019; Decarli et al. 2020; Walter et al. 2020; Wang et al. 2022). This trend suggests that the RAMS is in place at least since $z \sim 3$.

6. Discussion

In this section we caution the reader about some caveats and limitations related to the RAMS (Sect. 6.1). Then we further discuss the main implications of the existence of a RAMS in the framework of AGN-galaxy co-evolution: the relative contribution of AGN versus SF-driven radio emission (Sect. 6.2), the triggering of radio-AGN activity in SFGs (Sect. 6.3), and the long-term imprinting of AGN feedback on galaxy SF (Sect. 6.4).

6.1. Possible caveats and limitations of the RAMS

The naming RAMS intentionally echoes both the star-forming MS (e.g. Speagle et al. 2014) as well as the AGN MS (e.g. Mullaney et al. 2012) obtained from X-ray data. However, we emphasise that, unlike the MS of SFGs and similar to that of X-ray AGN, our RAMS is not visible for individual galaxies as it is ‘hidden’ by intrinsic (radio) AGN variability.

Our $\langle L_{1.4}^{\text{AGN}} \rangle$ measurements should not be interpreted as the ‘typical’ (i.e. most likely) radio-AGN luminosity observed in a MS galaxy at (M_*, z) . To address this, we would need to compute the probability distribution that a galaxy of a given (M_*, z) hosts a radio AGN with that luminosity. Assuming that radio AGN triggering is a stochastic process, a sample-averaged AGN emission should match a time-averaged emission. Therefore, each $\langle L_{1.4}^{\text{AGN}} \rangle$ estimate could be interpreted as a time-averaged luminosity of the AGN over the entire galaxy’s life cycle. This is further discussed in Sect. 6.4.

The strong M_* dependence reported in Eq. (9) is not artificially induced by the M_* -evolving 1.4 GHz luminosity limit from SF ($L_{1.4}^{\text{SF}}$). On the contrary, we do find enhanced radio AGN activity in more massive galaxies on top of increasing SF-driven radio emission. As a consequence, taking a fixed q_{IRRC} (e.g. 2.64 from the local Universe; Bell 2003) would have led to a higher number of radio-excess AGN in high- M_* galaxies than in this work, hence steepening the correlation between $\langle L_{1.4}^{\text{AGN}} \rangle$ and galaxy M_* .

We assume a fixed luminosity function shape from Mauch & Sadler (2007) throughout the full (M_*, z) range studied in this work. As detailed in Sect. 4.3, any functional form with a faint (bright)-end slope flatter (steeper) than unity would yield a converged integral, and hence the cumulative and average radio-AGN luminosity (Sects. 5.1 and 5.2) would remain stable. Since we do not find systematic deviations from Mauch & Sadler’s luminosity function across the observed $L_{1.4}$ range, we do not explore alternative functional forms.

It is perhaps confusing that the mean radio AGN power, at each M_* and redshift, is much lower than the observed $L_{1.4}^{\text{AGN}}$ range covered by our RLF data points (e.g. Fig. 5). We clarify

that this is purely the result of averaging the cumulative radio-AGN luminosity across the entire sample of M_* -selected galaxies in each bin, which outnumber radio detections by >100 at the lowest M_* (see Table 4). However, an important remark is that the majority of radio AGN feedback is originated from a small fraction of galaxies close to L^* , implying relatively short AGN phases at high radio power and much longer periods at low AGN power.

Stellar mass incompleteness might affect the real N_{SFG} at low M_* . However, we note that correcting for missing low- M_* galaxies would decrease the resulting $\langle L_{1.4}^{\text{AGN}} \rangle$, strengthening the observed super-linear trend with M_* .

Missing radio AGN might alter the RAMS shape and normalisation. It is, indeed, possible that we are underestimating the identified number of radio-faint AGN within the IRRC, since they do not feature a radio excess (Sect. 3.1). Although constraining their demography is critical for obtaining a full AGN census, in our analysis we have conservatively restricted our $L_{1.4}$ -range to above $+2\sigma$ from the luminosity at the IRRC, as in this regime we reach the highest degree of AGN purity and completeness. Then, by extrapolating the best-fit AGN RLF down to $\langle L_{1.4}^{\text{SF}} \rangle$, we are implicitly factoring in the statistical contribution of radio-faint AGN in SF-dominated sources, even if formally undetected at 3 GHz.

We acknowledge that the most massive galaxies ($M_* > 10^{11} M_\odot$) at $z \sim 3$ could be under-represented, especially if extremely dusty and, thus, not fully captured by an optical-NIR counterpart catalogue (Laigle et al. 2016). This might partly explain the slightly offset data point in Fig. 8 at the highest (M_*, z) .

6.2. Radio AGN emission is usually sub-dominant compared to SF

Having determined a representative (i.e. time-averaged) mean radio-AGN luminosity across a wide M_* and redshift range, we can compare our RAMS with the shape and evolution of radio emission arising from SF processes. The latter is taken from D21 and calibrated for the same M_* -selected ($NUVrJ$ -based) SFGs. Relating the average power of both phenomena enables us to assess the dominant process of radio emission in SFGs. Eq. (9) describes the mean radio AGN power at fixed (M_*, z) , while the mean SF-related emission $\langle L_{1.4}^{\text{SF}} \rangle$ is implicitly given in Eq. (1), though in the form of q_{IRRC} . For consistency with the formalism adopted in the AGN part, we use the same analytical function of Eq. (9) to also fit the AGN-to-SF luminosity ratio,

$$\log\left(\frac{\langle L_{1.4}^{\text{AGN}} \rangle}{\langle L_{1.4}^{\text{SF}} \rangle}\right) = Q_1 + Q_2 \cdot \log(1+z) + (Q_3) \cdot \left(\log\frac{M_*}{M_\odot} - 10\right). \quad (10)$$

We find the following best-fit parameters: $Q_1 = (-0.63 \pm 0.15)$, $Q_2 = (-0.05 \pm 0.31)$, $Q_3 = (0.61 \pm 0.08)$. The best-fit yields $\chi^2_\nu = 0.71$. Figure 9 shows the average AGN-to-SF luminosity ratio (squares), as a function of M_* and coloured by redshift. The dashed horizontal line indicates equal contributions from AGN and SF. Coloured lines indicate the best-fit ratio (solid) expressed in Eq. (10).

The redshift evolution of $\langle L_{1.4}^{\text{AGN}} \rangle / \langle L_{1.4}^{\text{SF}} \rangle$ is quite weak and consistent with a null slope (Eq. (10)). This is clearly illustrated in Fig. 9 by the nearly redshift-invariant behaviour. Our findings seem to suggest that radio AGN activity and galaxy SF, at fixed M_* , broadly evolve over time with a similar pace. This is probably only part of the full story, since we are mapping radio AGN activity solely inside SFGs. Indeed, we have seen




Fig. 9. Logarithmic ratio between AGN-related and SF-related radio emission ($\langle L_{1.4}^{\text{AGN}} \rangle$ and $\langle L_{1.4}^{\text{SF}} \rangle$, respectively), as a function of M_* , coloured by redshift. Each data point (square) represents the ratio between average AGN and SF luminosities for M_* -selected SFGs in the same bin. The dividing threshold between AGN and SF-dominated regions (dashed black line) is crossed at $M_* \sim 10^{11} M_\odot$, above which galaxy radio emission is mainly powered by AGN jets. Filled squares mark bins in which the mean $\langle L_{1.4}^{\text{AGN}} \rangle$ is above the 5σ VLA 3 GHz luminosity limit ($L_{1.4}^{\text{lim}}$, scaled to 1.4 GHz). Coloured lines indicate the best-fit ratio (solid) from Eq. (10).

in Sect. 5.1 that the cumulative energy exerted by radio AGN in SF+passive galaxies (Smolčić et al. 2017c) or via radio-mode feedback (Croton et al. 2016) is notably dominant over that produced in SFGs alone. However, at $z > 1$ ($NUVrJ$ -based) SFGs vastly outnumber passive galaxies, and contain the bulk of radio-excess AGN. Therefore, our findings should be providing a representative radio view of SMBH-galaxy growth at the cosmic noon.

The positive M_* slope obtained in Eq. (10) suggests a steeper M_* dependence of $\langle L_{1.4}^{\text{AGN}} \rangle$ than $\langle L_{1.4}^{\text{SF}} \rangle$. This is not surprising, given the super-linear relationship of the RAMS (Eq. (9)) and the typically sub-linear M_* dependence for the (radio) star-forming MS (D21)³. This enhancement of AGN activity with M_* is further discussed in Sect. 6.4.

Another take-away message from Fig. 9 is that, at any redshift, radio AGN emission is usually sub-dominant relative to that from SF. The only exception comes from the most massive galaxies at $M_* > 10^{11} M_\odot$ (at $z \lesssim 2$), whose radio emission is mainly ($\approx 65\%$ on average) powered by AGN jets. This result is in line with previous studies (e.g. Sabater et al. 2019) finding widespread radio AGN activity in these massive galaxies, albeit without separating between passive and star-forming systems.

However, as mentioned in Sect. 5.2, the bulk of radio AGN activity originates from relatively short phases in which the AGN emission is dominant over SF. This is what still allows us to calculate the mean $\langle L_{1.4}^{\text{AGN}} \rangle$ despite this being on average (i.e. across the galaxy’s lifetime) sub-dominant compared to $\langle L_{1.4}^{\text{SF}} \rangle$.

We emphasise that the regions labelled as ‘AGN-dominated’ or ‘SF-dominated’ do not reflect the probability of a radio AGN being turned on in a given galaxy, but rather a luminosity-

³ The consensus range for the MS slope is about 0.5–0.9 (Speagle et al. 2014 for a review; see also Leslie et al. 2020 from VLA-COSMOS 3 GHz data), albeit this is affected by the bending at the highest M_* , possibly linked to the transition from cold-to-hot gas accretion (e.g. Dekel et al. 2013; Daddi et al. 2022a,b; Popesso et al. 2022).

weighted probability. In other words, our average measurements carry a degeneracy between timescales and intensity of each episode of AGN activity. This distinction is supported by multiple observations of intermittent jet activity in radio galaxies (e.g. Jurlin et al. 2020; Brienza et al. 2021), which can occur over timescales comparable to those of galaxy SF (10^{7-8} yr; see e.g. Konar et al. 2013). Breaking such a degeneracy in observations is challenging at high redshifts since jet-driven emission appears intrinsically more compact (<1 kpc; e.g. Costa et al. 2018) than in luminosity-matched AGN at $z \sim 0$ (e.g. Bondi et al. 2018 based on the same VLA-COSMOS 3 GHz data), and also because radio-faint jets can easily be washed out by stellar-driven radio emission, which is enhanced at high redshifts (Padovani 2016; Magliocchetti et al. 2018; Delvecchio et al. 2018). Nonetheless, our statistical approach tied to the AGN RLF allows us to dissect the role of radio-AGN duty cycle and intensity of AGN activity, as discussed in Sects. 6.3.1 and 6.3.2.

