



UNIVERSITÀ
DEGLI STUDI
FIRENZE

FLORE

Repository istituzionale dell'Università degli Studi di Firenze

Fundamental differences in the radio properties of red and blue quasars: insight from the LOFAR Two-metre Sky Survey (LoTSS)

Questa è la Versione finale referata (Post print/Accepted manuscript) della seguente pubblicazione:

Original Citation:

Fundamental differences in the radio properties of red and blue quasars: insight from the LOFAR Two-metre Sky Survey (LoTSS) / D J Rosario; V A Fawcett; L Klindt; D M Alexander; L K Morabito; S Fotopoulou; E Lusso; G Calistro Rivera. - In: MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY. - ISSN 0035-8711. - ELETTRONICO. - 494:(2020), pp. 3061-3079. [10.1093/mnras/staa866]

Availability:

This version is available at: 2158/1220382 since: 2021-01-04T10:18:08Z

Published version:

DOI: 10.1093/mnras/staa866

Terms of use:

Open Access

La pubblicazione è resa disponibile sotto le norme e i termini della licenza di deposito, secondo quanto stabilito dalla Policy per l'accesso aperto dell'Università degli Studi di Firenze (<https://www.sba.unifi.it/upload/policy-oa-2016-1.pdf>)

Publisher copyright claim:

(Article begins on next page)

Fundamental differences in the radio properties of red and blue quasars: Insight from the LOFAR Two-metre Sky Survey (LoTSS)

D. J. Rosario¹, V. A. Fawcett¹, L. Klindt¹, D. M. Alexander¹, L. K. Morabito¹, S. Fotopoulou², E. Lusso^{3,4}, and G. Calistro Rivera⁵

¹*Centre for Extragalactic Astronomy, Department of Physics, Durham University, South Road, DH1 3LE, Durham, UK*

²*School of Physics, University of Bristol, UK*

³*Dipartimento di Fisica e Astronomia, Università di Firenze, via G. Sansone 1, 50019 Sesto Fiorentino, Firenze, Italy*

⁴*INAF – Osservatorio Astrofisico di Arcetri, Largo Enrico Fermi 5, 50125 Firenze, Italy*

⁵*European Southern Observatory, Karl-Schwarzschild-Strasse 2, 85748, Garching bei München, Germany*

7 April 2020

ABSTRACT

Red quasi-stellar objects (QSOs) are a subset of the luminous end of the cosmic population of active galactic nuclei (AGN), most of which are reddened by intervening dust along the line-of-sight towards their central engines. In recent work from our team, we developed a systematic technique to select red QSOs from the Sloan Digital Sky Survey (SDSS), and demonstrated that they have distinctive radio properties using the Faint Images of the Radio Sky at Twenty centimeters (FIRST) radio survey. Here we expand our study using low-frequency radio data from the LOFAR Two-metre Sky Survey (LoTSS). With the improvement in depth that LoTSS offers, we confirm key results: compared to a control sample of normal “blue” QSOs matched in redshift and accretion power, red QSOs have a higher radio detection rate and a higher incidence of compact radio morphologies. For the first time, we also demonstrate that these differences arise primarily in sources of intermediate radio-loudness: radio-intermediate red QSOs are $\times 3$ more common than typical QSOs, but the excess diminishes among the most radio-loud and the most radio-quiet systems in our study. We develop Monte-Carlo simulations to explore whether differences in star formation could explain these results, and conclude that, while star formation is an important source of low-frequency emission among radio-quiet QSOs, a population of AGN-driven compact radio sources is the most likely cause for the distinct low-frequency radio properties of red QSOs. Our study substantiates the conclusion that fundamental differences must exist between the red and normal blue QSO populations.

Key words: quasars: general – galaxies: star formation – radio continuum: galaxies – surveys

1 INTRODUCTION

Quasi-stellar objects (QSOs; often also called quasars) are powerful Type 1 Active Galactic Nuclei (AGN) that mark the sites of the most substantial accretion-powered growth of supermassive black holes in the Universe. They can be identified by their extreme luminosities and unique spectral features (Croom et al. 2004; Schneider et al. 2010) even in the early Universe (e.g., Fan et al. 2001).

Most QSOs display blue ultraviolet-optical spectral slopes, consistent with effective temperatures of their inner accretion discs of $\sim 10^5$ K. This allows the efficient identi-

fication of QSOs from multi-band imaging (e.g., Sandage & Wyndham 1965; Richards et al. 2002), making them an optimal target population for large spectroscopic surveys such as the Sloan Digital Sky Survey (SDSS). While most QSOs have blue colours, a minority ($\approx 10\%$; Richards et al. 2003; Klindt et al. 2019) have red optical colours indicative of mild or moderate dust reddening along the line of sight ($E(B-V) \gtrsim 0.04$). These red QSOs likely represent the detectable end of a population of more heavily extinguished luminous AGN, some of which have been uncovered through large-area near-infrared surveys (Glikman et al. 2007; Banerji et al. 2012).

Within the scope of the classical unified model of AGN,

in which the accretion disc and broad-line region are obscured by a canonical obscuring structure, the nuclear dusty ‘torus’ (Netzer 2015, and references therein), red QSOs are simply normal blue QSOs that are viewed at grazing incidence to the torus. However, this view has been challenged by numerous lines of evidence which demonstrate that red QSOs have distinctive properties that cannot be explained by differences in torus orientation. Compared to typical QSOs, their bolometric luminosity functions are flatter (Banerji et al. 2015), their host galaxies are more often in major mergers (Urrutia, Lacy & Becker 2008; Glikman et al. 2015), and they may show a higher incidence of strong AGN outflows (Urrutia et al. 2009).

The classical radio bands offer some unique advantages towards the study of QSOs, particularly in tests of AGN unification. Most radio emission from galaxies arises from optically thin synchrotron-emitting plasma, and is not substantially absorbed by intervening dust or cold gas. In the context of AGN, radio emission is therefore an excellent orientation-insensitive tracer of the nuclear output. Indeed, AGN dominate the bright radio source population that comprise the majority of large-area or all-sky surveys (e.g., Helfand, White & Becker 2015). At fainter fluxes, non-AGN processes, such as star formation (SF), contribute significantly to the radio emission of detectable sources (Panessa et al. 2019, and references therein).

Early work has suggested a link between the occurrence of radio emission and the optical colours of QSOs (Ivezić et al. 2002; Richards et al. 2003; White et al. 2003, 2007; Tsai & Hwang 2017), but until recently, there lacked a detailed and systematic study which controlled for important systematics such as survey selection effects, AGN luminosity, and cosmic evolution of the QSO population. Harnessing the extensive spectroscopic sample and multi-wavelength constraints available over the SDSS, our team has devised a redshift-insensitive approach to select red QSOs which mitigates these biases. Using the SDSS DR7 QSO catalogue and the FIRST 1.4 GHz radio survey, our foundational study identified a number of essential radio properties that are preferentially found among red QSOs (Klindt et al. 2019). These systems show a detection rate in FIRST that is up to three times higher than normal QSOs of equivalent luminosity at the same redshifts. The radio sources responsible for the enhancement are inevitably compact, showing no extended structure beyond the FIRST imaging beam of 5". They are also often faint in the radio; the most substantial enhancement in the radio source population among red QSOs occurs among “radio-quiet” sources (see Klindt et al. (2019) or Section 3.3 for a definition of radio-loudness).

A number of large-area modern radio surveys are currently in operation, giving us the opportunity to validate and expand upon the key work of Klindt et al. (2019). Especially interesting in this regard are the next-generation low-frequency radio surveys from the International Low-Frequency Array (LOFAR), particularly the wide-area component of the LOFAR Two-metre Sky Survey (LoTSS), which will eventually cover a contiguous area of 20,000 deg². For sources with typical synchrotron spectra, LoTSS is an order of magnitude more sensitive than FIRST (Section 3.1). In addition, the spectral baseline available between the observing bands of LoTSS (≈ 144 MHz) and FIRST (≈ 1.4 GHz) yield additional useful diagnostics into the radio prop-




Figure 1. Observed $g - i$ colour plotted against redshift for the ≈ 43000 SDSS DR14 QSOs that lie in the region covered by the LoTSS DR1 field (see Section 2.3). Only QSOs in the redshift range of $0.5 < z < 2.5$ are shown. Red QSOs (rQSOs) and control QSOs (cQSOs) from the ‘parent sample’ (see Section 2.1) are shown as red and cyan points. These constitute the ‘colour-selected’ QSOs used in this work, while the black points indicate other QSOs excluded from our analysis. Note that the range of colours used to define these samples is redshift-dependent as the observed g and i bands sample different parts of a typical QSO spectral energy distribution (SED) with redshift. rQSOs span a much larger range in colour than cQSOs.

erties of QSOs. LoTSS also allows us to ameliorate potential biases related to the fact that a fraction of the SDSS QSOs were themselves selected for spectroscopic follow-up using the FIRST radio survey (Richards et al. 2002).

In this work, we take advantage of the expanded, deeper, and more complete database of spectroscopic QSOs from the SDSS DR14 QSO catalogue (Pâris et al. 2018), combining it with low-frequency radio imaging and photometry from the LoTSS first full-resolution data release (DR1). Section 2 presents the datasets and the methodology of our analysis. Taking an approach that parallels Klindt et al. (2019), we contrast the radio detection rates (Section 3.1), luminosities and spectral indices (Section 3.2), radio-loudness (Section 3.3) and morphologies (Section 3.4) of both red and normal QSOs. In Section 4, we explore the origin of the radio emission in QSOs and conclude this work by folding together our results into a discussion of the likely origin of the unique radio properties of red QSOs.

Throughout this work, we assume a concordance cosmology (Planck Collaboration et al. 2016) and implement it in our calculations using software built into the `AstroPy cosmology` module.

2 DATA AND METHODS

2.1 SDSS colour-selected QSOs

The most recent edition of the SDSS QSO catalogue (DR14; Pâris et al. 2018) contains more than 500K entries from a number of different QSO spectroscopic surveys under the SDSS umbrella, including those from the classical SDSS I/II surveys, BOSS, e-BOSS, and dedicated follow-up programs

such as SPIDERS (Pâris et al. 2018, and references therein). In addition to spectroscopic redshifts, the catalogue includes photometry in the five SDSS bands (*ugriz*) and the amount of foreground Galactic extinction in each band. When SDSS fluxes and colours are discussed in this work, the reader should assume, unless otherwise stated, that they have been corrected for Galactic extinction following the recommendations of Pâris et al. (2018).

We selected red QSOs (hereafter, rQSOs) following the approach described in Klindt et al. (2019) and illustrated in Figure 1. For our selection, we considered all SDSS DR14 QSOs in the redshift range $0.5 < z < 2.5$, motivated by the desire to exclude lower luminosity Seyfert 1 galaxies which populate the SDSS DR14 catalogue at low redshifts, and to avoid the effects of the Lyman break on our colour selection at $z > 2.5$. Working in contiguous, equally-populated redshift bins of 1000 objects each, we selected rQSOs as the 10% of population with the reddest $g - i$ colours. This process identifies a set of rQSOs that is unbiased by the systematic observed-frame colour variation of the SDSS QSO population with redshift. From the same binned samples, we also selected the 50% of QSOs with colours that spanned the median $g - i$ colour (hereafter, cQSOs). These systems represent the typical “blue” QSO population and serve as a set of controls for much of our further analysis.

From the 43124 SDSS DR14 QSOs in the region of sky with LoTSS coverage (see Section 2.3), we arrive at a set of 3568 rQSOs and 17316 cQSOs, which we refer to as the ‘parent sample’ in the rest of this work. Figure 1 summarises the effect of our observed-frame colour-based selection as it pertains to these QSOs.

