

More Numerous and Bluer: Companions of the Brightest Little Red Dots Differ from Control Galaxies at $> 3\sigma$, Hinting at Synchronized Black Hole Seed Formation

FABIO PACUCCI ¹ AND C. MEGAN URRY ²

¹*Center for Astrophysics | Harvard & Smithsonian, 60 Garden St, Cambridge, MA 02138, USA*

²*Department of Physics, Yale University, New Haven, CT, USA*

ABSTRACT

The Little Red Dots (LRDs) are a population of compact, red sources, detected abundantly at $4 \lesssim z \lesssim 9$. While their nature is still debated, many models suggest they are powered by accreting massive black holes. Multiple studies have detected blue companions in their vicinity, motivating our analysis: we compare 253 LRDs uniformly selected across several JWST deep fields against > 7000 matched controls in the same mosaics. We find that 27.7% of the brightest LRD quartile host a blue ($\beta < -1$) companion at $0.5 - 2$ kpc, versus 11.6% of redshift-matched controls: a $\times 2.39$ excess, significant at 3.3σ . This phenomenon is limited to the brightest LRDs, as the full LRD sample pairs at a rate indistinguishable from controls. Moreover, the excess is confined to the close-pair window (< 2 kpc), outside which LRD environments are ordinary. The companions of the brightest LRDs are overwhelmingly (92%) blue, and significantly bluer, with a median $\beta \approx -2.32$, versus $\beta \approx -1.92$ for those of controls, rejecting a common parent distribution at 3.4σ . The signal is thus directional in three independent ways: luminosity, separation, and companion color, all expected for synchronized-pair seed formation. Therefore, we conclude that the association between the brightest LRDs and their close blue companions is physical, i.e., the two phenomena are causally connected: either one drives the other, or both emerge from the same synchronized formation event. A natural explanation for why the most luminous LRDs have nearby, young, star-forming neighbors is that the radiation from these neighbors enabled the direct collapse of the massive black holes powering the LRDs themselves.

Keywords: Active galaxies (17) — Early universe (435) — Galaxy evolution (594) — Supermassive black holes (1663) — Surveys (1671)

1. INTRODUCTION

Little Red Dots (LRDs) are a class of objects discovered by the *James Webb Space Telescope* (JWST) during its first deep-field observations (D. D. Kocevski et al. 2023; J. Matthee et al. 2024; R. Maiolino et al. 2024). The observational features of LRDs are peculiar. First, they are very compact, typically unresolved, sources of light with effective radii $\lesssim 100$ pc (J. F. W. Baggen et al. 2023). Second, they are characterized by a V-shaped spectral energy distribution (SED) with increased emission at longer and shorter wavelengths (V. Kokorev et al. 2024; D. D. Kocevski et al. 2025; D. J. Setton et al. 2025). These observational properties are used to effectively select LRDs (D. D. Kocevski et al. 2025).

Originally interpreted as compact galaxies with high stellar densities (I. Labbé et al. 2023), the paradigm began to shift when it was discovered that most LRDs are characterized by broad emission lines (J. E. Greene et al. 2024), typically a signature of a massive black hole. These black holes appear extremely overmassive compared to the stellar mass of their hosts (F. Pacucci et al. 2023; I. Juodžbalis et al. 2024; B. L. Jones et al. 2025; A. R. Gupta et al. 2026), although many unknowns remain regarding the origin of their broad lines, leading to the possibility that their masses may be overestimated (V. Rusakov et al. 2026; E. Modelevsky et al. 2026).