To provide a rough idea of how incomplete is our current picture of radio AGN activity in SFGs, in Fig. 9 we highlight as filled squares those bins in which the mean $\langle L_{1.4}^{\text{AGN}} \rangle$ is formally above the 5σ VLA 3 GHz luminosity limit ($L_{1.4}^{\text{lim}}$, corresponding to $20 \mu\text{Jy beam}^{-1}$ at 1.4 GHz). This check allows us to test the detectability of the mean radio AGN activity across the galaxy’s life cycle by current deep radio surveys. Unsurprisingly, among AGN-dominated bins, only those at $M_* > 10^{11} M_\odot$ (and $z \lesssim 2$) would be formally detectable. Even pushing radio sensitivity to below μJy levels (e.g. with SKA1-MID or ngVLA), isolating the sub-dominant ($\approx 10\text{--}25\%$ at $M_* \sim 10^{10} M_\odot$, based on Fig. 9) AGN-related emission at 1.4 GHz will be out of reach also at sub-arcsecond angular resolution ($\geq 0.1''$ at 0.95–1.76 GHz in SKA1-MID Band-2; e.g. Braun et al. 2019; see also Sweijen et al. 2022 for LOFAR imaging). Thus, a promising alternative comes from very long baseline interferometry (VLBI) techniques, which are crucial for pinning down radio-AGN emission even in the SF-dominated regime (Maini et al. 2016; Herrera Ruiz et al. 2018; Muxlow et al. 2020; Radcliffe et al. 2021a,b). In this context, the added VLBI capability to the Square Kilometer Array (SKA) will be critical (Paragi et al. 2015) to reach a full radio-AGN census.

6.3. What drives the super-linear RAMS with M_* ?

The super-linear (1.41 ± 0.09) $\log M_*$ dependence of the RAMS suggests enhanced (radio) AGN activity compared to SF in massive SFGs. To make this trend explicit in terms of SFR, we rewrite Eq. (10) by converting $\langle L_{1.4}^{\text{SF}} \rangle$ into $\langle \text{SFR} \rangle$. To this end, we employ the z -evolving $\text{SFR}\text{--}M_*$ fitting form proposed by Lee et al. (2015) but applied to the data from D21 (see Table 1 of Daddi et al. 2022a). The resulting expression reads as follows:

$$\log \left(\frac{\langle L_{1.4}^{\text{AGN}} / \text{W Hz}^{-1} \rangle}{\langle \text{SFR} / M_\odot \text{ yr}^{-1} \rangle} \right) = R_1 + R_2 \cdot \log(1+z) + (R_3) \cdot \left(\log \frac{M_*}{M_\odot} - 10 \right). \quad (11)$$

With a reduced chi-square of $\chi_\nu^2 = 0.70$, the best-fit parameters are: $R_1 = (20.73 \pm 0.15)$, $R_2 = (0.08 \pm 0.32)$, $R_3 = (0.77 \pm 0.08)$. Because the $\text{SFR}\text{--}M_*$ relation is flatter than the $L_{1.4}^{\text{SF}}\text{--}M_*$ relation⁴, the above ratio $\langle L_{1.4}^{\text{AGN}} \rangle / \langle \text{SFR} \rangle$ exhibits a slightly steeper M_* dependence than in Eq. (10). This is also clearly shown in Fig. 10.

⁴ The difference in slope is ≈ 0.15 , that is, the net M_* dependence of the $L_{\text{IR}}/L_{1.4}^{\text{SF}}$ ratio (or q_{IRRC}) found in D21.




Fig. 10. Logarithmic ratio between $\langle L_{1.4}^{\text{AGN}} \rangle$ and $\langle \text{SFR} \rangle$, as a function of M_* , coloured by redshift. As in Fig. 9, all parameters are averaged across the entire sample of M_* -selected SFGs in the same bin. Coloured lines indicate the best-fit ratio (solid) obtained from Eq. (11).

6.3.1. The radio-AGN duty cycle across SFGs

A super-linear trend of $\langle L_{1.4}^{\text{AGN}} \rangle$ with M_* might be attributed either to a higher radio-AGN duty cycle in more massive galaxies or to an intrinsically brighter radio-AGN phase in more massive galaxies, or to both effects. Breaking such a degeneracy is important to assess the role of the host-galaxy M_* in driving radio AGN activity.

We leveraged the best-fit RLF obtained at each (M_*, z) in order to calculate the fraction of all SFGs hosting a radio AGN. For consistency with the formalism used so far, we set $\langle L_{1.4}^{\text{SF}} \rangle$ as the lower radio-luminosity limit for a galaxy to be in the AGN phase at each (M_*, z) . Under this assumption, the fraction of galaxy’s lifetime spent as AGN (i.e. the radio-AGN duty cycle) corresponds to the fraction of all SFGs above this limit ($f_{\text{gal}}(> \langle L_{1.4}^{\text{SF}} \rangle)$),

$$f_{\text{gal}}(> \langle L_{1.4}^{\text{SF}} \rangle) = \frac{\int_{L_{1.4}^{\text{SF}}}^{L_{1.4}^{\text{max}}} \Phi(L_{1.4}^{\text{AGN}}) \cdot \langle V_{\text{max}} \rangle \cdot \text{sd}(\log L_{1.4}^{\text{AGN}})}{N_{\text{SFG}}}, \quad (12)$$

where the numerator is the total number of galaxies in radio-AGN phase, scaled from the AGN RLF in a given (M_*, z) bin. Figure 11 (left) shows the radio-AGN duty cycle as a function of M_* , colour-coded by redshift. We find a close-to-linear best-fit relation ($\propto M_*^{0.97 \pm 0.07}$). A significant evolution is also observed, although weaker than that of the RAMS ($\propto (1+z)^{1.40 \pm 0.27}$). This result suggests that more massive and/or more distant SFGs undergo a progressively higher radio-AGN duty cycle. The duty cycle reaches close to 10% at $M_* \gtrsim 10^{11} M_\odot$. Other studies have also reported similar results. For instance, Best et al. (2005) found that the radio-AGN duty cycle in local passive galaxies strongly rises as $\propto M_*^{2.5}$ (see also Sabater et al. 2019). When selecting only star-forming hosts based on different cuts in the specific SFR ($\text{sSFR} = \text{SFR}/M_*$), Kondapally et al. (2022) found a slightly shallower behaviour, with the radio-AGN fraction scaling as $\propto M_*^{1.37 \pm 0.57}$ at redshift $0.3 < z \leq 1.5$ (see their Fig. 9). Consistently with our work, Kondapally et al. (2022) found that the fraction of SFGs hosting a radio AGN reaches $\approx 10\%$ at $10^{11.5} M_\odot$, albeit using a different AGN nomenclature (i.e. ‘SF-LERGs’; see details in




Fig. 11. Interpreting the evolving RAMS with (M_*, z) . *Left:* fraction of all SFGs in the radio-AGN phase (f_{gal}), i.e. having radio luminosity above $\langle L_{1.4}^{\text{SF}} \rangle$, at a given (M_*, z) . This fraction is equivalent to the radio-AGN duty cycle, as written in Eq. (12). Symbols and colours are as in Fig. 8. *Right:* mean radio luminosity in the radio-AGN phase ($\langle L_{1.4}^{\text{active}} \rangle$) as a function of (M_*, z) , as written in Eq. (13).

Sect. 6.4.1). Instead, the marginally steeper dependence of the duty cycle on M_* from Kondapally et al. (2022), can be partly explained by their different integration limit: the authors computed the duty cycle as the fraction of SFGs hosting a radio AGN with $\log[L_{150 \text{ MHz}}/\text{W Hz}^{-1}] \geq 24$, which corresponds to $\log[L_{1.4 \text{ GHz}}/\text{W Hz}^{-1}] \geq 23.3$ (if assuming $\gamma = -0.75$). This threshold is roughly consistent with our lower integration limit, $L_{1.4}^{\text{SF}}$, at $M_* \gtrsim 10^{11} M_\odot$ and $z \sim 1$, whereas it is systematically higher (by 1 dex) than $L_{1.4}^{\text{SF}}$ at $M_* \sim 10^{10} M_\odot$ and/or at lower redshifts (see Table 4). Hence, we stress that adopting a (M_*, z) -dependent integration limit enables us to reach a globally higher incidence of radio-faint AGN, especially in lower- M_* galaxies, which a more conservative luminosity cut would likely miss. This explains our flatter M_* trend compared to that found in Kondapally et al. (2022), despite obtaining consistent numbers for the most massive galaxies.

6.3.2. The AGN luminosity in the radio-bright phase

Given the linear correlation between radio-AGN duty cycle and M_* (Fig. 11, left), we conclude that the super-linear M_* trend of the RAMS (Fig. 8, 1.41 ± 0.09) cannot be explained by a larger incidence of radio AGN alone, but it must be steepened by the fact that the typical AGN luminosity during a radio-bright phase is also higher in more massive than in less massive galaxies.

As a sanity check, we calculate the mean AGN luminosity during a radio-AGN phase ($\langle L_{1.4}^{\text{active}} \rangle$), again based on the best-fit RLF:

$$\langle L_{1.4}^{\text{active}} \rangle = \frac{\int_{L_{1.4}^{\text{SF}}}^{L_{1.4}^{\text{max}}} \Phi(L_{1.4}^{\text{AGN}}) \cdot L_{1.4}^{\text{AGN}} \cdot d(\log L_{1.4}^{\text{AGN}})}{\int_{L_{1.4}^{\text{SF}}}^{L_{1.4}^{\text{max}}} \Phi(L_{1.4}^{\text{AGN}}) \cdot d(\log L_{1.4}^{\text{AGN}})}. \quad (13)$$

This mean AGN luminosity is obtained as a weighted-average over the RLF. A similar formalism has been used in Aird et al. (2022) for X-ray AGN. Figure 11 (right) displays $\langle L_{1.4}^{\text{active}} \rangle$ across all (M_*, z) bins. Not surprisingly, we again observe a positive correlation with M_* (as $\propto M_*^{0.63 \pm 0.07}$), though weaker than that seen for the duty cycle, but a more significant evolution with redshift (as $\propto (1+z)^{1.96 \pm 0.27}$).

We acknowledge that part of the above M_* dependence could be explained by a roughly constant (kinetic) Eddington ratio and a fixed M_{BH}/M_* ratio. In this assumption, more massive galaxies would simply appear more (radio) luminous because they host more massive black holes, while the Eddington ratio is, in fact, constant. A more detailed analysis of the Eddington ratio distributions in radio AGN is postponed to a future work.

Based on these results, we argue that our empirical RAMS can be explained only by a combination of both a higher duty cycle, and a brighter radio-AGN phase in more massive and/or higher- z galaxies.