While broadly based on the approach of Klindt et al. (2019), our selection is different in detail. Klindt et al. (2019) identify QSOs from the SDSS DR7 ‘uniform’ sample, a set of objects subject to magnitude and colour cuts. Their work determined that the restriction to a uniform sample did not strongly influence their results, and therefore, we relax this criterion. In addition, Klindt et al. (2019) identified a control sample as the 10% of QSOs around the median redshift-dependent colour, while we use a substantially more relaxed 50% cQSO cut. This again relies on the assertion from Klindt et al. (2019) that the radio properties were only deviant among the reddest of the SDSS QSO population. Our more liberal cQSO selection approach provides us with a larger control sample allowing for more detailed statistical tests. An accompanying study by Fawcett et al. (2020), which examines the radio properties of SDSS QSOs in the SDSS Stripe82 and the COSMOS fields, also employs a similar colour-based selection strategy as in this paper.

2.2 WISE photometry

Rest-frame mid-infrared (MIR) emission in the vast majority of SDSS QSOs is dominated by AGN-heated dust (e.g. Richards et al. 2006), which can serve as a useful extinction-free measure of the intrinsic power of the AGN¹. This prop-

erty is particularly important when comparing cQSOs and rQSOs: without the application of poorly-constrained extinction corrections the additional dust extinction in red QSOs leads to depressed estimates of their UV/optical-based AGN luminosities (Klindt et al. 2019).

We use photometry from the *WISE* all-sky survey to estimate the rest-frame 6 μm luminosity ($L_{6\mu\text{m}}$) of the QSOs. To this end, we imposed on our sample a $S/N > 2$ detection in the three shorter *WISE* bands (W1 [3.5 μm], W2 [4.6 μm], and W3 [12 μm]). We compiled this photometry from the ALLWISE catalogue available on IRSA, using a cone search with a tolerance of $2''.7$, which ensures a false-match probability of $< 1\%$ (Lake et al. 2012; Klindt et al. 2019).

$L_{6\mu\text{m}}$ was calculated using W2 and W3 fluxes, expressed per logarithmic flux interval (i.e., νf_ν) by multiplying the flux densities by the pivot frequencies of the respective filter profiles. These fluxes were interpolated/extrapolated in log-linear fashion to the redshift-dependent observed-frame wavelength corresponding to rest-frame 6 μm . The resultant flux was converted to a luminosity using the cosmological luminosity distance. While we did not explicitly use the W1 flux in these calculations, we require a detection in this band to ensure clean counterpart associations with the LoTSS radio sources (see Section 2.3). About 56% of QSOs from the parent sample have the requisite photometry in the three *WISE* bands. Henceforth, we refer to the DR14 colour-selected QSOs in the redshift range $0.5 < z < 2.5$ with $L_{6\mu\text{m}}$ estimates (i.e., a $S/N > 2$ detection in W1, W2 & W3) as the ‘working sample’: a total of 2542 rQSOs and 9212 cQSOs.

Our working sample covers a large range in redshift. To minimise the confounding effects of Malmquist bias, we divide the QSOs in the working sample into four redshift bins for most of our analysis: $0.5 \leq z < 0.8$, $0.8 \leq z < 1.5$, $1.5 \leq z < 2.0$, and $2.0 \leq z < 2.5$.

For the rest of this work, we assume that the 6 μm bolometric corrections of cQSOs and rQSOs are indistinguishable, i.e., an identical average fraction of the intrinsic luminosities of their AGNs are reprocessed by their dusty tori. For this to be the case, rQSOs and cQSOs must have the same distribution of torus covering factors. If the structures of the tori were very different between the two colour-selected QSO subsets, one might expect them to display systematic differences in their MIR colours, since the relative proportion of dust of different temperatures is sensitive to torus geometry. However, Klindt et al. (2019) demonstrated that the *WISE* colours of the different QSO subsets are very similar, supporting our working assumption. Future work from our team aims to robustly test this notion using detailed spectroscopic and SED modelling, but for now, we proceed with the assumption that the cQSOs and rQSOs have intrinsically similar tori, and therefore, similar 6 μm bolometric corrections.

Figure 2 compares the $L_{6\mu\text{m}}$ distributions of cQSOs and rQSOs from the working sample in our four main redshift bins. In the three higher redshift bins ($z > 0.8$), the median $L_{6\mu\text{m}}$ of the rQSOs is higher than the cQSOs with high statistical significance: two-sided Kolmogorov-Smirnov (KS) tests indicate that the respective distributions are inconsistent at a probability of $> 99\%$. As discussed in Klindt et al. (2019), this can be understood as a consequence of the optical incompleteness suffered by red QSOs in the SDSS. In the

¹ Based on the Draine (2003) IR interstellar extinction law, the MIR extinction expected from the dust column inferred among rQSOs is < 0.05 mag, essentially negligible for the purposes of this work.




Figure 2. The distributions of rest-frame 6 μm luminosity ($L_{6\mu\text{m}}$) of cQSOs (cyan) and rQSOs (red) from the working sample (defined in Section 2.1) in four redshift bins. All distributions have been normalised to unity for ease of visual comparison. The median $L_{6\mu\text{m}}$ and its error for cQSOs (cyan text) and rQSOs (red text) are written within each panel. The dashed vertical lines give the 6 μm luminosities of powerful starbursts with the indicated SFRs, scaled from the local SF-dominated ULIRG Arp 220. In our lowest redshift bin, a non-negligible fraction of $L_{6\mu\text{m}}$ from the lowest-luminosity QSOs could potentially arise from strong SF rather than AGN-heated dust. At higher redshifts, our sample consists of more luminous AGN that are too powerful to be contaminated by SF in the MIR.

rest-frame optical/UV, rQSOs are fainter than equally powerful cQSOs due to their higher dust extinction. Therefore, the rQSOs that are optically bright enough for spectroscopic targeting and characterisation by the SDSS will generally be more luminous than cQSOs when using an extinction-insensitive measure such as $L_{6\mu\text{m}}$.

At $0.5 \leq z < 0.8$, rQSOs and cQSOs are equally luminous. Among the lower luminosity QSOs at these redshifts, emission from dust heated by strong starbursts can significantly contaminate $L_{6\mu\text{m}}$. The vertical dashed lines in Figure 2 show a characteristic MIR luminosity from extreme starbursts with different star formation rates (SFRs). If a significant fraction of rQSOs at low redshifts are reddened not by dust, but through the contamination of their optical and MIR light from a powerful on-going starburst, this may explain why their MIR luminosities are similar to normal bright QSOs with blue colours. This is consistent with the conclusions of Klindt et al. (2019) that, at lower redshifts, there exists a subpopulation in which a lower luminosity Type 1 AGN in a massive host galaxy satisfies our simple colour-based selection of rQSOs. Towards higher redshifts, the nuclear luminosities involved are too high to suffer significantly from this effect, leading to a much purer dust-reddened rQSO sample.

2.3 LoTSS DR1

The LOFAR Two-metre Sky Survey is an on-going programme to image the entire northern sky in the 120–168 MHz radio band at a resolution of 6'' (Shimwell et al. 2017). The first full resolution data release (DR1; Shimwell et al. 2019) covers 424 deg² in the HETDEX spring field (RA: 10h45m – 15h30m, Dec: 45° – 57°), which lies completely within the SDSS footprint. Calibrated and synthesised radio images, in the form of science mosaics and root-mean-square (RMS) noise maps, and a combined source catalogue are available from the LOFAR surveys public releases page.

We selected SDSS DR14 QSOs in the LoTSS DR1 area by identifying those that have a measured value in the RMS maps at their sky positions. This approach is superior to a simple footprint-based selection, as it allows us to account

for the patchy, non-uniform sky coverage of the radio survey. While the LoTSS radio imaging has a median sensitivity of 71 $\mu\text{Jy beam}^{-1}$, the noise varies considerably over the field, and is particularly high around bright sources due to side lobe residuals that remain after deconvolution of the maps (Shimwell et al. 2019). We perform the majority of the analysis in this paper – radio detection statistics and morphological analysis – on all QSOs in the LoTSS DR1 area, irrespective of the local noise level. However, we have checked that our primary results hold after restricting the sample to those in areas of low noise. While this restriction is better for uniform statistics, it biases our study against those QSOs with the brightest radio counterparts since these sources alter the local RMS in their regions of the maps.

The LoTSS images do not have a uniform point-spread function (PSF) due to ionospheric variations, astrometric errors, the deconvolution masking strategy, and time- and frequency-dependent station beams and smearing. As a consequence, it is not trivial to determine whether a source is unresolved or not, and occasionally absolute astrometric errors can be $> 1''$.

To help ameliorate positional uncertainties and optimise the use of the radio survey for multi-wavelength work, the LoTSS DR1 release includes a catalogue of associations to optical and infrared (IR) counterparts (Williams et al. 2019). This catalogue is the product of likelihood-ratio matching using colour and magnitude information, as well as the visual examination of 13,000 extended or complex sources, to identify the best *WISE* W1 and/or PanSTARRS counterparts to each radio source. Multiple components for some large radio sources were reduced to single entries, integrating the fluxes of the components and tying the central or core component with the best optical/IR counterpart. While the catalogue tabulates various multi-wavelength and derived properties, such as photometric redshifts, for the purposes of this work we only use positions of optical/IR counterparts and integrated LoTSS radio flux measurements.

We linked the SDSS QSOs to LoTSS sources using the counterpart positions from this catalogue. All the QSOs in our working sample have *WISE* W1 detections (by construction; see Section 2.2) and are bright enough to be detectable

in PanSTARRS. Therefore, we expect negligible incompleteness to be introduced into our sample by using the associations from the LoTSS DR1 catalogue. Indeed, we found that essentially all cross-matches between the positions of the radio source associations and those of our QSOs had offsets $< 0''.75$. Taking this as the cross-matching threshold, we linked 1868 QSOs to LoTSS radio sources with an integrated $S/N > 5$ ($\approx 16\%$ of the working QSO sample). In Section 3.1, we present a more detailed analysis of the radio detection statistics of the QSOs.

2.4 FIRST

We also searched for associations between QSOs from our working sample and radio sources detected by the Faint Images of the Radio Sky at Twenty centimeters (FIRST) survey, a legacy VLA programme that imaged the entire SDSS footprint at 1.4 GHz with a resolution of $5''$ and a 5σ detection limit of 1 mJy (Becker, White & Helfand 1995; Helfand, White & Becker 2015). Following the approach taken by Klindt et al. (2019), we assumed that all detected radio sources with centroids within $10''$ of a QSO are associated with it. In the cases of multiple radio sources within the $10''$ search radius, we add the FIRST measurements of the integrated fluxes of all components to yield a total flux for such multi-component sources. This simple criterion has a low false-match probability ($\approx 0.2\%$) but identifies the vast majority of extended and multi-component radio sources found among SDSS QSOs, missing only $\approx 2\%$ of associations from large double-lobed radio sources (Lu et al. 2007; Klindt et al. 2019).

The LoTSS data offer a means to reliably recover the fluxes of very extended radio sources, owing to the high sensitivity of the LOFAR images to low surface brightness steep-spectrum emission. The LoTSS DR1 associations catalogue identifies radio sources that are composed of multiple components, and provides positions for each sub-component (Williams et al. 2019). Based on our visual assessment of the LoTSS images (Section 3.4), we determined that we could reliably associate FIRST counterparts to these large extended sub-components using a secondary search tolerance of $3''$. In this way, we obtained more accurate total fluxes for 68 FIRST-detected radio QSOs, including emission well beyond $10''$ from their central positions.

Of the QSOs in the working sample, only 544 ($\approx 4.6\%$) have a FIRST association. The difference in overall detection rate with respect to LoTSS is a testament to the significantly better sensitivity of the LOFAR surveys to radio emission from QSOs.