Currently, many interpretations for the nature of the LRDs rely on the presence of a massive black hole at their center (R. P. Naidu et al. 2025; F. Pacucci et al. 2026), although alternative explanations exist, such as involving globular clusters in formation (J. Chisholm et al. 2026), supermassive stars with temperatures that can reproduce the LRD spectral features (D. Nandal &

A. Loeb 2026), or repeating tidal disruption events (J. Bellovary 2025).

An important piece of the puzzle in uncovering the nature of the LRD was recently unearthed with the discovery of abundant blue companions at projected separations of $0.5 - 5$ kpc (J. F. W. Baggen et al. 2026; A. J. Barger & L. L. Cowie 2026). This finding may be crucial: if the blue companions occur around LRDs much more often than in standard galaxies, there may be a causal connection between their presence and the LRD phenomenon. In other words, the LRDs may *cause* the presence of the blue companions, or *be caused* by it.

In particular, J. F. W. Baggen et al. (2026) suggested that UV radiation from star formation, active in the blue companions, would enable the direct collapse of a massive black hole inside the nearby LRD. Direct collapse black holes (DCBHs, A. Loeb & F. A. Rasio 1994; S. P. Oh & Z. Haiman 2002; V. Bromm & A. Loeb 2003; M. Volonteri & M. J. Rees 2005; M. C. Begelman et al. 2006; G. Lodato & P. Natarajan 2006) were originally proposed as a fast formation pathway to explain the detection of $\gtrsim 10^9 M_\odot$ supermassive black holes (SMBHs) at $z \gtrsim 6$ (X. Fan et al. 2001a; D. J. Mortlock et al. 2011; F. Wang et al. 2021), which would otherwise require challenging and sustained accretion rates, possibly at super-Eddington levels. DCBHs, formed from the collapse of pristine (i.e., not metal-enriched), atomic-cooling (K. Omukai 2001) halos, would have an initial mass in the range $\sim 10^4 - 10^6 M_\odot$ (A. Ferrara et al. 2014), significantly speeding up their growth process to reach the $\gtrsim 10^9 M_\odot$ SMBH masses.

However, a condition *sine qua non* for the formation of a DCBH is the irradiation of the primordial, atomic-cooling halo by a significant flux of Lyman-Werner (LW) photons ($11.2 - 13.6$ eV, B. T. Draine & F. Bertoldi 1996), which efficiently dissociates H_2 molecules via the Solomon process, hence preventing cooling and fragmentation. Three ingredients are thus necessary to form a DCBH: (i) pristine gas, (ii) an atomic-cooling halo, (iii) irradiation by a sufficient flux of LW photons. Although estimates of the minimum required LW flux vary widely (K. Omukai 2001; M. Dijkstra et al. 2008; M. A. Latif et al. 2014; M. Habouzit et al. 2016; F. Pacucci et al. 2017), this condition is considered the most constraining for the collapse, as it would limit the locations where DCBHs can form. Recent simulations, however, indicate that substantially lower fluxes may suffice (R. Hazlett et al. 2026).

Notably, J. F. W. Baggen et al. (2026) found that the majority of the blue companions in their LRD sample already provide LW fluxes above the classical threshold; if the lower threshold of R. Hazlett et al. (2026) is adopted,

every companion in their sample would qualify. The presence of these blue companions is reminiscent of the “synchronized pair” scenario of E. Visbal et al. (2014), in which two atomic-cooling halos collapse nearly simultaneously, and the intense LW radiation from the first, star-forming halo triggers the formation of the DCBH in the companion halo. Independently, F. Pacucci et al. (2026) recently suggested, using a series of radiation-hydrodynamic simulations, that the observational properties of LRDs can be entirely explained by identifying them with DCBHs.

However, normal galaxies tend to be clustered. Hence, the key question is: are the LRDs special? It is entirely possible that any object—LRD or ordinary galaxy at the same redshift—might have blue companions nearby. To confirm the exciting possibility that blue companions actually provide the necessary LW photons to trigger the DCBH formation, it is crucial to show that LRDs are significantly more likely to have blue companions nearby compared to a matched control sample of galaxies.