6.4. Mapping integrated AGN feedback across the galaxy population

The strong M_* stratification seen in the best-fit AGN RLF (Fig. 5) is reflected into the super-linear trend reported in our RAMS (Fig. 8). This behaviour corroborates the idea that radio AGN activity is enhanced relative to SF in the most massive SFGs, at least since $z \sim 3$. In this context, we compare our findings with recent radio and X-ray studies to discuss the role of star-forming host galaxies in triggering different types of AGN feedback.

6.4.1. Radio AGN at the crossroad: Towards a single AGN population in SFGs?

It has been argued that more massive galaxies undergo a higher radio-AGN duty cycle (e.g. Sabater et al. 2019, at $z < 0.3$), that is, that the fraction of the galaxy's lifetime in which a radio AGN is switched on, on average, strongly increases with M_* . At higher redshift ($z < 2.5$), a similar analysis has been presented in Kondapally et al. (2022), who exploited 150 MHz data from the first data release of the LOFAR Two-meter Sky Survey Deep Fields (LoTSS-Deep) survey (Tasse et al. 2021; Sabater et al. 2021; Kondapally et al. 2021; Duncan et al. 2021).

Following the historical dichotomy observed in the local Universe between radiatively efficient and inefficient radio AGN

(e.g. Best & Heckman 2012), Kondapally et al. (2022) separate AGN between low-excitation and high-excitation radio galaxies (LERGs and HERGs, respectively), studying their space density and incidence as a function of galaxy M_* . The authors classify LERGs based on the presence of (>0.7 dex, i.e. 3σ) radio-excess from a ridgeline set by Best et al. (in prep.) and the absence of optical-MIR signatures of AGN activity from SED-fitting. On the contrary, HERGs in their study consist of both radio-excess and SED-based AGN. The population of LERGs is further split among quiescent and star-forming LERGs, based on the sSFR of the host galaxy. Therefore, our radio-excess AGN in SF hosts are broadly consistent with the combined (SF-)HERG and SF-LERG population in their study. Interestingly, Kondapally et al. (2022) find that quiescent LERGs show systematically different properties and evolution compared to SF-LERGs, which instead resemble more closely the HERG population. For instance, they observe a flatter M_* dependence of the fraction of SF-LERGs (slope $\sim 1.37 \pm 0.57$), compared to that of quiescent LERGs (slope ~ 2.5 at all redshifts). The incidence of SF-LERGs increases with redshift following the evolution of the cold gas fraction ($\propto(1+z)^{-2.5}$; e.g. Tacconi et al. 2018), while quiescent LERGs do not show hints of evolution. This behaviour suggests that a different fuelling mechanism, likely associated with the availability of cold gas supply, is responsible for triggering SF-LERGs. The link with the HERG population is further reinforced by the self-similar fractions of SF-LERGs and HERGs at all redshifts.

Therefore, a plausible scenario is that HERGs and SF-LERGs are both triggered by cold gas accretion, which is more largely available in star-forming than in passive systems. Unsurprisingly, HERGs are mostly hosted in SFGs (Delvecchio et al. 2017), and X-ray stacking of radio-excess AGN reveals systematically ($>3\times$) higher BHARs in star-forming than in quiescent hosts (split by *NUVrJ* colours), at fixed radio power and redshift (Delvecchio et al. 2018; see also Magliocchetti et al. 2018). While it is true that local LERGs can also be fuelled by cold gas accretion (e.g. Ruffa et al. 2019), this is possibly driven by sporadic cold gas filaments due to radiative cooling of the hot halo (Hardcastle 2018; see Hardcastle & Croston 2020 for a review) or chaotic accretion (Gaspari et al. 2015, 2017), which in any case do not switch LERGs to ‘radiative mode’ AGN.

We reiterate that the main difference between (SF-)HERGs and SF-LERGs in the literature is the presence of excess emission in their X-ray-MIR-optical data relative to pure SF. However, SMBH accretion is a stochastic process and can potentially vary over $\lesssim 1$ Myr timescales (e.g. Schawinski et al. 2015), biasing these criteria to ‘instantaneous’ rather than ‘long-term’ AGN activity. Additionally, the above criteria are sensitive to the contrast of AGN versus host light, and hence they can potentially miss relatively faint AGN in highly star-forming systems. In this framework, splitting AGN by host-galaxy type (as well as $M_{*,z}$) might be more effective in capturing the historical SMBH fuelling due to cold gas accretion in the host, as compared to using conventional ‘single-epoch’ AGN diagnostics. In other words, radio-excess AGN in SF hosts (both SF-LERGs and SF-HERGs) and radio-excess AGN in quiescent hosts (both Q-LERGs and Q-HERGs) might behave broadly as two distinct AGN populations, irrespective of their on-going SMBH growth.

This simple picture, in which the long-term AGN activity is primarily modulated by the host-galaxy type (but internally varying with $M_{*,z}$) is further corroborated by the comparison between average radio and X-ray AGN activity in SFGs (Sect. 6.4.2).

6.4.2. Radio and X-ray AGN in SF hosts: A common fuelling scenario?

Our analysis has demonstrated that the incidence, evolution, integrated and mean luminosity of radio-excess AGN follow a strong trend with galaxy M_* . Here we discuss the quantitatively similar M_* dependence seen in the average X-ray properties of M_* -selected samples of SFGs.

A well-known study by Mullaney et al. (2012) put forward the idea that the average BHAR/SFR ratio in MS galaxies is both redshift and M_* -invariant at $M_* > 10^{10} M_\odot$ and $0.5 < z < 2.5$. This ‘hidden AGN MS’ lies at BHAR/SFR $\sim 10^{-3}$, consistent with a constant M_{BH}/M_* ratio, and thus in line with the empirical black hole–bulge mass scaling relations at $z \sim 0$ (Kormendy & Ho 2013). In the past decade, similar studies on larger samples and wider M_* ranges have been favouring an increasing BHAR/SFR ratio with M_* (Rodighiero et al. 2015; Yang et al. 2018; Aird et al. 2019; Bernhard et al. 2019; Carraro et al. 2020; Delvecchio et al. 2020), suggesting that BHAR is enhanced relative to SFR in more massive galaxies. The BHAR/SFR ratio evolves as $\propto M_*^{0.5-0.7}$ and, when adopting the same bending MS, is consistent with a redshift-invariant ratio. Pulling out the SFR term, the effective BHAR scales as $\propto M_*^{1.5}$, which is remarkably similar to that observed for the mean $\langle L_{1.4}^{AGN} \rangle$ (Sect. 5.2).

We note that X-ray based BHAR estimates are usually scaled from the mean X-ray luminosity $\langle L_X \rangle$ obtained either from X-ray stacking (e.g. Rodighiero et al. 2015; Yang et al. 2018; Carraro et al. 2020) or via Bayesian modelling of X-ray detections and non-detections (e.g. Aird et al. 2019; Bernhard et al. 2019). These techniques can smooth over short-term fluctuations due to AGN flickering (\lesssim Myr; e.g. Chen et al. 2013; Hickox et al. 2014). Recurrent AGN activity is also seen in the radio (e.g. Jurin et al. 2020; Brienza et al. 2021), albeit over longer timescales (10^{7-8} yr; see e.g. Konar et al. 2013). Therefore, our sample-average $\langle L_{1.4}^{AGN} \rangle$ measurements should be, if anything, less affected by AGN variability than from the X-rays. Nonetheless, radio-faint AGN emission can suffer more from ‘host galaxy dilution’ (Padovani et al. 2017), which we addressed by adopting a $(M_{*,z})$ -dependent IRRC (D21).

Figure 12 displays the logarithmic ratio between $\langle L_{1.4}^{AGN} \rangle$ and $\langle L_X \rangle$ (scaled to the same units) as a function of M_* , colour-coded by redshift. X-ray measurements are taken from the SFG sample of Carraro et al. (2020), for consistency with this work. Indeed, Carraro et al. (2020) stacked X-ray images of a M_* -selected sample, and identified SFGs via *NUVrJ* colour criteria as in this work. Other similar studies adopted different galaxy classifications based on sSFR (e.g. Aird et al. 2019; Ito et al. 2022), and thus the results are not fully comparable to those of Carraro et al. (2020), although they are qualitatively consistent with each other. A three-dimensional fitting of all $\langle L_{1.4}^{AGN} \rangle / \langle L_X \rangle$ data points with M_* and redshift yields a weak, poorly significant dependence on both M_* (0.14 ± 0.10) and redshift (0.00 ± 0.38). Imposing for simplicity a z -invariant relationship leads to an even weaker M_* trend (0.10 ± 0.10), and consistent with a constant ratio of $\log(\langle L_{1.4}^{AGN} \rangle / \langle L_X \rangle) \approx -3.5$ (black line in Fig. 12). The best-fit expression yields $\chi_{\text{red}}^2 = 1.48$. Intriguingly, this $(\langle L_{1.4}^{AGN} \rangle / \langle L_X \rangle)$ ratio resembles the ‘radio loudness’ parameter R_X typically used to separate ‘radio quiet’ from ‘radio loud’ AGN (Terashima & Wilson 2003; see also Lambrides et al. 2020).

Therefore, we conclude that mean X-ray and radio AGN luminosities in SFGs seem to evolve in a strikingly similar fashion with both M_* and redshift, suggesting that similar




Fig. 12. Logarithmic ratio between our $\langle L_{1.4}^{\text{AGN}} \rangle$ measurements and the $\langle L_X \rangle$ obtained from X-ray stacking (Carraro et al. 2020), as a function of M_* , coloured by redshift. The black line indicates the best-fit ratio with M_* , by imposing a z -invariant trend, which returns a M_* -slope of 0.10 ± 0.10 . The grey shaded area marks the $\pm 1\sigma$ confidence interval after propagating the parameter uncertainties. See Sect. 6.4.2 for details.

mechanisms trigger and sustain long-term X-ray and radio AGN activity in SFGs. Such a similar behaviour with M_* extends beyond average measurements, encompassing the global evolution of the AGN luminosity function. Indeed, the characteristic knee luminosity L^* of the X-ray AGN luminosity function (XLF), at fixed redshift, has been reported to increase with M_* in a qualitatively similar fashion to this work, although partly induced by a M_* -invariant characteristic Eddington ratio (λ_{Edd} , or $\text{sBHAR} \propto L_X/M_*$; Aird et al. 2012) assumed in the XLF modelling (Delvecchio et al. 2020). Moreover, the integrated AGN luminosity density in the X-rays is dominated by AGN in MS galaxies with $10.5 < \log(M_*/M_\odot) < 11$, showing a peak at $z \sim 2$ (e.g. Delvecchio et al. 2020). This is fully consistent with our reported evolutionary trend of the kinetic AGN luminosity density in SFGs (Sect. 5.1). Similarly, the average $\langle L_X \rangle/\text{SFR}$ ratio, at fixed M_* , evolves with redshift following the MS relation, closely resembling the evolution of $\langle L_{1.4}^{\text{AGN}} \rangle/\text{SFR}$ (Sect. 6.3).