A small number of QSOs with FIRST detections do not have entries in the LoTSS DR1 association catalogue because they have been excluded by the likelihood ratio matching technique due to complex structure or true mis-association. We visually examined the 57 colour-selected QSOs from our working sample that fell into this category, and identified reliable associations for 37. Only 20 QSOs are definite FIRST sources without a reliable LoTSS detection. These QSOs are candidates for inverted spectrum radio sources (see Figure 5), a rare and interesting population.

2.5 Matching subsamples of colour-selected QSOs

An important consequence of the differences in intrinsic luminosities between rQSOs and cQSOs (Section 2.2) is the effects of *WISE* selection on the redshift distributions of the colour-selected QSO subsets. By construction, cQSOs and rQSOs in the parent sample have essentially identical redshift distributions: due to the non-uniformity of SDSS QSO selection across the sky, there are subtle differences introduced by the restriction of the parent sample to the LoTSS survey area, but these are quite minor. Despite this, because the population of rQSOs are intrinsically brighter in the MIR, they become over-represented at higher redshifts after applying our *WISE* cuts. In the working sample, this effect leads to a long high redshift tail among rQSOs compared to cQSOs.

In order to minimise any biases in our results due to differences in redshifts and intrinsic luminosities of the two colour-selected QSO subsets, we employ a procedure of matching the rQSOs to cQSOs of similar redshift and $L_{6\mu\text{m}}$. For each rQSO in the working sample, we randomly choose a unique cQSO within a redshift tolerance of $0.02 \times (1+z)$ and a $L_{6\mu\text{m}}$ tolerance of 0.1 dex. We exclude from the analysis any rQSOs that do not have a matching cQSO from this process; in practice, however, $> 99\%$ of rQSOs are paired with cQSOs in any given iteration of the matching procedure. Henceforth, we refer to these matched subsamples as “L6z-matched” QSOs.

The matching procedure is a Monte-Carlo process. Due to its inherent randomness, all of the analysis in this work that relies on the L6z-matched QSO subsamples is subject to variance. We minimise the effects of variance in L6z-matching by performing a minimum of 30 independent iterations of each analysis. Results are quoted and figures are plotted using the median statistics of these iterations. As an example, consider the lower panel of Figure 4. The blue points show radio detection statistics of L6z-matched cQSOs in bins of redshift. The plotted blue stars are the median detection fractions from 30 iterations of the matching procedure, and the errorbars on the plotted points show the median upper and lower errors using binomial statistical calculations applied at each iteration, following the technique of Cameron (2011). A similar methodology is used to portray summary statistics throughout the paper, unless specifically noted otherwise.

In Figure 3, we illustrate how the steps in selection and matching alter the nature of the control QSOs used in this study. The hatched black histogram shows the distribution of I-band absolute magnitudes, calculated by Pâris et al. (2018), of all cQSOs selected from the SDSS DR14 QSO catalogue that lie in the region covered by LoTSS DR1 and in the redshift interval of $0.5 < z < 2.5$. With the requirement of a $S/N > 2$ detection in the three shorter *WISE* bands (W1/2/3), the number of cQSOs drops by about 50%, the remainder having a median luminosity that is brighter by half a magnitude (cyan histogram). These constitute the cQSOs in the working sample. Finally, after matching to the $L_{6\mu\text{m}}$ and redshifts of the rQSOs, we are left with an even smaller subset of cQSOs that are more luminous again by 0.3 mag (blue open histogram).

Figure 3 demonstrates that as we settle towards a control set of normal QSOs that fully captures the joint dis-




Figure 3. Distributions of rest-frame I-band absolute magnitude (M_I) of cQSOs in the redshift interval $0.5 < z < 2.5$. The hatched histogram shows all cQSOs from the SDSS DR14 parent sample, the filled cyan histogram shows those cQSOs with the additional requirement of a $\text{SNR} > 2$ measurement in the *WISE* 1, 2, and 3 bands, and the open blue histogram indicates the cQSOs left after a further restriction imposed by matching to rQSOs in redshift and $L_{6\mu\text{m}}$. The black, cyan, and blue filled circular points at the top of the graph show the median M_I of the distributions in the respective order noted above. cQSOs are typically less intrinsically luminous than rQSOs, and the L6z-matching process accounts for this. We only show the results of a single iteration of L6z-matching in this Figure, but the blue histogram is essentially unchanged by matching variance.

tribution in redshift and AGN luminosity of red QSOs, we select progressively more luminous systems. This highlights the importance of our fully matched approach: we can be confident that any differences we see in the radio properties of rQSOs with respect to cQSOs are not due to simple selection effects arising from luminosity or redshift mismatches.

3 RESULTS

3.1 LoTSS detection statistics

The simplest comparison we can make between the 144 MHz radio properties of rQSOs and cQSOs is in terms of their detection rates in LoTSS DR1. We consider a source to be detected if the integrated S/N of the radio source (including all elements of a multi-component source) is greater than 5.

In Table 1, we show the fraction of rQSOs and cQSOs in each of our four primary redshift bins that are detected in LoTSS DR1. We report separate detection fractions for the working sample and the L6z-matched subsample. Errors in the detection fractions are calculated using the methodology of Cameron (2011). These detection fractions are also plotted in the lower panel of Figure 4, where we only show the points for L6z-matched rQSOs since they are practically the same objects as the rQSOs from the working sample.

We find that rQSOs are detected at a substantially higher rate in LoTSS compared to the cQSOs. The detection fraction of rQSOs remains constant at $\approx 25\%$ across

all redshifts considered in this study, and is always higher than the $\approx 13\%$ detection fraction of the cQSOs. The effect of matching in $L_{6\mu\text{m}}$ and redshift only marginally increases the cQSO detection fraction, confirming the conclusion from Klindt et al. (2019) that differences in intrinsic accretion rate are not the primary determinant of this result.

Since the FIRST 1.4 GHz survey covers the LoTSS DR1 area, we can also compare the FIRST detection statistics of the colour-selected QSOs between the two radio datasets. Table 1 also lists the fractions of L6z-matched colour-selected QSOs that have $\text{S/N} \geq 5$ detections in the FIRST survey. At all redshifts, these fractions are consistently factors of ~ 3 lower than the detection fractions in LoTSS DR1 for the same sample of QSOs.

LoTSS is considerably more sensitive than FIRST to the emission from radio QSOs across a wide range of spectral indices. Only FIRST sources with strongly inverted metre-wave spectra, or those that display strong variability on decade time-scales, will have no LoTSS DR1 counterpart in the overlapping area of the two surveys. After visually reconciling mis-associated sources (see Section 2.4), and excluding FIRST sources at the edges of the LoTSS maps, we find that $\approx 97\%$ of FIRST-detected QSOs are also detected in LoTSS. In Section 3.2, we take advantage of the mutual completeness of LoTSS and FIRST to examine and compare the radio spectral indices of colour-selected QSOs.

LoTSS offers us an important opportunity to evaluate the radio selection biases of SDSS rQSOs. Since FIRST predated the bulk of the SDSS Legacy spectroscopic survey, radio detections were used as a way to distinguish between stars and QSOs in the early targeting strategy (Richards et al. 2002). This pre-selection of radio-bright QSOs, particularly in parts of the QSO multi-colour selection space with low completeness (overlapping the colours of stars or at high redshifts), will preferentially boost the incidence of radio QSOs with abnormal colours, such as rQSOs. Klindt et al. (2019) minimise such biases by only studying SDSS DR7 QSOs that satisfy the principal multi-colour selection of the Legacy SDSS survey, i.e., those flagged as UNIFORM in Schneider et al. (2010). The SDSS DR14 QSO catalogue is much larger and much more heterogeneous than the DR7 edition. rQSOs selected from this catalogue are therefore more susceptible to FIRST pre-selection biases.

However, the LoTSS survey pushes into a regime of radio luminosity that was not covered by FIRST (see Section 3.2). As a much more recent survey, the sources in LoTSS are not entangled with the selection function of SDSS DR14 QSOs. Even if we exclude all FIRST-matched QSOs from the L6z-matched sample, effectively removing any residual bias introduced by FIRST pre-selection, we still find that rQSOs have $\times 1.6$ higher LoTSS detection rates compared to cQSOs. The higher proportion of rQSOs in LoTSS are an important and independent validation of the key result from Klindt et al. (2019): the population of red QSOs do produce significantly larger amounts of detectable radio emission than normal QSOs.

3.2 Radio luminosities

A comparison of the distributions of 144 MHz radio luminosities (L_{144}) between the two colour-selected QSO subsets can inform us on the possible causes for the higher incidence

Table 1. Radio detection statistics for SDSS DR14 colour-selected QSOs

Subset	$0.5 \leq z < 0.8$	$0.8 \leq z < 1.5$	$1.5 \leq z < 2.0$	$2.0 \leq z < 2.5$
Working sample ^a : LoTSS DR1 detection statistics and detection fractions				
cQSOs	158 / 1046 (14% – 16%)	505 / 3906 (12% – 13%)	334 / 2599 (12% – 13%)	195 / 1657 (11% – 12%)
rQSOs	67 / 285 (21% – 26%)	273 / 968 (26% – 29%)	177 / 687 (24% – 27%)	158 / 598 (24% – 28%)
L6z-matched sample ^b (30 iterations): LoTSS DR1 detection fraction				
cQSOs	12% – 16%	14% – 16%	13% – 16%	12% – 15%
rQSOs	20% – 25%	26% – 29%	23% – 27%	24% – 27%
L6z-matched sample ^b (30 iterations): FIRST detection fractions				
cQSOs	2% – 4%	3% – 5%	2% – 4%	2% – 3%
rQSOs	4% – 7%	7% – 9%	6% – 8%	7% – 9%

Detection fractions, shown as percentages, span the 16th to 84th percentile confidence intervals.

^aLeft / Right numbers are radio-detected / total

^bAbsolute numbers are subject to L6z-matching variance. Therefore, only detection fractions are shown.




Figure 4. Upper panel: The meter-wave radio luminosity (L_{144}) against redshift for QSOs from the working sample. Red/cyan coloured points represent rQSOs/cQSOs. Filled points are those sources detected with $S/N \geq 5$ in both LoTSS DR1 and FIRST surveys. The grey band shows the range in L_{144} from star-forming galaxies with a SFR of $100 M_{\odot} \text{ yr}^{-1}$ based on the FIR-radio correlation and its evolution from Calistro Rivera et al. (2017). Lower panel: The LoTSS DR1 detection fractions for rQSOs (red points), cQSOs from the working sample (cyan points), and cQSOs from the L6z-matched sample (blue points). The fractions are evaluated for QSOs in four redshift bins, delineated by vertical dotted lines in both panels.




Figure 5. Radio spectral indices between 144 MHz and 1.4 GHz (α_R , defined such that the monochromatic radio luminosity $L_\nu \propto \nu^{\alpha_R}$) for all FIRST-detected QSOs from our working sample. Red/cyan colours represent rQSOs/cQSOs. Left panel: α_R vs. redshift. The dashed line is the spectral index of a source that lies at the nominal 5σ detection limits of both FIRST and LoTSS DR1. Only FIRST-detected sources with very flat or inverted radio spectra can be too faint to be detected by LoTSS. Right panel: Distributions of α_R for QSOs detected in both surveys.

of radio emission among rQSOs. Because of the large range in redshifts covered by our working sample, the calculation of monochromatic radio luminosities requires the application of a k-correction to the observed-frame LoTSS fluxes to bring them to a fixed rest-frame estimate. We take the standard approach of assuming a single spectral slope for the meter-wave radio continuum emission parameterised by the spectral index α_R , such that the monochromatic radio luminosity $L_\nu \propto \nu^{\alpha_R}$.