This comparison must use an appropriate projected separation. The $0.5 - 5$ kpc range over which companions are typically reported (e.g., J. F. W. Baggen et al. 2026) mixes two physically distinct regimes: a close-pair regime, where two halos are directly interacting (or even co-forming), and a wider environment, more connected to ordinary galaxy clustering. The distinction is crucial for the DCBH interpretation, because the LW flux received by the prospective collapsing halo at a distance d scales as d^{-2} : only the closest separations enable the collapse.

Another important aspect deserving careful design is the choice of the control sample. The abundance of detected neighbors depends on many parameters, including redshift, brightness, and source compactness. Remarkably, these are properties in which the LRDs differ systematically from typical galaxies at the same cosmic epoch (V. Kokorev et al. 2024; D. D. Kocevski et al. 2025; A. J. Taylor et al. 2025).

In this Letter, we compare a sample of 253 LRDs to a carefully matched sample of galaxies, constructed as described in §2. Our results are described in §3 and discussed in §4. Throughout this Letter, we adopt a flat Λ CDM cosmology, with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.3$; across the redshift range of our sample, $1''$ corresponds to $4.6 - 7.0$ kpc of projected separation.

2. DATA AND METHODS

We design a controlled experiment to compare the occurrence of blue close companions around LRDs versus normal galaxies. Specifically, we extract identical companion statistics around: (i) LRDs, (ii) Tier-1 con-

trol galaxies, matched in redshift, (iii) Tier-2 control galaxies, matched in redshift, brightness, and compactness, and (iv) random blank-sky positions. All measurements are performed on the same mosaics, with the same pipeline and identical parameters, so any difference between samples cannot arise from the measurement itself. In the following, we describe the process in detail.

2.1. The LRD Sample

Our LRD targets are derived from the photometric catalog in [D. D. Kocevski et al. \(2025\)](#), which applies uniform compactness and color selection criteria across several deep JWST surveys: we use its sources in CEERS ([S. L. Finkelstein et al. 2025](#)), the JADES GOODS-S footprint ([D. J. Eisenstein et al. 2026](#)), NGDEEP ([M. B. Bagley et al. 2024](#)), and the PRIMER COSMOS and UDS mosaics ([J. S. Dunlop et al. 2021](#)).

We select sources in the redshift range $4.0 \leq z \leq 8.5$, within which the NIRCcam short-wavelength bands sample the rest-frame UV continuum of the companions. This selection allows a uniform measurement of the UV slope, β , across the full sample. The parent catalog contains 341 LRDs; we exclude the 23 in the strongly lensed Abell 2744 field, where magnification distorts projected separations, and the redshift selection leaves 273 objects. Of these, we remove 20 (all in the GOODS-S region) for incomplete mosaic coverage, leaving a final sample of 253 LRDs. Spectroscopic redshifts are available for 27 objects, while we adopt the photometric redshifts from [D. D. Kocevski et al. \(2025\)](#) for the remainder. From this LRD sample, we also define the brightest quartile: the 65 LRDs with $F444W \leq 25.20$. At the redshifts of our sample, this band probes the rest-frame optical, $\sim 4700 - 8900 \text{ \AA}$, which is the red side of the SED.

2.2. Imaging Data and Control Samples

All measurements for companion detection and photometry are performed on the public *grizli* ([G. Brammer 2023](#)) reductions of the JWST deep fields, using every short-wavelength filter available.

We construct two control tiers. In Tier-1, we draw up to 30 control galaxies per LRD from the photometric catalog of the same field, matched in redshift within $|\Delta z|/(1+z) \leq 0.05$, and, for added safety, excluding any source within $2''$ of a detected LRD. The overlapping halves of the PRIMER mosaics duplicate 163 controls, which we remove, leaving 7427 unique galaxies in 253 redshift-matched subsets. Tier-2 is defined by matching each LRD to the subset of controls satisfying a redshift, brightness, and compactness criterion. The redshift criterion in this case is $|\Delta z| \leq 0.50$, which is more permissive than Tier-1 by design. Most redshifts entering the

match are photometric (§2.1), with uncertainties that approach this tolerance at these redshifts; a tighter cut would thus select controls primarily on noise.