Such a degree of consistency in the evolution of the X-ray and radio AGN population in SFGs is surprising given the completely different methodologies adopted in each of the above studies. However, these results add to the evidence discussed earlier (Sect. 6.4.1) that the availability of cold gas supply modulates the long-term fuelling onto the SMBH, broadly independent of the (instantaneous) AGN diagnostics at a given wavelength. The small overlap (10–15%) previously reported between X-ray and radio AGN populations (e.g. Goulding et al. 2014; Azadi et al. 2015; Delvecchio et al. 2017; Ji et al. 2022), even for the same galaxy type, could be attributed to intrinsic AGN variability, which globally reaches a duty cycle of 10% in the most massive ($M_* \gtrsim 10^{11} M_\odot$) SFGs (see Fig. 11, left), consistently with the X-ray–radio AGN overlap in M_* -matched galaxies.

We reiterate that both $\langle L_{1.4}^{\text{AGN}} \rangle$ and $\langle L_X \rangle$ are averaged across the entire SFG population, at each M_* and redshift, thus smoothing over the entire AGN duty cycle. We note that studying galaxies at fixed redshift, stellar mass and galaxy type enables us to roughly lock the expected SFR of the host, and hence

the baryonic accretion rate modulated by the dark matter halo (Daddi et al. 2022a). At fixed redshift, more massive SFGs are embedded in more massive halos, and hence it is plausible to expect that stochastic gas accretion triggers more frequent and/or more luminous AGN activity.

In support of our arguments, widespread X-ray and radio AGN activity in massive galaxies ($M_* > 10^{10} M_\odot$) has been recently reported in Ito et al. (2022). They exploited the most recent optical-to-MIR photometry from the COSMOS2020 catalogue (Weaver et al. 2022), stacking radio (VLA-3 GHz; Smolčić et al. 2017b) and X-ray (Chandra; Civano et al. 2016; Marchesi et al. 2016) images of the underlying M_* -selected sample in bins of M_* and redshift. Though the strongest M_* dependence was found for quiescent galaxies, also in SFGs Ito et al. (2022) found a rising trend of both $\langle L_X \rangle$ and $\langle L_{1.4}^{\text{AGN}} \rangle$ with M_* , out to $z \sim 3$.

Hence, in this simple framework the average radio and X-ray nuclear activity might broadly trace each other because they are fuelled via similar mechanisms (i.e. stochastic cold gas accretion). This ‘hidden’ connection, however, does not imply that radio and X-ray AGN are switched on at the same time. Clearly, the small overlap between detected X-ray and radio AGN suggests that these phases of AGN feedback are broadly unsynchronised, and only smoothing over the SMBH duty cycle reveals their long-term connection.

As a final caveat, by ‘cold gas accretion’ we are neither implying that gas fuelling of SF and black hole growth happens at the same time, nor that the physical mechanisms funneling the gas inwards are necessarily the same. We simply mean that the long-term growth of both black holes and galaxies is possibly modulated by the amount of usable gas already in the host (Harrison 2017), irrespective of the internal-external mechanisms that did channel the gas inwards. Because more massive and more distant SFGs have more cold gas available than the rest of the population (Tacconi et al. 2018; 2020; Liu et al. 2019; Wang et al. 2022), we interpret the higher average (radio and X-ray) AGN activity in these galaxies as broadly driven by larger supply of cold gas in the host, on statistical basis. However, we also acknowledge the fact that SMBH and SF fuelling could be further promoted by other aspects, such as the compactness or geometry of SF regions, and the morphology of the host galaxy (e.g. Gómez-Guijarro et al. 2022; Aird et al. 2022), which might affect the efficiency of channelling gas towards the smallest ($\ll \text{kpc}$) scales, and thus the effective gas accretion rate. For this reason, we caution that understanding the higher AGN-to-SF luminosity ratios seen in more massive SFGs (Figs. 9 and 10) might require additional, possibly non-linear, processes happening during cold gas accretion.

7. Summary and conclusions

We have presented a novel approach for assessing the relationship between radio AGN activity and galaxy stellar mass across M_* -selected SFGs. In particular, we performed this analysis on radio-excess AGN, factoring in the statistical contribution of radio-faint AGN within the IRRC, which is critical for computing a representative sample-averaged radio AGN power across the galaxy population. To achieve this goal, we exploited deep VLA-COSMOS 3-GHz data to identify bona fide radio-excess AGN, that is, AGN that show ($>2\sigma$) excess radio emission relative to that expected from pure SF (i.e. the IRRC; see Sect. 1), at each (M_*, z) . From this, we built the 1.4 GHz luminosity function (AGN RLF) of radio-excess AGN in SFGs at $0.1 \leq z \leq 4.5$, following previous works (e.g. Smolčić et al. 2017c;

Novak et al. 2018; Ceraj et al. 2018; Butler et al. 2019; Kono & Takeuchi 2021). Then, for the first time, we decomposed the AGN RLF at each redshift into different M_* bins over the range $9 < \log(M_*/M_\odot) < 12$, fitting and integrating each luminosity function down to the minimum $L_{1.4}^{\text{SF}}$ set by the IRRC at the same (M_*, z) .

Our main results can be summarised as follows:

1. The integrated radio-AGN luminosity density across SFGs is mostly driven by massive galaxies with $10.5 < \log(M_*/M_\odot) < 11$ and peaks at $z \sim 2$ (Sect. 5.1).
2. When averaging this cumulative radio AGN power across all M_* -selected galaxies at each (M_*, z) , we obtain a super-linear (slope of 1.41 ± 0.09) RAMS that links the mean (i.e. time-averaged across the galaxy's life cycle) radio AGN power ($\langle L_{1.4}^{\text{AGN}} \rangle$) and galaxy M_* from $z \sim 3$ (Sect. 5.2).
3. The mean radio AGN power $\langle L_{1.4}^{\text{AGN}} \rangle$ at fixed M_* evolves with redshift in a similar fashion to the MS relation ($\propto (1+z)^{2.5}$; see e.g. Speagle et al. 2014), suggesting that long-term radio AGN activity and galaxy SF have proceeded at a similar pace through cosmic time, at least since $z \sim 3$ (Sect. 5.2).
4. The comparison between radio emission from AGN versus SF reveals that AGN emission is typically dominant over SF only at $M_* > 10^{11} M_\odot$ (Sect. 6.2). This is also the only M_* range in which current deep radio-continuum surveys (e.g. VLA-COSMOS 3 GHz; Smolčić et al. 2017b) are able to formally detect individual radio AGN with luminosity $\langle L_{1.4}^{\text{AGN}} \rangle$. Probing typical radio AGN activity in less massive galaxies requires deeper, new-generation radio surveys (e.g. SKA1-MID or ngVLA), ideally complemented with VLBI techniques.
5. We find that the bulk of radio AGN activity originates from relatively short phases in which the AGN emission is dominant over SF (Sect. 5.2). This is what still allows us to calculate the mean $\langle L_{1.4}^{\text{AGN}} \rangle$ despite this being on average (i.e. across the galaxy's lifetime) sub-dominant compared to $\langle L_{1.4}^{\text{SF}} \rangle$.
6. The super-linear dependence of $\langle L_{1.4}^{\text{AGN}} \rangle$ on M_* suggests that radio AGN activity is strongly enhanced in more massive SFGs, as compared to the shallower evolution of galaxy SF along the star-forming MS (e.g. Schreiber et al. 2015).
7. We dissect the effects of the radio-AGN duty cycle (Sect. 6.3.1) and AGN luminosity in the radio-bright phase (Sect. 6.3.2) to explain the shape and evolution of the RAMS. Our analysis suggests that more massive (and higher- z) galaxies have higher mean $\langle L_{1.4}^{\text{AGN}} \rangle$ due to a combination of both a higher duty cycle and an intrinsically brighter radio-AGN phase.
8. Intriguingly, the super-linear trend of $\langle L_{1.4}^{\text{AGN}} \rangle$ with M_* closely resembles the evolution of X-ray AGN activity with M_* (e.g. Carraro et al. 2020; Ito et al. 2022), possibly suggesting a common fuelling scenario (Sect. 6.4.2). While our analysis favours a long-term X-ray–radio connection, these AGN phases are likely ‘unsynchronised’ due to a relatively short X-ray–radio AGN duty cycle – $\sim 10\%$ even in the most massive galaxies ($M_* \gtrsim 10^{11} M_\odot$) – which is consistent with the small X-ray–radio AGN overlap (10–15%) seen in M_* -matched SFGs.

Our findings support the idea that the global feeding and feedback cycle of SMBHs in SFGs is enhanced over SF in more massive systems, but it is broadly independent of the (instantaneous) AGN diagnostics at a given wavelength. Smoothing over the entire SMBH duty cycle, as done here on a statistical basis, can reveal the historical, long-term SMBH growth in their hosts.

Acknowledgements. We thank the anonymous referee for a positive report. I.D. thanks K. Ito for useful discussions. The Cosmic Dawn Center (DAWN) is funded by the Danish National Research Foundation under grant No. 140. S.J. acknowledges the Villum Fonden research grants 37440, 13160 and the financial support from European Union's Horizon research and innovation program under the Marie Skłodowska-Curie grant agreement No. 101060888. C.S. acknowledges financial support from the Italian Ministry of University and Research - Project Proposal CIR01_00010.