For the QSOs with detections in both LoTSS DR1 and FIRST, we calculate a representative α_R from the ratio of the integrated fluxes at 144 MHz and 1.4 GHz. In Figure 5, we plot this measured spectral index against redshift, differentiating between rQSOs and cQSOs. The spectral indices show a broad distribution with a peak at $\alpha_R \approx -0.7$, a typical value for the cosmic synchrotron-dominated radio source population.

The vast majority of LoTSS detected QSOs are too faint to be detected in FIRST. For these objects, we estimate L_{144} assuming a fixed $\alpha_R = -0.7$. We have verified that none of the results presented in this paper depend on the exact value for the spectral index assumed for the FIRST-undetected radio sources.

In the upper panel of Figure 4, we plot L_{144} against redshift for all LoTSS detected QSOs from the working sample, split by colour-selection. Filled/open symbols indicate which sources are detected/undetected in the FIRST survey. Interestingly, there is a substantial increase in the number counts of radio sources just below the FIRST limit, indicating that LoTSS is probing a different regime of the radio QSO population that was not well-characterised by FIRST (e.g., Gürkan et al. 2019). A similar result has been noted at higher frequencies from deep targeted radio observations of QSOs (e.g., Kimball et al. 2011), and can be understood




Figure 6. Metre-wave radio luminosity (L_{144}) vs. $6\mu\text{m}$ MIR luminosity ($L_{6\mu\text{m}}$) for cQSOs (left) and rQSOs (right). The colours of the plotted points indicate their redshifts (see colourbar). FIRST-detected sources are shown with filled symbols, while open symbols show sources that are only detected in LoTSS. The black solid line marks a constant ‘radio-loudness’ of $\mathcal{R} = -4.5$, the approximate threshold between radio-loud and radio-quiet systems. Dotted black lines show different values of \mathcal{R} that are 0.5 dex apart, and the system of black lines delineate bins of \mathcal{R} featured in Figure 7.

as an increasing proportion of star-formation powered systems among low luminosity radio QSOs. We will explore the nature of this population in more detail in Section 4.1.

3.3 Radio loudness

Additional insight into the nature of this population comes from the relationship between MIR and radio power, which we examine in Figure 6 by plotting L_{144} against $L_{6\mu\text{m}}$ for LoTSS-detected QSOs at $0.8 < z < 2.5$. We exclude the lowest redshifts, those from the working sample at $0.5 < z < 0.8$, to minimise any contamination from strong starbursts (see Section 2.2 and Figure 2). We treat cQSOs and rQSOs separately, and distinguish between those with FIRST detections (filled symbols) and those without (open symbols).

Both cQSOs and rQSOs span a wide range in $L_{6\mu\text{m}}$ and L_{144} . The FIRST-detected systems, being brighter in the radio, occupy the upper scatter of radio luminosities (filled symbols), while those detected solely in LoTSS are inherently less radio-luminous. Among these weaker sources

(open symbols), a correlation appears to exist between $L_{6\mu\text{m}}$ and L_{144} . However, based on an examination of their redshifts (the colour scheme used to plot the points in Figure 6), we conclude that the correlation is driven primarily by the strong redshift-dependent luminosity bias of our sample. At any given redshift, the trend between the MIR and radio luminosities is not obvious, suggesting at best a weak intrinsic correlation.

The diagonal lines in Figure 6 delineate bins of constant L_{144} - $L_{6\mu\text{m}}$ ratio. As we can be assured that $L_{6\mu\text{m}}$ for SDSS QSOs at these redshifts is dominated by reprocessed AGN radiation (Section 2.2), this ratio is a measure of the relative power of synchrotron emission to accretion disc emission, often referred to as ‘radio-loudness’ (\mathcal{R}) in the common literature. For our purposes, we define \mathcal{R} in dimensionless fashion as follows:

$$\mathcal{R} = \log \frac{1.44 \times 10^{13} L_{144}}{L_{6\mu\text{m}}} \quad (1)$$

with L_{144} in W Hz^{-1} and $L_{6\mu\text{m}}$ in erg s^{-1} .

Radio-loudness has a number of different definitions in the literature. Traditionally, it is expressed as a ratio of rest-frame luminosity densities (e.g. Kellermann et al. 1989) or even observed-frame flux densities (e.g. Ivezić et al. 2002). For such definitions, the typical divide between ‘radio-loud’ and ‘radio-quiet’ systems has a value of ≈ 10 . For our purposes, we instead use a ratio of rest-frame monochromatic powers (or luminosities per logarithmic frequency interval), which is a better representation of the relative amount of energy released in the synchrotron component by the AGN (see Section 4.2).

In Klindt et al. (2019), a boundary value was derived between radio-loud and radio-quiet systems for the ratio of 1.4 GHz and *WISE*-based MIR luminosities, calibrated to be equivalent to the more traditional divide laid out by early SDSS studies (Ivezić et al. 2002). As our study deals with 144 MHz luminosities, we translate the calibration from Klindt et al. (2019) by assuming a radio spectral index of -0.7, which results in a boundary at a value of $\mathcal{R} = -4.5$. A source with a canonical spectral index that lies on this boundary in our study will also lie on the equivalent boundary at 1.4 GHz, but would have radio-loudness parameters of -4.2 in the works of Klindt et al. (2019) and Fawcett et al. (2020).

In Figure 6, we plot lines of constant \mathcal{R} as diagonal lines, incremented in steps of 0.5 dex. The solid black line in Figure 6 marks $\mathcal{R} = -4.5$ the separation between radio-loud and radio-quiet systems discussed above. Taking this fiducial separation, we can compare the relative number of radio-loud and radio-quiet QSOs in both colour-selected subsets. We find that 24% of LoTSS-detected rQSOs are radio-loud, as opposed to $\approx 18\%$ of L6z-matched cQSOs, a small but significant difference.

Following Klindt et al. (2019), we take this analysis further by examining whether the relative enhancement of the radio detection rates of rQSOs over cQSOs depends on \mathcal{R} . In Figure 7, we plot the ratio of LoTSS-detected rQSOs and L6z-matched cQSOs in the same bins of \mathcal{R} as set forth in Figure 6. At any value of \mathcal{R} , a value of this ratio > 1 indicates that rQSOs with that degree of ‘radio-loudness’ are preferentially detected in LoTSS. We find a definite peak in the Figure at $\mathcal{R} \approx -4.5$ spanning the divide between radio-quiet and radio-loud sources. At both large and small values

of \mathcal{R} , rQSOs are as likely to be detected in LoTSS as cQSOs, but they are ≈ 3 times more numerous than cQSOs if they are of intermediate radio-loudness.

The grey filled polygon in Figure 7 shows the detection rate enhancement among L6z-matched rQSOs and cQSOs from the FIRST-based of Klindt et al. (2019). Their data have been shifted along the \mathcal{R} axis by 0.3 dex to account for the differences in the radio bands between FIRST and LoTSS, allowing a comparison to our own data points. The rise in the detection rate enhancement with decreasing \mathcal{R} is completely consistent between our two studies, but we find a sharper and more significant dropoff among radio-quiet systems. Given the differences in the QSO parent samples, the redshift range under consideration, the statistical methodology, and the fact that FIRST is only barely sensitive enough to detect radio-quiet AGN (Figure 6), we consider the comparison between our two studies to be quite favourable.

Due to the much-improved sensitivity afforded by LoTSS over FIRST, our analysis is able to track the transition clearly for the first time and identify that the enhanced radio emission in rQSOs is preferentially found among radio-intermediate QSOs. In the parallel work of Fawcett et al. (2020), a similar result is evident from a joint analysis of the Stripe82 and COSMOS radio surveys.

3.4 Radio source morphologies

The observed morphologies of radio QSOs are determined by the power and accretion mechanisms of the central engines, their orientation to the line-of-sight, as well as their environments (Miley 1980; Best & Heckman 2012). Through an examination of the range of 144 MHz morphologies exhibited by the colour-selected QSOs, we expect to gain some insight into the nature of the excess radio emission found in rQSOs.

QSOs with an integrated LoTSS S/N = 5–15, as determined from integrated flux measurements in the LoTSS DR1 associations catalogue (Williams et al. 2019), are often too faint for a reliable visual classification. Some lie in regions of the LoTSS field which have not been deconvolved sufficiently for a good image. In our morphological analysis, we place these sources into a ‘Faint’ category determined purely by their overall S/N.

We refer to the remaining LoTSS radio sources from the working sample with S/N > 15 as the ‘Bright’ subset, and these were the targets of our assessment of visual morphology. By their selection, they comprise a more luminous subset of the colour-selected radio QSOs. However, we have verified that the LoTSS detection fraction of rQSOs in the Bright subset (11%) remains high compared to cQSOs in this subset ($\approx 5\%$), and we can therefore consider them representative of the full working subsample.

For the purposes of visual morphology classification, we created 1.5 arcminute square cutouts centred on the SDSS positions of the 652 QSOs from the Bright subset, and visualised them using a large dynamic range stretch. We assessed the cutout images guided by reference circles with diameters of $6''$ (the generalised PSF width of LoTSS) and $20''$ (within which sources are considered single associations in FIRST). Figure 9 features the stretch, colour table, and overall layout of the images used for classification.

The LoTSS sensitivity varies across the field and around




Figure 7. The ratio of the detection rate in LoTSS for rQSOs vs. L6z-matched cQSOs plotted in bins of the ‘radio-loudness’ parameter \mathcal{R} , defined in Equation 1 as the dimensionless ratio of the MIR and meter-wave radio luminosities. The coloured numbers next to the points indicate the number of LoTSS-detected rQSOs (red) and cQSOs (blue) that contribute to the ratio, a median from 30 bootstrap draws of L6z-matching. The vertical dashed line at $\mathcal{R} = -4.5$ marks the threshold between systems that are defined as ‘radio-loud’ (to the right) and ‘radio-quiet’ (to the left). Note the broad peak around this divide at which rQSOs are more common than cQSOs among LoTSS radio sources. The grey filled polygon shows the detection rate ratio for rQSOs vs. L6z-matched cQSOs from the FIRST-based study of Klindt et al. (2019), appropriately shifted by $\Delta\mathcal{R} = -0.3$ dex to account for the differences between the 1.4 GHz and 144 MHz bands.

bright sources, which could translate into correlations in the appearance of the noise level in the images if they were examined in the order of a catalogue sorted by sky coordinates. To avoid this, we randomised the order of the stamps before visual classification.

We classified the radio source in each stamp into four categories:

- **Compact:** The source size is visually consistent with the LoTSS average PSF and there are no other nearby clearly associated lobes.

- **Marginally Extended:** The source is mildly extended beyond the LoTSS beam, with all of its flux contained within a $10''$ radius of the QSO’s position. The source’s extended appearance or asymmetry may be due to factors such as local noise, the non-uniformity of the local PSF in different parts of the LoTSS DR1 area, or due to genuine source extension at the scale of $6''$. We treat this category as a tran-

sition between compact and well-extended morphologies; it likely contains a mixture of radio sources of both types.

- **Extended:** The source is clearly extended at the resolution of the LoTSS images, with emission apparent beyond $10''$ of the QSO’s position. However, it does not demonstrate the classical bi-polar geometry typical of radio galaxies. This category includes sources with a diverse range of sizes and structures, spanning the divide between compact radio morphologies and some of the largest radio sources in LoTSS. It may also include highly asymmetrical radio galaxies in which only one of the jets or lobes is bright enough to be visible in the LoTSS images.

- **Extended Radio Galaxy:** The source clearly shows a bi-polar jet/lobe geometry with the QSO at the radio core. These are often the largest and brightest radio QSOs.