The brightness criterion requires first placing controls and LRDs on the same photometric system. Total F444W magnitudes for the controls are taken from the field catalogs (total fluxes where available, aperture fluxes otherwise) and converted to AB magnitudes. To place controls and LRDs on the same photometric system, we use the LRDs themselves, which appear in both catalogs: in each field, a positional cross-match yields the median magnitude difference between the two systems ($+0.15$ to $+0.28$ mag), which we subtract from the control magnitudes. The object-to-object scatter around this median (i.e., $0.10 - 0.15$ mag) is the residual calibration uncertainty of any individual comparison. We exclude the 322 control galaxies lacking a valid flux in their field catalog from Tier-2 matching, but keep them for Tier-1. For the LRDs, we adopt the F444W values directly from the [D. D. Kocevski et al. \(2025\)](#) catalog, which defines the reference photometric system.

The brightness criterion is $|\Delta m_{444}| \leq 0.50$ mag, conservatively set at $3 - 5\times$ the residual cross-calibration scatter above, while still constraining controls to be within a factor of $\lesssim 1.6$ in flux. Statistical photometric errors at the flux levels of our sample are far smaller than the tolerance. In addition, a control is defined as compact if its catalog half-light radius does not exceed 1.15 times the 90th percentile of the LRD radii measured in the same field, so that the threshold is insensitive to the absolute size calibration of each catalog. LRDs with fewer than three matched controls are excluded from this Tier-2 analysis, to avoid small-number statistics.

Matching both in brightness and in compactness is fundamental because spurious compact “neighbors” can be artificially manufactured from the host’s light via two distinct mechanisms. The first mechanism relates to the algorithm we use: at fixed brightness, an extended light profile is far more likely to fragment into multiple segments during deblending than a point-like one. The second mechanism is astrophysical: our detection bands sample the rest-frame UV, where extended star-forming galaxies resolve into kpc-scale clumps ([A. Claeysens et al. 2023](#)); such clumps are compact and UV-blue, and thus satisfy the criteria (§2.3) designed to validate real companions. Both effects scale with the host’s brightness and extent and are minimal for point sources (§2.4). Since LRDs are mostly unresolved, controls matched in brightness alone would be systematically more contaminated by fragments of themselves.

2.3. Companion Measurement

We detect and measure companion properties with a pipeline that performs source detection directly on the short-wavelength imaging in cutouts centered on each target. We run detection on the signal-to-noise-weighted stack of the short-wavelength bands (F115W, F150W, and F200W, plus F090W where available) at a threshold of 2.5σ above the local background, with a deblending contrast of 10^{-3} so that close pairs are split wherever the data allow, following standard practice for JWST photometric catalogs (M. J. Rieke et al. 2023; A. Weibel et al. 2024; J. R. Weaver et al. 2024).

Each companion detection is characterized by its UV slope, obtained from a least-squares fit:

$$\log f_\lambda = \beta \log \lambda + \text{const}, \quad (1)$$

to aperture photometry in every available short-wavelength band for which the central wavelength falls in the rest-frame range $0.13 - 0.30 \mu\text{m}$ at the redshift of the target. We estimate background and noise locally in a $0''.75$ half-width box centered on the companion (J. R. Weaver et al. 2024).