References

- Aird, J., Coil, A. L., Moustakas, J., et al. 2012, *ApJ*, 746, 90
Aird, J., Coil, A. L., Moustakas, J., et al. 2013, *ApJ*, 775, 41
Aird, J., Coil, A. L., & Georgakakis, A. 2018, *MNRAS*, 474, 1225
Aird, J., Coil, A. L., & Georgakakis, A. 2019, *MNRAS*, 484, 4360
Aird, J., Coil, A. L., & Kocevski, D. D. 2022, *MNRAS*, 515, 4860
Algera, H. S. B., Hodge, J. A., Riechers, D., et al. 2021, *ApJ*, 912, 73
Allevato, V., Shankar, F., Marsden, C., et al. 2021, *ApJ*, 916, 34
Azadi, M., Aird, J., Coil, A. L., et al. 2015, *ApJ*, 806, 187
Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013, *ApJ*, 770, 57
Bell, E. F. 2003, *ApJ*, 586, 794
Bernhard, E., Grimmert, L. P., Mullaney, J. R., et al. 2019, *MNRAS*, 483, L52
Best, P. N., & Heckman, T. M. 2012, *MNRAS*, 421, 1569
Best, P. N., Kauffmann, G., Heckman, T. M., et al. 2005, *MNRAS*, 362, 25
Bîrzan, L., Rafferty, D. A., McNamara, B. R., Wise, M. W., & Nulsen, P. E. J. 2004, *ApJ*, 607, 800
Bîrzan, L., McNamara, B. R., Nulsen, P. E. J., Carilli, C. L., & Wise, M. W. 2008, *ApJ*, 686, 859
Bonato, M., Prandoni, I., De Zotti, G., et al. 2021, *A&A*, 656, A48
Bondi, M., Zamorani, G., Ciliegi, P., et al. 2018, *A&A*, 618, L8
Braun, R., Bonaldi, A., Bourke, T., Keane, E., & Wagg, J. 2019, arXiv e-prints [arXiv:1912.12699]
Brienza, M., Shimwell, T. W., de Gasperin, F., et al. 2021, *Nat. Astron.*, 5, 1261
Butler, A., Huynh, M., Kapińska, A., et al. 2019, *A&A*, 625, A111
Caplar, N., Lilly, S. J., & Trakhtenbrot, B. 2015, *ApJ*, 811, 148
Caplar, N., Lilly, S. J., & Trakhtenbrot, B. 2018, *ApJ*, 867, 148
Carraro, R., Rodighiero, G., Cassata, P., et al. 2020, *A&A*, 642, A65
Cavagnolo, K. W., McNamara, B. R., Nulsen, P. E. J., et al. 2010, *ApJ*, 720, 1066
Ceraj, L., Smolčić, V., Delvecchio, I., et al. 2018, *A&A*, 620, A192
Ceraj, L., Smolčić, V., Delvecchio, I., et al. 2020, *A&A*, 642, A125
Chabrier, G. 2003, *PASP*, 115, 763
Chen, C.-T. J., Hickox, R. C., Alberts, S., et al. 2013, *ApJ*, 773, 3
Civano, F., Marchesi, S., Comastri, A., et al. 2016, *ApJ*, 819, 62
Collet, C., Nesvadba, N. P. H., De Breuck, C., et al. 2016, *A&A*, 586, A152
Condon, J. J. 1992, *ARA&A*, 30, 575
Costa, T., Rosdahl, J., Sijacki, D., & Haehnelt, M. G. 2018, *MNRAS*, 479, 2079
Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *MNRAS*, 365, 11
Croton, D. J., Stevens, A. R. H., Tonini, C., et al. 2016, *ApJS*, 222, 22
Daddi, E., Delvecchio, I., Dimauro, P., et al. 2022a, *A&A*, 661, L7
Daddi, E., Rich, R. M., Valentino, F., et al. 2022b, *ApJ*, 926, L21
Daly, R. A., Sprinkle, T. B., O’Dea, C. P., Kharb, P., & Baum, S. A. 2012, *MNRAS*, 423, 2498
Davidzon, I., Ilbert, O., Laigle, C., et al. 2017, *A&A*, 605, A70
Decarli, R., Aravena, M., Boogaard, L., et al. 2020, *ApJ*, 902, 110
Dekel, A., Zolotov, A., Tweed, D., et al. 2013, *MNRAS*, 435, 999
Del Moro, A., Alexander, D. M., Mullaney, J. R., et al. 2013, *A&A*, 549, A59
Delhaize, J., Smolčić, V., Delvecchio, I., et al. 2017, *A&A*, 602, A4
Delvecchio, I., Smolčić, V., Zamorani, G., et al. 2017, *A&A*, 602, A3
Delvecchio, I., Smolčić, V., Zamorani, G., et al. 2018, *MNRAS*, 481, 4971
Delvecchio, I., Daddi, E., Aird, J., et al. 2020, *ApJ*, 892, 17
Delvecchio, I., Daddi, E., Sargent, M. T., et al. 2021, *A&A*, 647, A123
Di Matteo, T., Springel, V., & Hernquist, L. 2005, *Nature*, 433, 604
Donley, J. L., Rieke, G. H., Rigby, J. R., & Pérez-González, P. G. 2005, *ApJ*, 634, 169
Duncan, K. J., Kondapally, R., Brown, M. J. I., et al. 2021, *A&A*, 648, A4
Elbaz, D., Dickinson, M., Hwang, H. S., et al. 2011, *A&A*, 533, A119
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
Förster Schreiber, N. M., & Wuyts, S. 2020, *ARA&A*, 58, 661
Gaspari, M., Brighenti, F., & Temi, P. 2015, *A&A*, 579, A62
Gaspari, M., Temi, P., & Brighenti, F. 2017, *MNRAS*, 466, 677
Gehrels, N. 1986, *ApJ*, 303, 336
Girdhar, A., Harrison, C. M., Mainieri, V., et al. 2022, *MNRAS*, 512, 1608
Gitti, M., Brighenti, F., & McNamara, B. R. 2012, *Adv. Astron.*, 2012, 950641
Gobat, R., Daddi, E., Magdis, G., et al. 2018, *Nat. Astron.*, 2, 239
Godfrey, L. E. H., & Shabala, S. S. 2016, *MNRAS*, 456, 1172
Gómez-Guijarro, C., Elbaz, D., Xiao, M., et al. 2022, *A&A*, 659, A196

- Goulding, A. D., Forman, W. R., Hickox, R. C., et al. 2014, *ApJ*, 783, 40
- Grimmett, L. P., Mullaney, J. R., Jin, S., et al. 2019, *MNRAS*, 487, 4071
- Grylls, P. J., Shankar, F., Zanisi, L., & Bernardi, M. 2019, *MNRAS*, 483, 2506
- Hardcastle, M. J. 2018, *MNRAS*, 475, 2768
- Hardcastle, M. J., & Croston, J. H. 2020, *New Astron. Rev.*, 88, 101539
- Harrison, C. M. 2017, *Nat. Astron.*, 1, 0165
- Heckman, T. M., & Best, P. N. 2014, *ARA&A*, 52, 589
- Heesen, V., Staffehl, M., Basu, A., et al. 2022, *A&A*, 664, A83
- Helou, G., Soifer, B. T., & Rowan-Robinson, M. 1985, *ApJ*, 298, L7
- Herrera Ruiz, N., Middelberg, E., Deller, A., et al. 2017, *A&A*, 607, A132
- Herrera Ruiz, N., Middelberg, E., Deller, A., et al. 2018, *A&A*, 616, A128
- Heywood, I., Jarvis, M. J., Hale, C. L., et al. 2022, *MNRAS*, 509, 2150
- Hickox, R. C., Mullaney, J. R., Alexander, D. M., et al. 2014, *ApJ*, 782, 9
- Ibar, E., Ivison, R. J., Biggs, A. D., et al. 2009, *MNRAS*, 397, 281
- Ibar, E., Ivison, R. J., Best, P. N., et al. 2010, *MNRAS*, 401, L53
- Ito, K., Tanaka, M., Miyaji, T., et al. 2022, *ApJ*, 929, 53
- Jarvis, M., Taylor, R., Agudo, I., et al. 2016, in *MeerKAT Science: On the Pathway to the SKA (MeerKAT2016)*, 6
- Jarvis, M. E., Harrison, C. M., Thomson, A. P., et al. 2019, *MNRAS*, 485, 2710
- Jarvis, M. E., Harrison, C. M., Mainieri, V., et al. 2020, *MNRAS*, 498, 1560
- Ji, Z., Giavalisco, M., Kirkpatrick, A., et al. 2022, *ApJ*, 925, 74
- Jin, S., Daddi, E., Liu, D., et al. 2018, *ApJ*, 864, 56
- Jones, M. L., Hickox, R. C., Mutch, S. J., et al. 2019, *ApJ*, 881, 110
- Jurlin, N., Morganti, R., Brienza, M., et al. 2020, *A&A*, 638, A34
- Kennicutt, R. C., & Evans, N. J. 2012, *ARA&A*, 50, 531
- Konar, C., Hardcastle, M. J., Jamroz, M., & Croston, J. H. 2013, *MNRAS*, 430, 2137
- Kondapally, R., Best, P. N., Hardcastle, M. J., et al. 2021, *A&A*, 648, A3
- Kondapally, R., Best, P. N., Cochrane, R. K., et al. 2022, *MNRAS*, 513, 3742
- Kono, K. T., & Takeuchi, T. T. 2021, arXiv e-prints [arXiv:2112.01672]
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, 51, 511
- Laigle, C., McCracken, H. J., Ilbert, O., et al. 2016, *ApJS*, 224, 24
- Lambrides, E. L., Chiaberge, M., Heckman, T., et al. 2020, *ApJ*, 897, 160
- Lee, N., Sanders, D. B., Casey, C. M., et al. 2015, *ApJ*, 801, 80
- Leslie, S. K., Schinnerer, E., Liu, D., et al. 2020, *ApJ*, 899, 58
- Liu, D., Daddi, E., Dickinson, M., et al. 2018, *ApJ*, 853, 172
- Liu, D., Schinnerer, E., Groves, B., et al. 2019, *ApJ*, 887, 235
- Liu, D., Daddi, E., Schinnerer, E., et al. 2021, *ApJ*, 909, 56
- Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
- Magdis, G. E., Gobat, R., Valentino, F., et al. 2021, *A&A*, 647, A33
- Magliocchetti, M., Popesso, P., Brusa, M., & Salvato, M. 2018, *MNRAS*, 473, 2493
- Magnelli, B., Ivison, R. J., Lutz, D., et al. 2015, *A&A*, 573, A45
- Maini, A., Prandoni, I., Norris, R. P., Giovannini, G., & Spitler, L. R. 2016, *A&A*, 589, L3
- Marchesi, S., Civano, F., Elvis, M., et al. 2016, *ApJ*, 817, 34
- Marshall, H. L. 1985, *ApJ*, 299, 109
- Matthews, A. M., Condon, J. J., Cotton, W. D., & Mauch, T. 2021, *ApJ*, 914, 126
- Mauch, T., & Sadler, E. M. 2007, *MNRAS*, 375, 931
- McAlpine, K., Jarvis, M. J., & Bonfield, D. G. 2013, *MNRAS*, 436, 1084
- McNamara, B. R., & Nulsen, P. E. J. 2007, *ARA&A*, 45, 117
- Mendez, A. J., Coil, A. L., Aird, J., et al. 2013, *ApJ*, 770, 40
- Merloni, A., & Heinz, S. 2007, *MNRAS*, 381, 589
- Molnár, D. C., Sargent, M. T., Leslie, S., et al. 2021, *MNRAS*, 504, 118
- Mullaney, J. R., Alexander, D. M., Goulding, A. D., & Hickox, R. C. 2011, *MNRAS*, 414, 1082
- Mullaney, J. R., Daddi, E., Béthermin, M., et al. 2012, *ApJ*, 753, L30
- Murphy, E. J., Condon, J. J., Schinnerer, E., et al. 2011, *ApJ*, 737, 67
- Muxlow, T. W. B., Thomson, A. P., Radcliffe, J. F., et al. 2020, *MNRAS*, 495, 1188
- Nesvadba, N. P. H., De Breuck, C., Lehnert, M. D., Best, P. N., & Collet, C. 2017, *A&A*, 599, A123
- Noeske, K. G., Weiner, B. J., Faber, S. M., et al. 2007, *ApJ*, 660, L43
- Novak, G. S., Ostriker, J. P., & Ciotti, L. 2011, *ApJ*, 737, 26
- Novak, M., Smolčić, V., Delhaize, J., et al. 2017, *A&A*, 602, A5
- Novak, M., Smolčić, V., Schinnerer, E., et al. 2018, *A&A*, 614, A47
- Oke, J. B. 1974, *ApJS*, 27, 21
- O'Sullivan, E., Giacintucci, S., David, L. P., et al. 2011, *ApJ*, 735, 11
- Padovani, P. 2016, *A&ARv*, 24, 13
- Padovani, P., Bonzini, M., Kellermann, K. I., et al. 2015, *MNRAS*, 452, 1263
- Padovani, P., Alexander, D. M., Assef, R. J., et al. 2017, *A&ARv*, 25, 2
- Panessa, F., Baldi, R. D., Laor, A., et al. 2019, *Nat. Astron.*, 3, 387
- Paragi, Z., Godfrey, L., Reynolds, C., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 143
- Popesso, P., Concas, A., Cresci, G., et al. 2022, *MNRAS*, accepted [arXiv:2203.10487]
- Prandoni, I., & Seymour, N. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 67
- Radcliffe, J. F., Garrett, M. A., Muxlow, T. W. B., et al. 2018, *A&A*, 619, A48
- Radcliffe, J. F., Barthel, P. D., Garrett, M. A., et al. 2021a, *A&A*, 649, L9
- Radcliffe, J. F., Barthel, P. D., Thomson, A. P., et al. 2021b, *A&A*, 649, A27
- Rinaldi, P., Caputi, K. I., van Mierlo, S. E., et al. 2022, *ApJ*, 930, 128
- Rodighiero, G., Brusa, M., Daddi, E., et al. 2015, *ApJ*, 800, L10
- Ruffa, I., Prandoni, I., Laing, R. A., et al. 2019, *MNRAS*, 484, 4239
- Sabater, J., Best, P. N., Hardcastle, M. J., et al. 2019, *A&A*, 622, A17
- Sabater, J., Best, P. N., Tasse, C., et al. 2021, *A&A*, 648, A2
- Sadler, E. M., Cannon, R. D., Mauch, T., et al. 2007, *MNRAS*, 381, 211
- Saintonge, A., Catinella, B., Tacconi, L. J., et al. 2017, *ApJS*, 233, 22
- Sanders, R. L., Shapley, A. E., Jones, T., et al. 2022, *ApJ*, submitted [arXiv:2204.06937]
- Schawinski, K., Koss, M., Berney, S., & Sartori, L. F. 2015, *MNRAS*, 451, 2517
- Schmidt, M. 1968, *ApJ*, 151, 393
- Schreiber, C., Pannella, M., Elbaz, D., et al. 2015, *A&A*, 575, A74
- Shankar, F., Bernardi, M., Richardson, K., et al. 2019, *MNRAS*, 485, 1278
- Smith, D. J. B., Haskell, P., Gürkan, G., et al. 2021, *A&A*, 648, A6
- Smolčić, V., Zamorani, G., Schinnerer, E., et al. 2009, *ApJ*, 696, 24
- Smolčić, V., Delvecchio, I., Zamorani, G., et al. 2017a, *A&A*, 602, A2
- Smolčić, V., Novak, M., Bondi, M., et al. 2017b, *A&A*, 602, A1
- Smolčić, V., Novak, M., Delvecchio, I., et al. 2017c, *A&A*, 602, A6
- Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, 214, 15
- Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, *ApJS*, 148, 175
- Spingola, C., Dallacasa, D., Belladitta, S., et al. 2020, *A&A*, 643, L12
- Sudoh, T., Linden, T., & Beacom, J. F. 2021, *Phys. Rev. D*, 103, 083017
- Sweijen, F., van Weeren, R. J., Röttgering, H. J. A., et al. 2022, *Nat. Astron.*, 6, 350
- Tacconi, L. J., Genzel, R., Saintonge, A., et al. 2018, *ApJ*, 853, 179
- Tacconi, L. J., Genzel, R., & Sternberg, A. 2020, *ARA&A*, 58, 157
- Tasse, C., Shimwell, T., Hardcastle, M. J., et al. 2021, *A&A*, 648, A1
- Terashima, Y., & Wilson, A. S. 2003, *ApJ*, 583, 145
- van der Vlugt, D., Hodge, J. A., Algera, H. S. B., et al. 2022, *ApJ*, submitted [arXiv:2204.04167]
- Venturi, G., Cresci, G., Marconi, A., et al. 2021, *A&A*, 648, A17
- Walter, F., Carilli, C., Neeleman, M., et al. 2020, *ApJ*, 902, 111
- Wang, T.-M., Magnelli, B., Schinnerer, E., et al. 2022, *A&A*, 660, A142
- Weaver, J. R., Kauffmann, O. B., Ilbert, O., et al. 2022, *ApJS*, 258, 11
- Weigel, A. K., Schawinski, K., Caplar, N., et al. 2017, *ApJ*, 845, 134
- Willott, C. J., Rawlings, S., Blundell, K. M., & Lacy, M. 1999, *MNRAS*, 309, 1017
- Yang, G., Chen, C. T. J., Vito, F., et al. 2017, *ApJ*, 842, 72
- Yang, G., Brandt, W. N., Vito, F., et al. 2018, *MNRAS*, 475, 1887
- Yun, M. S., Reddy, N. A., & Condon, J. J. 2001, *ApJ*, 554, 803
- Zubovas, K., & King, A. R. 2012, *MNRAS*, 426, 2751