Figure 9 show typical examples of the morphology cat-




Figure 8. Radio morphology distributions of L6z-matched rQSOs (red) and cQSOs (blue). The Y-axis is the fraction of QSOs of a certain morphology (marked on the X-axis) out of the total number of QSOs of the colour-selected subset in the redshift range $0.8 < z < 2.5$. 30 bootstrap draws of L6z-matching were used to generate the data for this Figure. The coloured numbers next to the points indicate the median number of radio QSOs that contribute to the fraction, while the coloured numbers on the top right of each panel give the median of the total number of QSOs analysed from all the draws. See Section 3.4 for a description of the morphological classes. The left panel considers the distributions for all LoTSS detections, while the right considers only the subset with a meter-wave radio luminosity $L_{144} > 10^{26}$ W Hz⁻¹ (see Section 4.3 for the associated discussion).




Figure 9. Typical examples of the radio morphology classifications that are used in the morphological analysis of Section 3.4. The QSOs featured here all lie at $1.0 < z < 1.5$. Each image is 1.5 arcminutes on a side and oriented North up and East to the left. The visual category is written in large black letters at the top of each panel, and the SDSS ID of the respective QSO, which incorporates its celestial position, is written in white letters at the bottom left. The concentric white circles provide reference diameters of 6'' and 10'' centred on the QSO position. The ‘Faint’ category is not a morphological category itself, but includes all LoTSS sources that have a S/N < 15 from the survey catalogue.

egories, including an example of ‘Faint’ class for which no explicit classification was made.

In the left panel of Figure 8, we plot the fraction of colour-selected QSOs in each visual category relative to all QSOs in the respective colour-selected subset. Only L6z-matched QSOs in the redshift range of $0.8 < z < 2.5$ are considered here so we can be confident that AGN luminosity-driven selection effects are minimal.

We find a significant difference in the incidence of faint, compact and marginally extended radio sources among rQSOs compared to equally luminous cQSOs. About 1 – 2% of cQSOs are found in each of the four visually-assessed categories. In contrast, $\approx 6\%$ of rQSOs show compact radio morphologies, a factor of 3 times higher than cQSOs. In essence, almost all of the difference in the LoTSS detection

rates of rQSOs arises through an excess of compact or faint radio sources.

This result independently confirms the findings of Klindt et al. (2019), in which a similar morphological analysis was performed on FIRST images of SDSS DR7 QSOs. With the deeper images and better sensitivity to extended emission offered by LoTSS DR1, and the inclusion of fainter QSOs from the SDSS DR14 catalogue, we probe a different regime of accretion and radio power among QSOs. The similarity of the results from both works is decisive validation of the special radio properties of red QSOs.




Figure 10. Comparisons of the normalised distributions of metre-wave radio luminosity (L_{144}) of L6z-matched rQSOs (red open histograms) and cQSOs (blue filled histograms), separated into the five morphological categories (see Section 3.4). Only LoTSS-detected QSOs at $0.8 < z < 2.5$ are considered. 30 bootstrap draws of L6z-matching were used to generate the data for this Figure. The numbers on the top left of each panel are the median number of colour-selected QSOs that comprise the histograms. The median K-S P value is written in the top-left corner of each panel: this is the probability that the distributions are drawn from the same parent population, a measure of their similarity.




Figure 11. Comparisons of the normalised distributions of radio spectral index (α_R) of L6z-matched rQSOs (red open histograms) and cQSOs (blue filled histograms), separated into the five morphological categories (see Section 3.4). Only LoTSS-detected QSOs at $0.8 < z < 2.5$ that are additionally detected in the FIRST survey are considered here. 30 bootstrap draws of L6z-matching were used to generate the data for this Figure. The numbers on the top left of each panel are the median number of colour-selected QSOs that comprise the histograms. The median K-S P value is written in the top-left corner of each panel: this is the probability that the distributions are drawn from the same parent population, a measure of their similarity.

3.5 Radio source properties as a function of morphology

With a complete set of morphological classifications for the Bright colour-selected QSOs, we are able to examine the properties of their radio sources as a function of morphology. We place particular emphasis on the ‘Compact’ category since it is among these sources that rQSOs differ most from cQSOs.

In Figure 10, we compare the distributions of radio luminosity for both subsets of colour-selected QSOs in the L6z-matched sample, split by morphology and including Faint sources as well. In this Figure, and also for Figure 11 described below, we aggregate QSOs over the entire $0.8 < z < 2.5$ redshift interval to improve overall number statistics, rather than divide them into redshift bins. Therefore, the distributions plotted are subject to redshift-dependent systematics. However, because of the use of redshift-matching, we expect such systematics to apply equally to both cQSOs and rQSOs, allowing a fair comparison to be made between them.

For both colour-selected subsets, we find that the L_{144} distributions change substantially from compact to extended

sources (from left to right in Figure 10). As expected, the Extended Radio Galaxy category encompasses the most luminous radio sources in our sample. Despite this, both subsets of QSOs appear to show very similar L_{144} distributions within each morphology category. We have verified this quantitatively using two-sided K-S tests, assuming that a significant difference would yield a K-S P value < 0.05 .

For the subset of QSOs with 1.4 GHz detections from FIRST, we can also examine the radio spectral indices (α_R) as a function of morphology, as shown in Figure 11. We find that the α_R distributions between the two subsets of colour-selected QSOs are indistinguishable (K-S P values > 0.05), especially among compact radio sources where we have the best statistics. We conclude, therefore, that rQSOs with compact radio sources, while significantly more common than normal cQSOs, are not uniquely different in terms of either their intrinsic radio luminosities or their radio spectra.

4 DISCUSSION

We have combined the SDSS DR14 catalogue of QSOs and the LoTSS DR1 meter-wave radio survey to undertake a

controlled study of red QSOs to $z = 2.5$. Our strategy is to compare the radio properties of rQSOs (the 10% of QSOs with the reddest $g - i$ colours at any redshift) with typical cQSOs (the 50% of QSOs around the median $g - i$ colour at any redshift), matched in both redshift and mid-infrared luminosity (Section 2.5).

LoTSS DR1 detects radio sources that are almost an order of magnitude less luminous than the FIRST survey, allowing us to explore the properties of a different population of red QSOs than earlier work from our team (Klindt et al. 2019). The depth of LoTSS also allows us to overcome the effects, however minor, of biases introduced into our sample by the FIRST-preselection of QSOs in SDSS.

Our study confirms the results of Klindt et al. (2019), many of which are also independently verified in the recent work of Fawcett et al. (2020). At all redshifts considered in our study, the population of rQSOs are associated with an enhanced level of detectable radio emission compared to cQSOs, by a factor of ≈ 2 (Section 3.1). The enhancement is found among a preferred subset of rQSOs, those that are radio-quiet, with a radio-to-MIR luminosity ratio (\mathcal{R}) in the range of -5 to -3 (Section 3.3), and display radio morphologies that are compact at the LoTSS spatial resolution (Section 3.4). Despite their preponderance among rQSOs, these compact, radio-quiet radio sources do not distinguish themselves in terms of their radio spectral index over a decade in frequency (between 1.4 GHz and 144 MHz; Section 3.5).

In the remainder of our discussion, we explore the origin of the meter-wave radio emission in SDSS QSOs. We start with a set of Monte-Carlo simulations that use the well-known ‘‘FIR-radio correlation’’ to constrain the synchrotron emission powered by star formation (Section 4.1). We then compare the observed radio emission to that expected given the SMBH masses and accretion rates of the AGN in the QSOs. This will set the stage for a deeper understanding of the possible origins of the particular radio properties of rQSOs, leading to our conclusive statements.

4.1 The origin of meter-wave radio emission in QSOs: star formation

In the merger-driven evolutionary paradigm, red QSOs represent a phase of SMBH growth that parallels the final stage of a massive gas-rich major merger and is co-eval with, or immediately follows, a massive starburst. In this context, one may expect rQSOs to be commonly found in hosts with very high levels of SF. Since this SF is also capable of producing low-frequency radio emission (e.g., Calistro Rivera et al. 2017; Gürkan et al. 2018; Wang et al. 2019), it is valuable to test whether the higher incidence of compact radio sources among rQSOs may be produced by an underlying population of extreme starbursts that boost their radio power.

To set the stage, we consider empirical measurements of the 144 MHz emission from SF from the LOFAR study of Calistro Rivera et al. (2017). The cosmic rays accelerated in supernova explosions are responsible for the synchrotron emission from star-forming galaxies, and this component correlates well with the SF-heated IR luminosity (the ‘‘FIR-radio correlation’’, e.g., Condon 1992). From a set of star-forming galaxies with extensive multi-wavelength data in the Boötes field, Calistro Rivera et al. (2017) estimated the SFR using spectral energy distribution fits, and cali-




Figure 12. Illustration of the two forms of the IR luminosity distributions assumed in the Monte-Carlo simulations of the radio emission from SF (Section 4.1.1 and Appendix A). The lognormal distribution, plotted in green, has a mode at $\mu_{\text{IR,SF}}$ (the vertical dotted green line) and a redshift-independent width of $w_{\text{IR,SF}}$ of 0.65 dex. The distribution based on the IR luminosity function (IRLF), shown in orange, takes its overall shape from the IRLFs in Magnelli et al. (2013), with redshift-independent low and high luminosity slopes, and a low-luminosity cutoff (L_{low}). Both $\mu_{\text{IR,SF}}$ and L_{low} are monotonically related to the average IR luminosity $\langle L_{\text{IR,SF}} \rangle$, shown by the solid vertical line, which is empirically determined for QSOs from the *Herschel*-based study of Stanley et al. (2017). Both distributions shown here share the same $\langle L_{\text{IR,SF}} \rangle$, and their particular parameters are representative of cQSOs at $z \sim 1.5$.

brated the scatter and redshift evolution of the relationship between SFR and L_{144} . As a visualisation of this relationship, we plot a shaded band in Figure 4 showing the expected range in L_{144} for star-forming galaxies with $\text{SFR} = 100 M_{\odot} \text{ yr}^{-1}$, equivalent to local ultraluminous infrared galaxies (ULIRGs). The L_{144} range evolves with redshift since the nominal zeropoint of the FIR-radio correlation has a redshift dependence (Calistro Rivera et al. 2017).

Based on Figure 4, we note that a substantial subset of LoTSS-detected QSOs at $z \lesssim 1$ could be powered by extreme starbursts ($\text{SFRs} > 100 M_{\odot} \text{ yr}^{-1}$). Among the more distant QSOs, SFRs would have to be an order of magnitude higher to substantially contribute to the emission of most of the radio sources. Such systems would be among the most luminous star-forming galaxies in the Universe, the extremely rare hyperluminous infrared galaxies. Nevertheless, the 400 deg^2 of sky covered by LoTSS DR1 samples a very large cosmic volume, so the field could cover a sufficient number of such rare systems to potentially account for the nature of the observed QSO population.

4.1.1 Simulating the SFR contribution in typical QSOs

The well-established infrared-radio correlation is a valuable tool to help us constrain the importance of strong starbursts among rQSOs. Earlier studies with the *Herschel* and ALMA observatories have constrained the typical IR luminosities from SF-heated dust in QSOs. Using Monte-Carlo techniques, we can assess the degree to which the observed

radio emission of the LoTSS QSOs could arise from a star-forming population of QSO hosts with a distribution of IR luminosities motivated by these studies.

We consider two approaches to describe the IR luminosity distributions of QSOs (Figure 12). The first approach assumes that QSOs exhibit a lognormal distribution of IR luminosities arising from SF. This assumption is motivated by ALMA continuum studies of distant AGN and QSOs (Scholtz et al. 2018; Schulze et al. 2019), as well as theoretical considerations from hydrodynamic simulations (Scholtz et al. 2018). The second approach modifies the redshift-dependent infrared luminosity functions (IRLFs) of galaxies to construct IR luminosity distributions tuned to QSOs. Both approaches yield qualitatively similar conclusions, and therefore we only present here our analysis based on the assumption of a log-normal distribution, but refer the interested reader to the Appendix for the analysis using the IRLFs.