A candidate is considered a validated blue companion if it satisfies two requirements: (i) its aperture flux exceeds $2\times$ the aperture noise in at least two of the rest-UV bands entering the fit, and (ii) its UV slope is in the range $-3.5 \leq \beta < -1$. The upper bound is our main criterion for selecting blue companions. The lower bound is a data-quality control. Stellar populations cannot be arbitrarily blue: even zero-age, dust-free, metal-free ones can only reasonably reach $\beta \approx -3.2$ (D. Schaerer 2003; R. J. Bouwens et al. 2010). Therefore, a measured $\beta < -3.5$ cannot correspond to a real spectrum and instead suggests noise-dominated photometry, which is typically caused by a faint flux fluctuation in one band tilting the fit (R. J. Bouwens et al. 2010). This constraint rejects 556 candidates with otherwise valid photometry ($\approx 11\%$ of the detections satisfying the aperture flux criterion within the $0.5 - 5$ kpc annulus); these rejections occur at statistically consistent rates around LRDs and controls.

If the candidate passes criterion (i) but has $\beta \geq -1$, its information is retained, and it is counted as a red companion; candidates below the blue bound are discarded. For every target, we record the number of valid companions, N , and the binary indicator of hosting at least one. This latter indicator is our primary statistic for this project, answering the question: does this galaxy/LRD have at least one blue companion?

We consider two nested separation windows: the full $0.5 - 5$ kpc annulus and a close-pair window at $0.5 - 2$ kpc. The complementary annulus (i.e., $2 - 5$ kpc) serves

as a larger-environment control window. The inner edge, in practice, is set by the short-wavelength angular resolution: companions can be counted only at separations $\geq 0''.15$, corresponding to ≈ 0.7 kpc at $z = 8.5$. The smallest separation recorded anywhere in the sample is, indeed, 0.713 kpc. All close-pair fractions are therefore lower bounds with respect to studies resolving smaller separations with, e.g., gravitational lensing (J. F. W. Baggen et al. 2026).

Finally, each companion is cross-matched against the field catalog and its segmentation map and classified as: (i) *substructure* when it falls within the same segment as the target (i.e., it is part of the target), (ii) *independent* when it matches a distinct catalog source, or (iii) *unmatched* when it has no catalog counterpart. Our primary statistics retain all valid companions, including the ones classified as substructures. In fact, at the small separations of interest, a genuine close pair is frequently nearly merged into a single segment (particularly around bright compact hosts). Hence, discarding a same-segment detection would preferentially remove real pairs.

2.4. Validation

Companion fractions include chance superpositions, but they affect the LRD and control arms identically, as the two samples are drawn from the same fields, at the same depths, and in matched redshift ranges. Random blank-sky positions quantify this floor on average at $f_0 \approx 0.20$ per field over the full $0.5 - 5$ kpc annulus; since both arms share it, any chance contribution biases the quoted ratios toward unity, making them conservative.

Two injection tests validate the measurements. First, the completeness test: we inject companion-like point sources around LRDs and matched controls, and recover them with the same pipeline. The fraction of injected sources that the pipeline re-detects is higher around LRDs than around controls, by $4 - 5$ percentage points at separations of $1 - 3$ kpc, with the two rates converging at 4 kpc. This asymmetry is expected, as LRDs are faint in the short-wavelength detection bands, so an injected source near an LRD is superimposed on less host light than the same source near a control. Companions are thus found slightly more easily around LRDs. Correcting for this asymmetry, at first order, multiplies the close-pair ratio by 0.95 (from $\times 2.39$ to $\times 2.27$): well within the quoted confidence interval. The completeness asymmetry, therefore, slightly amplifies the excess and cannot be its origin.

Second, the artifacts test: we inject four flux-matched synthetic twins of each LRD at empty mosaic positions




Figure 1. Close-pair statistics at 0.5 – 2 kpc. *Left:* Fraction of sources hosting at least one validated blue ($\beta < -1$) companion at 0.5 – 2 kpc, for the brightest LRD quartile ($F444W \leq 25.20$; 18/65), all LRDs (29/253), and redshift-matched control galaxies (Tier-1). Error bars are the 68% confidence interval. *Right:* Excess of the brightest LRD quartile against the no-excess line, measured against redshift-matched controls (Tier-1, $\times 2.39$) and against controls additionally matched in brightness and compactness (Tier-2, $\times 2.17$). When measured against Tier-2 controls in the 2 – 5 kpc environment, the LRDs show no excess ($\times 1.04$). The brightest-quartile excess for LRDs is detected at 3.3σ against coeval galaxies.