Appendix A: Extras on sample selection

We display some relevant distributions for our VLA 3 GHz detections, before and after applying the selection cuts outlined in Table 1. Particularly, Fig. A.1 shows VLA sources with optical-NIR counterpart (Laigle et al. 2016) within the Ultra Deep Visible and Infrared Survey Telescope for Astronomy (UltraVISTA) 1.5 deg² area (empty circles), as well as its subset classified as ‘star-forming’ (red crosses). From top to bottom, the panels show the distribution with redshift of $L_{1.4}$, L_{IR} , and galaxy M_* , respectively. As discussed in Appendix D, the removal of $NUVrJ$ -passive galaxies reduces our sample by a factor of 2–3 at $z \lesssim 1$. At these redshifts, such a cut leaves us with higher L_{IR} (middle panel) and slightly lower $L_{1.4}$ (top panel) and M_* (bottom panel) than the parent VLA sample. We refer the reader to Sect. 4.4 for a detail explanation on why $NUVrJ$ -passive galaxies were removed.

In Figure A.2, we show the L_{IR} -versus- $L_{1.4}$ distributions of our final sample of 5,658 radio-detected ($S/N > 5$ at 3 GHz) SFGs across 1.5 deg², and within the range $0.1 \leq z \leq 4.5$ and $9 \leq \log(M_*/M_\odot) \leq 12$ (see Table 1). Dashed lines mark constant q_{IR} lines as a reference, including the local value ($q_{IR} = 2.64$; Bell 2003). At high M_\odot , both L_{IR} and $L_{1.4}$ increase due to galaxies becoming more star-forming. Moreover, at high luminosities the distributions appear to shift to lower q_{IR} : this is caused by a redshift trend within our flux-limited VLA sample (i.e. q_{IR} apparently decreases with redshift and, hence, with luminosity; see D21). Controlling for such internal redshift dependence is needed to coherently compare the q_{IR} distributions in different M_* bins. This is why in Sect. 3.1 (and in Fig. 2) we rescaled each observed q_{IR} by the corresponding q_{IRRC} at the same (M_*, z) .




Fig. A.1. Distributions of $L_{1.4}$ (top), L_{IR} (middle), and galaxy M_* (bottom) of our sample, as a function of redshift. Grey empty circles are VLA sources with an optical-NIR counterpart (Laigle et al. 2016) within the UltraVISTA 1.5 deg² area, while red crosses mark the subset classified as ‘star-forming’ based on $NUVrJ$ criteria. Dotted lines in the bottom panel enclose the M_* - z space of the final sample (5,658 sources).




Fig. A.2. Distribution of our final sample of all 3 GHz detection classified as star-forming (5,658; see Table 1), as a function of L_{IR} and $L_{1.4}$ and split into different M_* bins (increasing from left to right). Dashed lines mark constant q_{IR} lines as a reference.

Appendix B: Statistical corrections to the AGN RLF

B.1. Correction for classification purity

The comoving number density $\Phi(L)$ obtained in four redshift bins is shown in the upper panels of Fig. B.1 (grey open squares). As highlighted previously, the 2σ threshold at $\Delta q_{IRRC} = -0.44$ dex was designed to minimise AGN contamination within ‘normal’ SFGs (D21). Reversing this approach for identifying radio-excess AGN is, however, at the expense of AGN purity (i.e. higher galaxy contamination). Correcting for misclassified AGN can be done on statistical basis as explained in Sect. 3.2 and Fig. 2. Briefly, instead of counting each radio-excess AGN as unity, we weigh it as the corresponding f_N^{AGN} in the $1/V_{max}$ calculation. This approach enables us to statistically correct for source misclassifications by assigning each object a purity $0 \leq f_N^{AGN} \leq 1$ of being a radio-excess AGN.

Corrected values of $\Phi(L)$ are marked in Fig. B.1 as red starred symbols. As expected, the number density remains unchanged at the bright end, while at low $L_{1.4}$ it drops by an amount proportional to $1 - f_N^{AGN}$, that is, ≈ 30 – 40% . Therefore, this correction is quite important to clean up our AGN sample. We clarify that each object is treated as both an AGN and a galaxy: in the AGN RLF each object is counted as f_N^{AGN} , while in the complementary galaxy RLF (not shown here), the same object is counted as $(1 - f_N^{AGN})$. As a reference, the vertical dotted line in each z -bin indicates the 1.4 GHz luminosity threshold corresponding to $+2\sigma$ from the IRRC. This $L_{1.4}$ threshold is converted from the Δq_{IRRC} space, by taking for simplicity the median L_{IR} and the median M_* of the sample in that z -bin. This is for visual purposes in Fig B.1. Instead, in practice the correction for AGN classification purity is applied to each object based on its observed Δq_{IRRC} value, irrespective of its $L_{1.4}$.

It is likely that radio-excess AGN are under-represented (or at least their fraction is highly uncertain) at radio luminosities below our $+2\sigma$ $L_{1.4}$ threshold, since we assumed that the peak of the Δq_{IRRC} distribution is entirely made by SFGs (see also Fig. 2), while radio AGN fill the residual part of the distribution. For this reason, throughout this manuscript we disregard all data points placed at $\Delta q_{IRRC} > -0.44$ dex (i.e. the -2σ boundary from the IRRC). This interval is indicated as the shaded area in Fig. B.1. After applying the above corrections, the final num-

bers and fractions relative to the parent VLA-COSMOS 3GHz sample are reported in each z -bin. See the text for details.