Our approach to connect the FIR distributions to the observed QSO population is as follows. Through an analysis of composite IR SEDs of SDSS QSOs from *Herschel* and *WISE* data, Stanley et al. (2017) measured the dependence of the mean IR luminosity attributable to SF as a function of redshift and AGN bolometric luminosity. In addition to a strong evolution with redshift, they reported a weak but definite trend in the average IR luminosities $\langle L_{\text{IR,SF}} \rangle$ of QSOs with bolometric luminosity, which is governed primarily through a secondary correlation with SMBH mass (Rosario et al. 2013; Stanley et al. 2017).

Using $L_{6\mu\text{m}}$ and a $6\ \mu\text{m}$ bolometric correction = 8 (Richards et al. 2006), we constrain the median $L_{6\mu\text{m}}$ of the cQSOs (i.e., typical SDSS QSOs) and read off characteristic values for $\langle L_{\text{IR,SF}} \rangle$ as a function of redshift directly from Figure 8 of Stanley et al. (2017). Since we are reliant on their measurements, we perform our own analysis in bins of redshift that parallel those used in that earlier work: $0.5 < z < 0.8$, $0.8 < z < 1.0$, $1.0 < z < 1.5$, $1.5 < z < 2.0$, $2.0 < z < 2.5$. Following this approach, for cQSOs in these respective redshift intervals, we adopt $\langle L_{\text{IR,SF}} \rangle$ of 44.6, 44.9, 45.1, 45.3, and 45.5 $\log \text{ erg s}^{-1}$.

$\langle L_{\text{IR,SF}} \rangle$ is related monotonically to the mode of the lognormal function ($\mu_{\text{IR,SF}}$) that we assume is an adequate description of the IR luminosity distribution of the cQSOs:

$$\mu_{\text{IR,SF}} = \log(\langle L_{\text{IR,SF}} \rangle) - 1.15 \times w_{\text{IR,SF}} \quad (2)$$

where $w_{\text{IR,SF}}$ is the logarithmic width of the distribution. We adopted $w_{\text{IR,SF}} = 0.65$ dex, fixed at all redshifts. This choice is motivated by the ALMA study of Scholtz et al. (2018) which constrained the empirical IR distributions of X-ray AGN in deep extragalactic fields. Very luminous QSOs at $z \sim 2$ may show a narrower distribution of SFRs (width ≈ 0.5 ; Schulze et al. 2019), but this level of difference in the assumed distribution width makes little qualitative difference to our results.

For any given interval in redshift, we identify all cQSOs in the LoTSS DR1 field, irrespective of their radio properties. Then, treating the IR distribution function in that redshift interval as a probability distribution function, we randomly assign IR luminosities to all these cQSOs. Adopting the FIR-radio correlation appropriate for 144 MHz from Calistro Rivera et al. (2017), including its redshift evolution and scatter, we are then able to convert these assigned IR

luminosities to radio luminosities, and compare the predictions of this simulated distribution of L_{144} to the actual L_{144} distribution measured for the cQSOs.

In the top panels of Figure 13, we plot the observed L_{144} distributions (filled histograms) and simulated L_{144} distribution (dashed histograms) for cQSOs in the LoTSS DR1 field. To allow for the LoTSS detection limits, we additionally calculate observed-frame 144 MHz fluxes from the simulated L_{144} of the cQSOs, using their actual redshifts and a fixed $\alpha_R = -0.7$. In the Figure, we plot using solid open histograms the simulated luminosities for the subset of cQSOs that are brighter than the nominal LoTSS DR1 5σ detection limit of 0.35 mJy.

An examination of top panels of Figure 13 indicates that, as per our nominal simulation, SF can account for as much as 50–100% of the lower luminosity radio sources ($L_{144} < 10^{26} \text{ W Hz}^{-1}$) at all redshifts under consideration. In other words, most of the radio emission from the low-luminosity radio cQSOs can arise from SF at a level consistent with the infrared luminosities of this population. Only the luminous tail of radio sources are bright enough that SF cannot power their luminosities.

In the two highest redshift intervals, the nominal simulation does not fully recover the peak of the distribution of L_{144} among radio-detected cQSOs (two right panels in the upper row of Figure 13). This may mean that pure SF is not enough to account for all the lower-luminosity radio sources among these QSOs. However, the simple application of an additional 0.2–0.3 dex to the nominal values of $\langle L_{\text{IR,SF}} \rangle$ in these redshift intervals is sufficient to reproduce the observed L_{144} distributions. Such small differences are within the scope of the systematic uncertainties inherent to the SED fitting approach of Stanley et al. (2017). We proceed with our analysis using the parameters of our nominal simulation, but note that such uncertainties on the empirical IR luminosity distributions – their averages and scatter – can quantitatively influence the comparisons of simulated and observed L_{144} distributions. Therefore, we restrict our conclusions to purely qualitative statements at this juncture, and refrain from making any quantitative measures of statistical similarity between the distributions.

We obtain similar results² if, instead of using the direct meter-wave FIR-radio correlation from Calistro Rivera et al. (2017), we adopt the 1.4 GHz FIR-radio correlation from Delhaize et al. (2017) and a canonical $\alpha_R = -0.7$. Since Calistro Rivera et al. (2017) and Delhaize et al. (2017) are independent studies that used different radio bands, different survey fields and different analytic approaches, the similarity between the results of both simulations indicates a consistent treatment of the SF contribution to the radio emission of the QSOs.

4.1.2 Adapting the simulation to rQSOs

Turning now to the red QSOs, it is clear that the very same simulation parameters cannot fully reproduce the in-

² We have to increase $\langle L_{\text{IR,SF}} \rangle$ by about 0.2 dex to reproduce the same relative numbers of LoTSS detections as our nominal simulation, but this is within the uncertainties of the measurements presented by Stanley et al. (2017)




Figure 13. Distributions of radio luminosity (L_{144}) comparing observed LoTSS-detected QSOs (filled histograms) and a simulation of the radio emission from SF within the same QSOs (open histograms). The results for cQSOs (cyan) are in the top panels and rQSOs (red) are in the bottom panels, while the populations are split into redshift bins spanning left to right as indicated. The dashed-line open histograms show the simulated L_{144} for all QSOs from the individual colour-selected samples, while the solid-line open histograms only show the subset that would be “detectable” in LoTSS at their redshifts and radio luminosities. The consistency between solid open and filled histograms indicate that SF can be an important contribution to L_{144} among the fainter end of the LoTSS radio QSOs. See Section 4.1 for more details. The dotted vertical black line marks a fixed luminosity of 10^{26} W Hz^{-1} .

created fraction of LoTSS-detected rQSOs. In order to do so, there would have to be more rQSOs than cQSOs with simulated L_{144} above the LoTSS detection threshold. Within the framework of the simulation, we achieve this by ascribing a higher average SFR to the rQSOs, by increasing the mode $\mu_{\text{IR,SF}}$ of the lognormal distribution that describes the IR luminosity of the rQSOs. After some experimentation, we found that a simple increase in $\mu_{\text{IR,SF}}$ of 0.4 dex over the nominal value of the cQSOs is enough to account for the redshift-averaged enhancement of the rQSO detection rate in LoTSS (a factor of ≈ 1.9). Implementing this increased SFR, we find that SF could also account for as much as 50–100% of the low-luminosity rQSOs (lower panels of Figure 13).

LoTSS detects the steeply declining exponential tail of the star-forming QSO population, the starburst systems. Is there a maximum radio luminosity that SF can realistically power among the QSOs? In each panel of Figure 13, we identify a radio luminosity of $L_{144} = 10^{26}$ W Hz^{-1} as a vertical dotted line. This luminosity is chosen as a benchmark for further analysis, because it marks a rough threshold beyond which SF cannot account for the majority of the radio emission. From the simulation, less than 5% of radio-detected rQSOs with $L_{144} > 10^{26}$ W Hz^{-1} at any redshift can be powered by SF, even with the tweaks to the IR luminosity distributions that are needed to reproduce the L_{144} distributions of low-luminosity radio sources. Therefore, if it can be demonstrated that cQSOs and rQSOs more luminous than this threshold still show substantial differences, SF is un-

likely to be root cause of the particular radio properties of red QSOs.

4.1.3 Behaviour as a function of radio-loudness \mathcal{R}

In our Monte-Carlo simulation, we probabilistically associated radio luminosities to the observed population of QSOs and applied the flux limits of the LoTSS survey to determine the detectable fraction. We can go a step further and explore whether the derived difference in the median SFRs of rQSOs over cQSOs (0.4 dex) can also reproduce the variation of the rQSO detection enhancement with the radio-to-MIR luminosity ratio (\mathcal{R} ; Section 3.3 and Figure 7).

In Figure 14, we show as a shaded region the predicted range of the radio detection enhancement of rQSOs over L6z-matched cQSOs as a function of \mathcal{R} . As a baseline, we also show the median predictions of a simulation in which the rQSOs share the same IR luminosity distributions as the cQSOs (solid magenta line). The observed enhancement is shown with the data points, in a similar fashion to Figure 7. Both the simulated and observed data points plotted here are summarised from 30 independent bootstrap draws of an L6z-matched sample, which allows us to display a representative range of enhancements shown by the L6z-matched QSOs.

In comparing the shaded regions to the points, it must be kept in mind that the simulation does not account for the luminous radio-loud sources which are undoubtedly powered by the AGN. Therefore, while the observed data points extend beyond the radio-quiet/loud boundary (the vertical




Figure 14. The ratio of the detection rate in LoTSS for rQSOs vs. L6z-matched cQSOs plotted in bins of the ‘radio-loudness’ \mathcal{R} , comparing the observed ratio (data points) against the predictions of our fiducial Monte Carlo simulation of SF in QSOs (shaded magenta regions; see Section 4.1 for details). These data are summarised from 30 bootstrap L6z-matching runs into the working sample in order to overcome some of the sample variance inherent in the L6z-matching process. The vertical dashed line at $\mathcal{R} = -4.5$ marks the threshold between systems that are ‘radio-loud’ (to the right) and ‘radio-quiet’ (to the left). For reference, the solid magenta line shows the median trend for a version of the simulation where rQSOs have the same IR distributions as the cQSOs.

dashed line), the simulation cannot reproduce these powerful systems. Our comparison is relevant only in the range of $\mathcal{R} = -6$ to -4 , where SF can make an important contribution to the detectable radio emission.

The simulation in which rQSOs have a increased modal SFR by 0.4 dex can broadly reproduce the enhancements that we observe, and capture some of the dropoff towards low values of \mathcal{R} . At the lowest values of $\mathcal{R} < -5$, the simulation does predict a higher enhancement than we see in the data. However, since the exact form of the dropoff is sensitive to details such as the assumed form of the IR luminosity distribution and the variation of $\langle L_{\text{IR,SF}} \rangle$ with QSO bolometric luminosity, we only concentrate here on broad trends rather than full quantitative agreement. It must also be remembered that some of this consistency is because the SFR offset was *tuned* to reproduce the observed redshift-averaged enhancement in the detection fraction of rQSOs (Section 4.1.2).

We conclude that the simulation is qualitatively successful at reproducing the behaviour of rQSOs as a function of \mathcal{R} , but caution that this is not an independent test since it is sensitive to the tuning of the simulation parameters.

4.2 The origin of meter-wave radio emission in QSOs: accretion-powered emission

Our results indicate that the most substantial enhancement of the radio detection fraction in rQSOs is found among compact radio sources (Section 3.4). We have examined SF as a possible driver of the radio emission in such systems in the previous sub-section. Here we consider whether the




Figure 15. Observed 144 MHz radio luminosities (Y-axis) vs. predicted radio luminosities (X-axis) using the supermassive black hole ‘Fundamental Plane’ relationship of Merloni, Heinz & di Matteo (2003). Only QSOs with compact radio morphologies in LoTSS and black hole masses from the SDSS DR12 QSO catalogue (Kozłowski 2017) are shown here. See Section 4.2 for more details.

observed radio emission is consistent with an origin from AGN-related processes.