and re-run the entire pipeline. The twins are recovered in 966 of 1012 cases (99% in the bright quartile), and zero spurious close pairs are produced, bounding the spurious pair rate below 4.5% per bright LRD (1.2% per LRD in the full sample) at 95% confidence. In addition, as a safety measure, we visually inspect every system entering the close-pair statistics to confirm the results.

3. RESULTS

In this Section, we present the results of our study. First, we investigate whether LRDs are more likely than matched control galaxies to host a blue companion, and at which separations. Second, we study whether the companions they host differ in their properties from those of ordinary galaxies.

3.1. More Numerous Blue Companions, and Closer

We begin our investigation with the close-pair window, 0.5 – 2 kpc, whose statistics are summarized in Fig. 1. Among the brightest quartile of our LRD sample (65 objects; hereafter, the bright LRDs), 18 host at least one validated blue companion, while several host ≥ 2 . This finding leads to a blue companion pair fraction for bright LRDs of 27.7%, with a 1σ confidence interval of 21.8 – 34.3%. The Tier-1 control galaxies (i.e., matched in redshift only) pair at 11.6%, when measured with an identical pipeline. The brightest LRDs are therefore $\mathcal{R}_z = 2.39$ times more likely to have a close-pair blue companion galaxy than redshift-matched field galaxies.

This companion excess corresponds to a one-sided binomial probability of $p = 3.2 \times 10^{-4}$, or a 3.3σ excess.

The amplitude of the bright-LRD excess is also stable against the most restrictive tier of controls. In fact, against Tier-2 galaxies (i.e., matched in redshift, $F444W$ brightness, and compactness), the ratio remains similar, at $\mathcal{R}_{z,m,c} = 2.17$. However, while the amplitude is essentially unchanged, the statistical constraint is weaker because the bright, compact control-galaxy population is much smaller. Importantly, the compactness criterion is what makes the comparison with equally bright galaxies fair, as LRDs are exceptionally bright for their small size (V. Kokorev et al. 2024; D. D. Kocevski et al. 2025).

This excess is a property of the brightest LRDs only. Across the full LRD sample, the pair fraction decreases to 29/253 = 11.5%, which is indistinguishable from the control rate (see Fig. 1). Hence, the physical connection between LRDs and their blue companions seems to operate in the most luminous ones, not across the whole population.

An instrumental origin is strongly disfavored (§2.4): the synthetic-twin injections constrain the spurious pair rate to $< 4.5\%$ per bright LRD, and the compact-control pair rate ($\approx 10\%$) bounds artifacts—both far below the observed 27.7%. The measured pair fraction also agrees quantitatively with comparable measurements in the literature, which find $\approx 21\%$ lensed LRDs with a companion in the separation range overlapping ours (A. J.




Figure 2. Cumulative fraction of the brightest LRD quartile hosting a blue companion as a function of projected separation (red; shading is the 1σ confidence interval) and their Tier-1 control galaxies (gray). Both curves are empty below ≈ 0.7 kpc: the $0''.15$ resolution floor corresponds to ≈ 0.7 kpc at $z = 8.5$, comparable to the smallest separation recorded in the sample, 0.713 kpc. The companion fraction rises steeply at < 2 kpc, where 18/65 (27.7%) of the brightest quartile host a companion, a $\times 2.39$ excess over controls, significant at 3.3σ . Beyond the close-pair window, the two populations accumulate companions at comparable rates, and the cumulative ratio declines to $\times 1.79$ by 5 kpc. When measured in the 2 – 5 kpc annulus against Tier-2 controls, the ratio is $\times 1.04$: at larger separations, the surroundings of the bright LRDs are indistinguishable from those of ordinary compact galaxies of the same brightness.