B.2. Correction for $L_{1.4}$ purity

The correction for AGN classification purity acts solely on the counting, not on the luminosity of each object. Thus, following the motivation presented in Sect. 3.2, we attempt at decomposing the total 1.4 GHz luminosity of each source between AGN- and galaxy-driven contributions.

As done in Ceraj et al. (2018) and Fig. 2 (dashed grey lines), we compute the AGN-related fraction at 1.4 GHz, $f_{1.4}^{AGN}$. Then we multiply each $L_{1.4}$ by the corresponding $f_{1.4}^{AGN}$ to obtain the AGN-related luminosity (i.e. $L_{1.4}^{AGN}$). The global effect on the AGN RLF is shown in Fig. B.1 (bottom panels). This correction for $L_{1.4}$ re-distributes the former dataset (grey stars) to lower luminosities. Specifically, except for the brightest $L_{1.4}$ bins that are unchanged, the (pure) AGN-related $\Phi(L)$ steepens in the faint end at the expense of bins at moderate $L_{1.4}$. Of course, in the faintest sources, this step shifts $L_{1.4}^{AGN}$ formally below the luminosity limit of the survey (or within the $\pm 2\sigma$ locus of the IRRC, grey shaded area), in which case the source is removed from our final sample.

In this study, we refrain from re-calculating the V_{max} of each source using the new AGN-scaled flux density, since the detectability of each object remains bounded to its combined AGN+galaxy flux density. Uncertainties on $\Phi(L)$ after applying the afore-mentioned corrections are discussed in Sect. 4.2.

Finally, we double-check that correcting only for AGN $L_{1.4}$ purity (i.e. without applying an AGN classification purity), but extending this to all radio-detected galaxies, returns a consistent luminosity function, although about 20% more AGN just above the $+2\sigma$ $L_{1.4}$ threshold (dotted black line in Fig. B.1). This is because a correction for AGN $L_{1.4}$ purity alone does not distinguish between a radio-excess AGN and a $>2\sigma$ outlier SFG, and hence we go ahead with this two-step correction.

B.3. AGN completeness and final numbers

A potential caveat of our approach is related to the fact that we correct for AGN classification and $L_{1.4}$ purity based on the offset

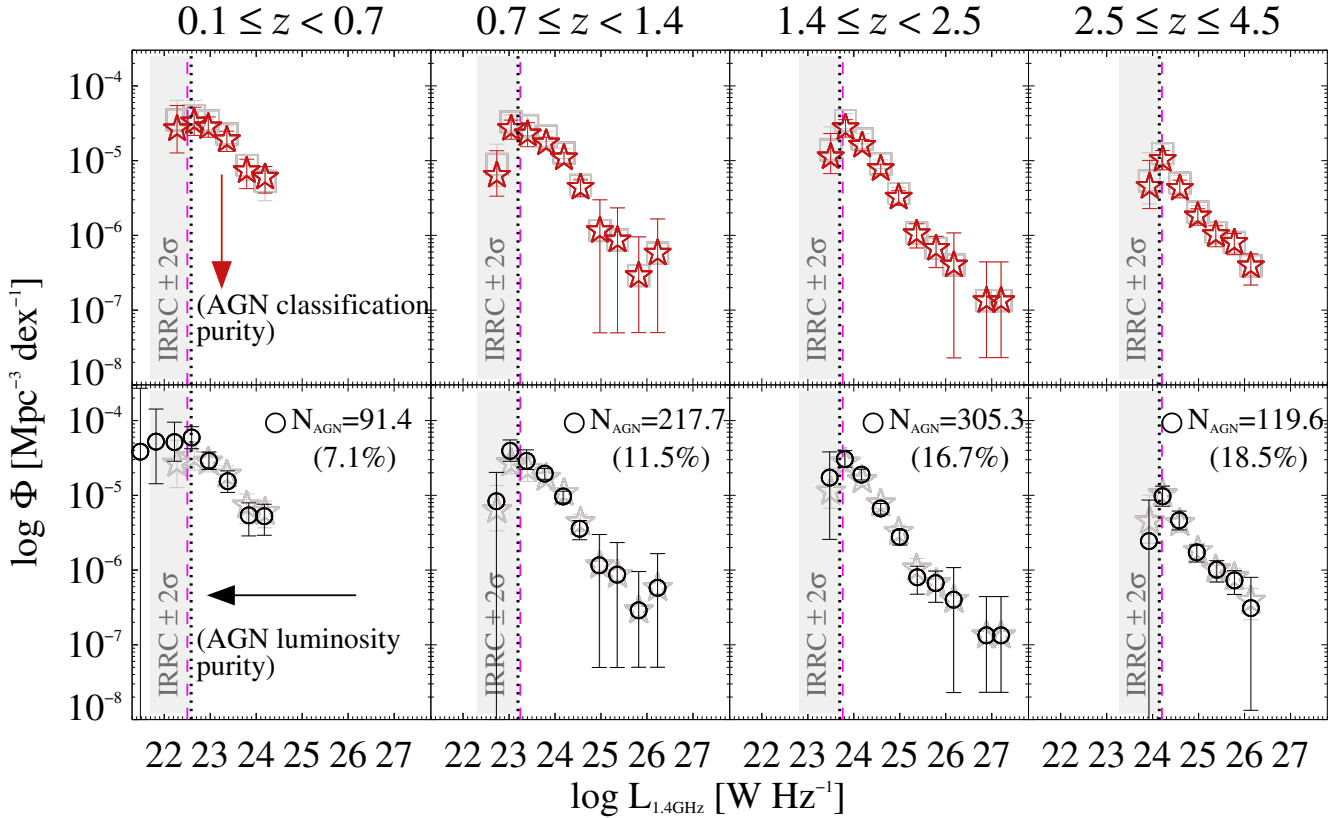


Fig. B.1. 1.4 GHz luminosity function of radio-excess AGN in (*NUVrJ*) SFGs (AGN RLF) at different redshifts (increasing from left to right). The grey shaded area encloses the $\pm 2\sigma$ locus (i.e. ± 0.44 dex) around the luminosity corresponding to the IRRC. All data points below the $+2\sigma$ boundary (dotted black line) are conservatively disregarded to minimise incompleteness. (*Top panels*): Former AGN RLF (grey squares) corrected for AGN classification purity according to f_N^{AGN} (see Sect. 3.2 and Fig. 2). This step lowers $\Phi(L)$ at the faint end (red stars). (*Bottom panels*): f_N^{AGN} -corrected AGN RLF (grey stars) further corrected for AGN luminosity purity at 1.4 GHz, based on $f_{1.4}^{\text{AGN}}$. This step mildly steepens $\Phi(L)$ at the faint end (black circles), while the brightest bins remain unchanged. The dashed magenta line indicates the luminosity corresponding to 90% AGN completeness ($L_{90\%}$) in each z -bin. After applying the above corrections, the final numbers and fractions relative to the parent VLA-COSMOS 3GHz sample are reported in each z -bin. See the text for details.

of each object from the IRRC (i.e. Δq_{IRRC}), but re-scaled to $L_{1.4}$ space based on the median L_{IR} of the sample. Thus, in principle, even a moderately bright radio AGN hosted by a starbursting SFG (i.e. well above average in L_{IR}) would display no radio excess, and hence it might not be counted in our final AGN RLF.

In order to quantify such incompleteness in $L_{1.4}^{\text{AGN}}$, we proceed as follows. We consider the subset of all radio-detected galaxies (black histogram in Fig. 2) with $\Delta q_{\text{IRRC}} > -0.44$ dex (i.e. those that do not show radio excess). Then we look at their cumulative distribution of $L_{1.4}^{\text{AGN}}$ (i.e. $L_{1.4} \times f_{1.4}^{\text{AGN}}$), and we set the luminosity corresponding to the 90th percentile as our completeness limit at 90% level ($L_{90\%}$; see the dashed magenta line in Fig. B.1). We redid this check in each z -bin. This $L_{90\%}$ nearly matches our $+2\sigma$ $L_{1.4}^{\text{AGN}}$ threshold⁵. This is because, when moving closer to $\Delta q_{\text{IRRC}} \sim 0$ (i.e. the IRRC), the distribution is populated by SFGs and progressively radio-fainter AGN, so the brightest ‘missing’ radio AGN will hardly be at above our $+2\sigma$ $L_{1.4}^{\text{AGN}}$ threshold (dotted black line in Fig. B.1). However, this check demonstrates that our approach delivers a $\approx 90\%$ complete sample of radio AGN.

The final numbers and fractions relative to the parent VLA-COSMOS 3GHz sample (5,658 sources) are reported in each z -bin for convenience (non-integer numbers reflect the sum over

⁵ We note that the 100% completeness limit would be placed only 0.1–0.15 dex above $L_{90\%}$.

f_N^{AGN}). The radio-excess AGN fraction is significantly redshift dependent. This is mostly a selection effect induced by the redshift-dependent luminosity cut corresponding to the limiting flux of the survey. As a consequence, at higher redshift only the brightest radio sources (i.e. more likely AGN) are detectable.

However, we acknowledge that our global fraction of radio-excess AGN (734/5,658 $\sim 13\%$) is significantly smaller than that found in Smolčić et al. (2017a) (1,814/7,729 $\sim 23\%$). This difference is partly introduced by our rather conservative corrections for classification and $L_{1.4}$ purity. Accounting for the radio AGN within $\pm 2\sigma$ of the IRRC (i.e. without radio excess) - in turn likely underestimated - would bring the radio AGN fraction to $\sim 19\%$. Another concomitant effect is played by the absence of passive galaxies in our sample. Despite being only $\approx 15\%$ of the parent VLA-3 GHz sample, about a third of those (*NUVrJ*) passive systems would be classified as ‘radio-excess AGN’ according to our criteria, increasing the global AGN fraction by an additional $\sim 5\%$, and thus in line with Smolčić et al. (2017a). The observed prevalence of radio AGN in passive systems (e.g. Gobat et al. 2018; Magdis et al. 2021; Kondapally et al. 2022; Ito et al. 2022) might be interpreted as both a quenching effect due AGN-driven jets hampering SF (e.g. Heckman & Best 2014), but also partly as a selection effect, due to passive galaxies simply showing a higher contrast between AGN versus host radio emission, at fixed $L_{1.4}$. We test the impact of source classification and sample selection (i.e. SF vs. passive) on the AGN RLF in Appendix C. The

take-away message from this test is that source classification methods do not significantly alter the shape and normalisation of the AGN RLF. Instead, the lack of passive galaxies in our sample induces a systematically 2–3× lower AGN RLF at $z \lesssim 1$ (nearly constant with $L_{1.4}$), as seen in Fig. C.1.