The most widely known mechanism of AGN-powered radio continuum emission is through bulk relativistic outflows (jets) ejected from the nucleus. Another mechanism recently proposed as an important contribution in radio-quiet AGN is through particle acceleration in fast shocks driven by thermal winds (e.g., Zakamska & Greene 2014). Though the importance of the latter process is still under debate, it is of particular interest in rQSOs because powerful obscured quasars are often associated with the kinematic signatures of strong winds (Brusa et al. 2015; Zakamska et al. 2016; Klindt et al. 2019). In the analysis below, we use energetic arguments based on studies of traditional jetted radio AGN, and briefly link in the expectations from the wind-shock scenario as well.

It is worth noting at this stage that strongly beamed synchrotron emission, such as that found in blazars or flat spectrum radio quasars, cannot account for the rQSO population (Klindt et al. 2019). If such synchrotron emission was bright enough to contaminate the optical band in our sources, their predicted radio luminosities would be orders of magnitude higher than we observe. Instead, among the most luminous radio QSOs we only find a low radio detection rate enhancement, in direct contrast to the expectations based on pervasive beamed-synchrotron emission.

In the well-known empirical work of Merloni, Heinz & di Matteo (2003) which studied correlations in emission properties across both stellar mass and SMBH accretion systems, the authors reported the discovery of a ‘Fundamental Plane’ of black hole activity plane, a hypersurface in the three-dimensional space of disc radiative emission (X-rays), core synchrotron emission (5 GHz) and black hole

mass (M_{BH}). While there is debate about the nature and origin of this Fundamental Plane, it serves a useful purpose in our study by providing a reference level for the empirical radio emission expected from radiatively-efficient AGN.

A subset of the LoTSS DR1 radio QSOs have black hole masses estimated by Kozłowski (2017), as part of a value-added spectral analysis of QSOs in the SDSS DR12 release. Concentrating only on the QSOs with compact radio morphologies, which encompass all core-dominated radio sources, we are able to translate their $L_{6\mu\text{m}}$ and L_{144} into the variables that specify the Fundamental Plane of Merloni, Heinz & di Matteo (2003). To do this, we assume a fixed ratio of 0.5 dex between the X-ray and 6 μm luminosities of our QSOs (Lutz et al. 2004) and take our usual approach of deriving or assuming a value of the radio spectral index. This calculation gives us an expectation for the typical amount of radio emission these AGN should generate if they were translating their power into relativistic plasma.

In Figure 15, we plot the expected L_{144} from the Fundamental Plane calculation against the measured L_{144} for radio QSOs from our working sample. A majority of both cQSOs and rQSOs with compact radio emission have radio luminosities that are comparable to or more powerful than the predictions of the fundamental plane. As the Fundamental Plane is only applicable to the nuclear or core radio emission, we interpret the scatter of QSOs with higher observed L_{144} as likely arising from extended jetted emission that is unresolved by LoTSS at the 6'' angular resolution of the survey. From the analysis summarised in Figure 15, we can conclude that the radio luminosities of the QSOs are at least broadly consistent with an origin from AGN.

In Section 3.3, we identified a peak in the radio detection enhancement among rQSOs at a preferred value of the radio-loudness parameter $\mathcal{R} \approx -4$. Under the hypothesis that the radio emission in QSOs is powered by an AGN at these intermediate values of radio-loudness, it is useful to case the characteristic value of \mathcal{R} in more physical terms. The radio luminosity of a relativistic outflow is believed to be connected to its total kinetic power, though the exact relationship is only known in dense cluster environments where both quantities can be independently measured. Nevertheless, if we adopt such a relationship to calculate the kinetic power of a radio source (e.g., Equation 6 from Merloni & Heinz 2007), and take a standard 6 μm bolometric correction for QSOs (≈ 8 , Richards et al. 2006), we can convert \mathcal{R} to a ratio of kinetic-to-bolometric power in a AGN-powered radio source as follows:

$$\log \frac{\text{Kinetic}}{\text{Bolometric}} \approx 0.8\mathcal{R} - 0.2L_{6\mu\text{m}} + 11.5 \quad (3)$$

The mild dependence of $L_{6\mu\text{m}}$ is because of the sub-linear relationship derived by Merloni & Heinz (2007) between the kinetic and synchrotron power of a radio source.

The wind-shock scenario of Zakamska & Greene (2014) assumes a conversion efficiency between the wind kinetic energy and synchrotron radio power that is equivalent to starburst superwinds, a typical value of ≈ 4.5 dex. This is quite close to the ratio in relativistic jetted outflows reported by Merloni & Heinz (2007). Therefore, our inferences are not very sensitive to the driving mechanism of the synchrotron emission, certainly within the large systematic uncertainties that go into such calculations.

From Equation 3, we estimate that a QSO with $\log L_{6\mu\text{m}} = 44.5\text{--}46.5 \text{ erg s}^{-1}$ lying on the $\mathcal{R} = -4.5$ division puts out about 3-10% of its power in kinetic form. Interestingly, this is approximately the range expected for the efficiency of AGN feedback in successful simulations of galaxy evolution (e.g. Bower et al. 2006; Anglés-Alcázar et al. 2017).

4.3 What is responsible for the distinct radio properties of red QSOs?

Our simulation of the radio luminosities attributable to SF in QSO hosts indicates that a substantial proportion of the lower-luminosity radio sources detected by LoTSS, including a majority of those that lie below the FIRST detection limit, could be powered by SF. On the other hand, the luminosities of the compact radio sources, the category most over-represented among rQSOs, are quite consistent with an AGN origin. In this subsection, we bring together various lines of evidence laid out in this work to explore whether the distinguishing properties of rQSOs favour an enhanced level of SF or arise instead from a special AGN-powered population found mostly among rQSOs.

From the simulation in Section 4.1, we concluded that a difference in the median amount of SF at the level of 0.4 dex is enough to distinguish the L_{144} distributions of rQSOs and cQSOs in LoTSS. The SF in these QSOs will, in the most part, be found only on scales of their host galaxies, i.e. a few – 10 kpc. Therefore, we expect the star-forming radio emission to be unresolved by LoTSS, consistent with our result that the enhanced rQSO detection rate is only found among compact radio sources, in which the central radio emission is the dominant component of the total intensity.

Meter-wave radio emission from SF is primarily optically-thin synchrotron emission from supernova remnants and shows a very similar range of meter-wave spectral indices as AGN radio jets. Therefore, we would not expect strong differences in the spectral indices of SF-dominated radio QSOs compared to those that may be dominated by compact radio jets (e.g., Calistro Rivera et al. 2017). This is consistent with the similar spectral indices as a function of morphology that we find in our sample (Figure 11).

However, a closer look at the L_{144} distributions may imply a different interpretation. A valuable bit of insight from the simulations of Section 4.1 is that radio sources with $L_{144} > 10^{26} \text{ W Hz}^{-1}$ are very unlikely to arise from SF. Yet, from Figure 10, approximately half of the colour-selected QSOs with compact radio sources have radio luminosities above this threshold. If SF was the only driver for the different radio properties of the rQSOs, the more luminous compact radio sources, which cannot be dominated by SF, should be found at roughly the same rate of incidence in cQSOs and rQSOs.

In the right panel of Figure 8, we compare the distribution of morphologies for colour-selected QSOs with $L_{144} > 10^{26} \text{ W Hz}^{-1}$. Since this restricts our analysis only to the brighter LoTSS radio QSOs, the fraction of ‘Faint’ sources drops drastically as expected, while the sources in the other morphological categories suffer smaller reductions in their fractions because they are generally brighter. The main point to note from this Figure is that, despite the application of our L_{144} threshold, the relative enhancement in the fraction of rQSOs over cQSOs remains about the same

in both the ‘Compact’ and the ‘Marginally Extended’ categories. Therefore, the radio luminosities of the special population of compact radio-quiet sources found among rQSOs extends beyond the range where SF can be reasonably expected to explain the differences. This supports the notion that the primary driver of the enhancements seen among rQSOs is from a population of compact, AGN-powered radio sources.

In Figure 7, we examined the observed variation of the radio detection enhancement in rQSOs vs. the radio-loudness parameter \mathcal{R} , finding a more pronounced peak in the enhancement at $\mathcal{R} \sim -4$, with a strong drop off towards lower values of \mathcal{R} . This suggests that the special population of rQSOs responsible for the enhancement occupies a preferred radio-loudness, slightly below the canonical boundary between radio-loud and radio-quiet systems. How do we interpret this result in terms of the interplay between SF and AGN-powered radio emission in QSOs?

Our analysis in Figure 14 shows that a scenario where rQSOs have the same level of AGN-powered radio emission but a higher level of SF (by 0.4 dex) can qualitatively explain the trend of increasing rQSO detection enhancement with \mathcal{R} seen among radio-quiet QSOs. On the other hand, from our morphological assessment (Figure 8 and discussion above), we find a substantial detection enhancement even among radio sources which are luminous enough that they cannot be reasonably powered by SF.

We can reconcile these conflicting points of evidence by considering an alternative explanation for the trends with \mathcal{R} : the differences between rQSOs and cQSOs arise from an underlying population of radio-intermediate AGN-powered radio sources that are over-represented among rQSOs. In this scenario, both colour-selected QSO subpopulations have similar SFR distributions, and the dropoff that we see towards low values of \mathcal{R} in Figure 7 is because very radio-quiet systems, whether red or blue, are increasingly dominated by SF.

Adequately distinguishing between these two explanations requires a careful assessment of the level of SF in rQSOs *vis-à-vis* cQSOs. In the accompanying study of Fawcett et al. (2020), we have used multi-wavelength constraints to identify SF-dominated SDSS QSOs within the 2 deg² COSMOS extragalactic deep field. We find that such systems are invariably very radio-quiet ($\mathcal{R} < -5$), and that AGN-dominated radio-quiet QSOs potentially still show an enhancement in their radio detection rates. This bolsters the argument that differences in SF are not the root cause of the enhanced radio emission among rQSOs.

On-going radio studies at high angular resolution, designed to resolve the kpc-scale emission in rQSOs and cQSOs, may be able to shed more light on this phenomenon. In addition, we are undertaking a deep rest-frame far-IR imaging study of SDSS colour-selected QSO with the Atacama Large Millimetre Array (ALMA) which will help us directly measure the IR-based SFRs in rQSOs and cQSOs, and conclusively determine whether red QSOs are associated with enhanced SF.

Regardless of the real situation, our results require that either the distributions of SFR or AGN-driven radio emission in red QSOs are fundamentally different from typical QSOs. This implies that environments of these systems, whether on nuclear, host galaxy, or larger scales, are also

fundamentally different. A thorough exploration of these differences is essential in understanding the role of the red QSO phase in the co-evolution of galaxies and supermassive black holes.

5 CONCLUSIONS

Using the SDSS DR14 QSO catalog and the LoTSS DR1 survey, we compare the low-frequency radio properties of red QSOs (rQSOs) and typical QSOs (cQSOs). We use *WISE*-based MIR photometry and a statistical matching technique to control for differences in redshift and intrinsic luminosities between the two subsamples. We find:

(i) The population of rQSOs are $\times 3$ more likely to be detected in LoTSS than equivalent cQSOs (Section 3.1), consistent with an earlier FIRST-based study from our team (Klindt et al. 2019).

(ii) The sources that drive the higher radio detection fractions of rQSOs have radio morphologies that are primarily compact (Section 3.4), further confirming the findings of Klindt et al. (2019). The radio luminosities and spectral indices of the compact rQSOs are indistinguishable from cQSOs (Section 3.5)

(iii) Much of the difference between rQSOs and cQSOs can be attributed to an excess population of radio-intermediate sources, with a radio-loudness parameter $\mathcal{R} \approx -4$ (Section 3.3, where \mathcal{R} is a measure of the MIR-to-radio power).

(iv) By statistically modelling the expected amount of star formation (SF) in the host galaxies of QSOs, we demonstrate that SF is a major source of radio emission in the radio-quiet regime in both rQSOs and cQSOs (Section 4.1)

(v) Bringing together various lines of evidence, including the characteristic radio luminosities at which rQSOs show the most pronounced radio excess, we conclude that their special properties likely arise from processes that are linked to their AGN, rather than due to higher levels of SF (Section 4.3). Our results confirm, quite spectacularly, that dust-reddened QSOs are a fundamentally different population of AGN than the more common blue QSOs.

ACKNOWLEDGEMENTS

We thank the anonymous referee for their valuable input in the final stages of this work. D.R. and D.M.A. acknowledge the support of the UK Science and Technology Facilities Council (STFC) through grants ST/P000541/1 and ST/T000244/1. L.K. acknowledges the support of a Faculty of Science Durham Doctoral Scholarship. Funding for the SDSS has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is www.sdss.org. SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration. We make use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the

Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration.

REFERENCES

- Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017, *MNRAS*, 464, 2840
- Banerji M., Alaghband-Zadeh S., Hewett P. C., McMahon R. G., 2015, *MNRAS*, 447, 3368
- Banerji M., McMahon R. G., Hewett P. C., Alaghband-Zadeh S., Gonzalez-Solares E., Venemans B. P., Hawthorn M. J., 2012, *MNRAS*, 427, 2275
- Becker R. H., White R. L., Helfand D. J., 1995, *ApJ*, 450, 559
- Best P. N., Heckman T. M., 2012, *MNRAS*, 421, 1569
- Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, *MNRAS*, 370, 645
- Brusa M. et al., 2015, *MNRAS*, 446, 2394
- Calistro Rivera G. et al., 2017, *MNRAS*, 469, 3468
- Cameron E., 2011, *Publications of the Astronomical Society of Australia*, 28, 128
- Condon J. J., 1992, *ARA&A*, 30, 575
- Croom S. M., Smith R. J., Boyle B. J., Shanks T., Miller L., Outram P. J., Loaring N. S., 2004, *MNRAS*, 349, 1397
- Delhaize J. et al., 2017, *A&A*, 602, A4
- Draine B. T., 2003, *ARA&A*, 41, 241
- Fan X. et al., 2001, *AJ*, 122, 2833
- Fawcett V. A., Alexander D. A., Rosario D. J., Klindt L., Fotopoulou S., Lusso E., Morabito L., Calistro Rivera G., 2020, *MNRAS*, submitted
- Glikman E., Helfand D. J., White R. L., Becker R. H., Gregg M. D., Lacy M., 2007, *ApJ*, 667, 673
- Glikman E., Simmons B., Mailly M., Schawinski K., Urry C. M., Lacy M., 2015, *ApJ*, 806, 218
- Gürkan G. et al., 2019, *A&A*, 622, A11
- Gürkan G. et al., 2018, *MNRAS*, 475, 3010
- Helfand D. J., White R. L., Becker R. H., 2015, *ApJ*, 801, 26
- Ivezić Ž. et al., 2002, *AJ*, 124, 2364
- Kellermann K. I., Sramek R., Schmidt M., Shaffer D. B., Green R., 1989, *AJ*, 98, 1195
- Kimball A. E., Kellermann K. I., Condon J. J., Ivezić Ž., Perley R. A., 2011, *ApJ*, 739, L29
- Klindt L., Alexander D. M., Rosario D. J., Lusso E., Fotopoulou S., 2019, *arXiv e-prints*, arXiv:1905.12108
- Kozłowski S., 2017, *ApJS*, 228, 9
- Lake S. E., Wright E. L., Petty S., Assef R. J., Jarrett T. H., Stanford S. A., Stern D., Tsai C. W., 2012, *AJ*, 143, 7
- Lu Y., Wang T., Zhou H., Wu J., 2007, *AJ*, 133, 1615
- Lutz D., Maiolino R., Spoon H. W. W., Moorwood A. F. M., 2004, *A&A*, 418, 465
- Magnelli B. et al., 2013, *A&A*, 553, A132
- Merloni A., Heinz S., 2007, *MNRAS*, 381, 589
- Merloni A., Heinz S., di Matteo T., 2003, *MNRAS*, 345, 1057
- Miley G., 1980, *ARA&A*, 18, 165
- Netzer H., 2015, *ARA&A*, 53, 365
- Panessa F., Baldi R. D., Laor A., Padovani P., Behar E., McHardy I., 2019, *Nature Astronomy*, 3, 387
- Pâris I. et al., 2018, *A&A*, 613, A51
- Planck Collaboration et al., 2016, *A&A*, 594, A13
- Richards G. T. et al., 2002, *AJ*, 123, 2945
- Richards G. T. et al., 2003, *AJ*, 126, 1131
- Richards G. T. et al., 2006, *ApJS*, 166, 470
- Rosario D. J. et al., 2013, *A&A*, 560, A72
- Sandage A., Wyndham J. D., 1965, *ApJ*, 141, 328
- Schneider D. P. et al., 2010, *AJ*, 139, 2360
- Scholtz J. et al., 2018, *MNRAS*, 475, 1288
- Schulze A. et al., 2019, *MNRAS*, 488, 1180
- Shimwell T. W. et al., 2017, *A&A*, 598, A104
- Shimwell T. W. et al., 2019, *A&A*, 622, A1
- Stanley F. et al., 2017, *MNRAS*, 472, 2221
- Tsai A. L., Hwang C. Y., 2017, *ApJ*, 842, 57
- Urrutia T., Becker R. H., White R. L., Glikman E., Lacy M., Hodge J., Gregg M. D., 2009, *ApJ*, 698, 1095
- Urrutia T., Lacy M., Becker R. H., 2008, *ApJ*, 674, 80
- Wang L. et al., 2019, *A&A*, 631, A109
- White R. L., Helfand D. J., Becker R. H., Glikman E., de Vries W., 2007, *ApJ*, 654, 99
- White R. L., Helfand D. J., Becker R. H., Gregg M. D., Postman M., Lauer T. R., Oegerle W., 2003, *AJ*, 126, 706
- Williams W. L. et al., 2019, *A&A*, 622, A2
- Zakamska N. L., Greene J. E., 2014, *MNRAS*, 442, 784
- Zakamska N. L. et al., 2016, *MNRAS*, 459, 3144

APPENDIX A: SIMULATING THE SFR CONTRIBUTION IN TYPICAL QSOs: THE IRLF APPROACH

In Section 4.1.1, we described our approach of simulating the radio emission from SF in QSOs. In our primary simulations, we assumed that the IR luminosity distributions of SF heated dust in QSOs could be adequately described with lognormal functions (Figure 12). Alternately, one could instead assume that QSO host galaxies share the shape of the IR luminosity function (IRLF) of the co-eval galaxy population, and share in the general evolution of the IRLF. Here we consider a simulation based on the galaxy IRLF, and apply it to the SDSS DR14 cQSOs and rQSOs in the LoTSS DR1 region.

In the last decade, several studies with the FIR *Herschel* Space Observatory have accurately constrained the redshift-dependent IRLFs of galaxies over areas of sky comparable to the LoTSS HETDEX field (hundreds of deg²). These IRLFs empirically capture the space density of the luminous tail of the infrared galaxy population, which are responsible for most of the SF dominated systems detectable in LoTSS beyond the local Universe (Figure 4). For the purposes of our calculations, we use the IRLFs from Magnelli et al. (2013), a composite study that combined *Herschel* constraints from both large-area and deep fields. These IRLFs are described with a simple broken power-law function, with fixed low-and high-luminosity slopes of -0.6 and -2.2 that connect at a characteristic break luminosity (L_{\star}^{IR}). Since we treat these IRLFs as probability distribution functions that integrate to unity, we do not use their absolute normalisations in our analysis.

The integral of the IRLF described above is unbounded




Figure A1. Distributions of radio luminosity (L_{144}) comparing observed LoTSS-detected QSOs (filled histograms) and a simulation of the radio emission from SF within the same QSOs (open histograms). This version of the simulation assumes a distribution of IR luminosities in QSOs that follows the IR luminosity function of galaxies with a variable low-luminosity cutoff (see Section A for details). The results for cQSOs (cyan) are in the top panels and rQSOs (red) are in the bottom panels, while the populations are split into redshift bins spanning left to right as indicated. The dashed-line open histograms show the simulated L_{144} for all QSOs from the individual colour-selected samples, while the solid-line open histograms only show the subset that would be “detectable” in LoTSS at their redshifts and radio luminosities. The dotted vertical black line marks a fixed luminosity of 10^{26} W Hz^{-1} .

if the lower integration limit $\rightarrow -\infty$. Therefore, in order to use it as a distribution function, we impose a finite lower limit to the integral, which functionally acts as a low luminosity cutoff on the IRLF (L_{low}). For a given L_{\star}^{IR} , L_{low} is monotonically related to the average IR luminosity of the modified IRLF $\langle L_{\text{IR,SF}} \rangle$ (see Figure 12 for a visualisation). Stanley et al. (2017) have empirically determined $\langle L_{\text{IR,SF}} \rangle$ for SDSS QSOs using *Herschel* and *WISE* stacked photometry. We use their relations to set L_{low} as a function of redshift and the median bolometric luminosities of given QSO subpopulation.

For the IRLF simulation, we follow an approach similar to our primary simulation to estimate the 144 MHz radio luminosity (L_{144}) distributions of QSOs. We refer the reader to Section 4.1.1 for the essential details. The key differences to the aforementioned approach is (a) that we use a different family of IR luminosity distribution functions, and (b) that we adopt a different redshift binning scheme, to remain consistent with the binning used by Magnelli et al. (2013) for their determination of the IRLF and its evolution: $0.5 < z < 0.7$, $0.7 < z < 1.0$, $1.0 < z < 1.3$, $1.3 < z < 1.8$, and $1.8 < z < 2.3$.

In Figure A1, we compare the distributions of L_{144} predicted by IRLF simulation, for the full sample and the ‘detectable’ subset, against the observed L_{144} distributions of cQSOs and rQSOs from the working sample. In this version of the simulation, we achieve the higher incidence of starbursts required among rQSOs by increasing the cutoff luminosity L_{low} in the rQSO simulation by 0.4 dex, which repro-

duces the higher redshift-averaged LoTSS detection fraction seen in this subset of colour-selected QSOs.

A direct comparison of Figure A1 and Figure 13 indicates that both simulations yield the same basic qualitative conclusion that most, if not all, of the lower luminosity radio QSOs in LoTSS ($L_{144} < 10^{26}$ W Hz^{-1}) can be powered by SF. They agree in this regard despite the different underlying IR luminosity distributions that we assumed. However, there are some important differences. The simulation based on the IRLF overestimates the number of detectable QSOs at $z < 1.3$ by factors of a few. This can be significantly mitigated by reducing the values of $\langle L_{\text{IR,SF}} \rangle$ used in this simulation by 0.3-0.4 dex. However, these changes to the mean IR luminosity of the QSOs would make them inconsistent with the empirical measurements of Stanley et al. (2017).

Another, probably more realistic, way to bring the simulated and observed distributions in better agreement is to allow a more gradual low luminosity drop off in the shape of the modified IRLF, in place of a sharp cutoff at L_{low} . However, as one may infer from Figure 12, this leads to a functional form that resembles the lognormal function adopted in the standard simulations of Section 4.1.1. For this reason, we do not pursue the IRLF-based simulations for any further scientific analysis beyond the simple confirmation of our inferences regarding the importance of SF in powering low-luminosity radio QSOs in LoTSS.