Appendix C: Impact of input assumptions on the AGN RLF

We test the impact of our sample selection and radio-excess criterion on the evolving AGN RLF since $z \sim 3$. To this end, we directly compare our observed RLF against that derived by Smolčić et al. (2017c) from the same VLA-COSMOS 3 GHz data. For sake of consistency with their study, here we show our un-corrected data points (i.e. before applying AGN purity and luminosity corrections; see Sect. 3.2).

Figure C.1 displays four realisations of the same AGN RLF, by changing either sample selection or radio-excess criterion, or both. Individual data points (black circles) are shown in the same four redshift intervals. In each panel, the best-fit RLF from Smolčić et al. (2017c, using PLE fitting model) is overlaid for comparison. Figure C.1 is split into four rows, as follows.

(a): Observed AGN RLF re-computed following Smolčić et al. (2017c) in each redshift bin. The sample includes both SF and passive galaxies, in which radio-excess AGN are identified based on the redshift-dependent criterion of Delvecchio et al. (2017). As expected the data points are in very good agreement with the best-fit solution obtained in Smolčić et al. (2017c).

(b): Same as (a), but applying the radio-excess criterion from D21, based on a 2σ offset from the IRRC at the M_* and redshift of each source. The choice of a different radio-excess threshold does not strongly affect the RLF, except at the faint end at $z < 1.4$, where the criterion from D21 rejects some additional AGN.

(c): Same as (a) but changing sample selection. Here only radio detections classified as *NUVrJ*-based SFGs are displayed. As further discussed in Sect. 4.3, the lack of radio AGN in passive galaxies induces a 2–3 times drop at $z < 1$ that is roughly independent of $L_{1.4}$, while at higher redshifts the RLF remains identical.

(d): Dataset from this work, which only includes SF radio detections, with radio-excess AGN identified from D21 (same as top panel of Fig. B.1). As hinted at from (c), removing passive galaxies still results in a drop at $z < 1$, while at higher redshifts the RLF is unchanged. The equivalent AGN RLF obtained by taking a fixed 1.4–3 GHz spectral index ($\gamma = -0.75$) is shown for comparison (empty squares). The global consistency among RLFs indicates that our spectral index assumption does not affect the shape and evolution of the RLF.

In summary, these comparisons suggest that our different radio-excess AGN criterion method does not significantly alter the shape and normalisation of the AGN RLF, at any redshift. Instead, the lack of passive galaxies in our sample induces a systematically 2–3× lower normalisation, especially at $z \lesssim 1$. This offset is roughly independent of $L_{1.4}$; hence, in this work we can still adopt the faint- and bright-end slopes inferred by Mauch & Sadler (2007) on a local sample of radio AGN.

Appendix D: AGN RLF data points

Table D.1 summarises all AGN RLF data points used in this work, split into redshift and M_* .

Table D.1. Luminosity function of radio-excess AGN in SFGs.

z -bin	$L_{1.4}^{\text{AGN}}(z)$ [log(W Hz ⁻¹)]	$\Phi(L, z) \cdot 10^{-6}$ [Mpc ⁻³ dex ⁻¹]	M_* -bin [log(M_*/M_\odot)]	$L_{1.4}^{\text{AGN}}(M_*)$ [log(W Hz ⁻¹)]	$\Phi(L, z, M_*) \cdot 10^{-6}$ [Mpc ⁻³ dex ⁻¹]	
$0.1 \leq z < 0.7$	22.60	$59.38^{+36.13}_{-41.96}$	9.0–10.0	22.59	$21.03^{+6.76}_{-11.43}$	
	22.97	$29.03^{+20.20}_{-22.14}$		22.93	$8.80^{+4.85}_{-4.85}$	
	23.38	$15.46^{+9.70}_{-10.94}$		23.28	$3.56^{+2.52}_{-0.16}$	
	23.84	$5.39^{+2.86}_{-2.86}$	10.0–10.5	23.00	$6.27^{+3.46}_{-4.00}$	
	24.18	$5.26^{+2.92}_{-2.92}$		23.36	$3.96^{+1.18}_{-1.18}$	
			10.5–11.0	22.99	$11.73^{+6.78}_{-6.78}$	
				23.43	$4.57^{+1.21}_{-2.45}$	
				23.87	$2.98^{+2.32}_{-0.01}$	
				24.19	$3.16^{+2.14}_{-0.19}$	
			11.0–12.0	23.40	$3.47^{+1.50}_{-1.50}$	
				23.82	$2.84^{+2.29}_{-0.03}$	
				24.17	$2.10^{+1.78}_{-0.20}$	
	$0.7 \leq z < 1.4$	23.38	$28.81^{+17.09}_{-22.11}$	9.0–10.0	23.02	$19.32^{+12.14}_{-12.76}$
		23.78	$19.50^{+16.11}_{-16.36}$		23.35	$3.76^{+0.44}_{-2.44}$
24.18		$9.65^{+7.94}_{-7.94}$	10.0–10.5	23.85	$2.07^{+1.23}_{-1.23}$	
24.54		$3.57^{+2.55}_{-2.55}$		24.04	$0.31^{+0.40}_{-0.05}$	
24.96		$1.16^{+0.67}_{-0.05}$		24.43	$0.29^{+0.38}_{-0.05}$	
25.35		$0.87^{+0.59}_{-0.05}$		23.36	$9.07^{+3.66}_{-5.47}$	
25.82		$0.29^{+0.38}_{-0.05}$		23.76	$5.89^{+4.27}_{-4.27}$	
26.23		$0.58^{+0.50}_{-0.05}$		24.09	$1.10^{+0.68}_{-0.02}$	
				25.29	$0.29^{+0.38}_{-0.05}$	
			10.5–11.0	23.77	$7.67^{+5.63}_{-5.84}$	
				24.20	$3.91^{+2.83}_{-2.83}$	
				24.48	$1.17^{+0.66}_{-0.06}$	
				25.00	$0.58^{+0.50}_{-0.05}$	
				25.46	$0.29^{+0.38}_{-0.05}$	
			11.0–12.0	25.82	$0.29^{+0.38}_{-0.05}$	
				26.23	$0.58^{+0.50}_{-0.05}$	
				23.78	$3.96^{+2.68}_{-2.68}$	
				24.18	$4.51^{+3.35}_{-3.35}$	
				24.57	$2.19^{+1.39}_{-1.39}$	
			24.93	$0.58^{+0.50}_{-0.05}$		
			25.30	$0.29^{+0.38}_{-0.05}$		

Note. The AGN RLF is split into redshift (Fig. 3) and M_* (Figs. 5, 6) bins. Only luminosity bins above the corresponding $>2\sigma$ cut from the IRRC (i.e. those used in the fitting) are shown. Error bars are given at 1σ level.

Table D.2. Continued.

z -bin	$L_{1.4}^{\text{AGN}}(z)$ [log(W Hz ⁻¹)]	$\Phi(L, z) \cdot 10^{-6}$ [Mpc ⁻³ dex ⁻¹]	\mathcal{M}_* -bin [log($\mathcal{M}_*/\mathcal{M}_\odot$)]	$L_{1.4}^{\text{AGN}}(\mathcal{M}_*)$ [log(W Hz ⁻¹)]	$\Phi(L, z, \mathcal{M}_*) \cdot 10^{-6}$ [Mpc ⁻³ dex ⁻¹]
$1.4 \leq z < 2.5$	23.80	$30.55^{+21.40}_{-23.80}$	9.0–10.0	23.48	$6.38^{+1.32}_{-2.90}$
	24.17	$18.90^{+15.53}_{-16.03}$		23.79	$6.88^{+4.66}_{-4.63}$
	24.58	$6.67^{+5.57}_{-5.58}$		24.11	$1.28^{+0.80}_{-0.79}$
	25.00	$2.78^{+2.16}_{-2.16}$		24.62	$0.28^{+0.24}_{-0.02}$
	25.38	$0.80^{+0.48}_{-0.48}$		25.76	$0.13^{+0.17}_{-0.02}$
	25.78	$0.67^{+0.37}_{-0.37}$	10.0–10.5	23.80	$13.73^{+8.24}_{-8.24}$
	26.17	$0.40^{+0.28}_{-0.02}$		24.15	$3.49^{+2.42}_{-2.72}$
	26.88	$0.13^{+0.17}_{-0.02}$		24.48	$0.75^{+0.53}_{-0.02}$
	27.20	$0.13^{+0.17}_{-0.02}$		24.97	$0.65^{+0.34}_{-0.34}$
				25.50	$0.13^{+0.17}_{-0.02}$
				26.15	$0.13^{+0.17}_{-0.02}$
			10.5–11.0	24.17	$9.72^{+7.79}_{-7.60}$
				24.59	$3.39^{+2.62}_{-2.62}$
				25.00	$0.92^{+0.57}_{-0.57}$
				25.31	$0.26^{+0.23}_{-0.02}$
				25.83	$0.27^{+0.23}_{-0.02}$
			11.0–12.0	24.18	$4.63^{+3.42}_{-3.44}$
				24.59	$2.27^{+1.66}_{-1.66}$
				25.00	$1.25^{+0.84}_{-0.84}$
				25.40	$0.40^{+0.27}_{-0.02}$
			25.74	$0.27^{+0.23}_{-0.02}$	
			26.18	$0.27^{+0.23}_{-0.02}$	
			26.88	$0.13^{+0.17}_{-0.02}$	
			27.20	$0.13^{+0.17}_{-0.02}$	
$2.5 \leq z \leq 4.5$	24.22	$9.71^{+6.18}_{-7.15}$	9.0–10.0	24.20	$1.19^{+0.87}_{-0.04}$
	24.58	$4.63^{+3.10}_{-3.45}$		24.48	$0.36^{+0.31}_{-0.04}$
	24.96	$1.72^{+1.28}_{-1.28}$		24.95	$0.11^{+0.14}_{-0.02}$
	25.40	$1.01^{+0.69}_{-0.69}$		25.27	$0.09^{+0.12}_{-0.02}$
	25.78	$0.73^{+0.47}_{-0.47}$	10.0–10.5	24.57	$0.88^{+0.49}_{-0.49}$
	26.14	$0.31^{+0.18}_{-0.01}$		24.94	$0.36^{+0.20}_{-0.02}$
				25.28	$0.08^{+0.10}_{-0.01}$
				25.79	$0.16^{+0.24}_{-0.07}$
			10.5–11.0	24.56	$2.68^{+1.58}_{-1.88}$
				24.96	$0.94^{+0.59}_{-0.59}$
				25.40	$0.42^{+0.23}_{-0.23}$
				25.79	$0.23^{+0.16}_{-0.01}$
				26.19	$0.08^{+0.10}_{-0.01}$
			11.0–12.0	24.64	$0.71^{+0.41}_{-0.44}$
				24.99	$0.34^{+0.16}_{-0.16}$
			25.45	$0.42^{+0.31}_{-0.02}$	
			25.77	$0.33^{+0.20}_{-0.01}$	
			26.14	$0.23^{+0.16}_{-0.01}$	