Barger & L. L. Cowie 2026). This value is consistent with our 27.7%.

In Fig. 2, we show how the dependence of pair fraction on separation localizes the phenomenon. In fact, the excess is established entirely within the innermost 2 kpc. Within the close-pair window (0.5 – 2 kpc), the ratio already reaches its maximum value of $\mathcal{R}_z = 2.39$, while by 5 kpc it has declined to $\mathcal{R}_z = 1.79$, as the bright LRDs and their controls accumulate blue companions at comparable rates. The outer annulus (i.e., 2 – 5 kpc) also confirms these conclusions: here, the bright-LRD ratio, measured against Tier-2 controls, is 1.04. In other words, at larger separations, the surroundings of LRDs are indistinguishable from those of coeval compact galaxies of the same brightness.

The blue companions of the bright LRDs are thus: (i) more numerous, and (ii) systematically closer than those of control galaxies. Such a specific association, which is significant at 3.3σ against controls that share the same fields, depth, and measurement, is extremely unlikely to arise by chance.

3.2. Overwhelmingly Blue, and Bluer, Companions

Thus far, we have established that the brightest LRDs have more close blue companions. Now, we ask whether the bright LRDs and ordinary galaxies possess the same kind of companions. As we show below, the answer is a categorical no.

With no color selection, the brightest quartile of the LRD sample (65 objects) hosts 26 close companions with a measured UV slope, compared with 328 around their




Figure 3. Rest-UV slopes of all close companions with a measured β , with no color selection applied. Companions of the brightest LRDs are bluer than those of control galaxies, with median slopes $\tilde{\beta} = -2.32$ and -1.92 , respectively (dashed lines). The difference is significant at 3.4σ ($p = 6 \times 10^{-4}$).

Tier-1 controls. Of these bright-LRD companions, 92% (i.e., 24) are UV-blue, with $\beta < -1$ (see Fig. 3). In addition, the LRD companions are characterized by a median slope of $\tilde{\beta} = -2.32$, while this value increases to $\tilde{\beta} = -1.92$ for those of the controls. A shift of 0.40 in β is remarkable: a two-sample Kolmogorov–Smirnov test rejects the null hypothesis that the two color distributions are drawn from the same parent population at $p = 6 \times 10^{-4}$: a 3.4σ significance. Hence, the companions of the bright LRDs are not only overwhelmingly blue, but also very significantly bluer than a random sample of companions of ordinary galaxies available at these redshifts.

These LRD companions are young, low-extinction sources: precisely the population capable of producing a significant flux of LW photons (J. F. W. Baggen et al. 2026). Figure 4 displays a spectacular system: PRIMER-COS-7236 at $z = 5.29$. This LRD is flanked by two blue companions at projected separations of ≈ 1.8 and ≈ 1.6 kpc, both inside the close-pair window, as shown in the structural map.

It is thus remarkable that each constraint on the signal points to the same conclusion: (i) it appears only in the brightest LRDs, (ii) it is confined to separations below 2 kpc, and (iii) it involves companions that are significantly bluer than those of ordinary galaxies. Clustering alone accounts for none of these restrictions. Each of them, however, is expected if a star-forming halo is irradiating the halo immediately next to it.

4. SUMMARY AND CONCLUSIONS

This study was motivated by the discovery of abundant blue companions in the immediate vicinity of LRDs (J. F. W. Baggen et al. 2026; A. J. Barger & L. L. Cowie 2026); this finding may be key for understanding their nature. In fact, if star-forming companions occur near LRDs significantly more often than near ordinary galaxies, the UV radiation they provide may enable the direct collapse of massive black hole seeds, in the “synchronized pair” scenario (E. Visbal et al. 2014). Galaxies, however, are also clustered, and the mere presence of a companion is not conclusive. Claiming a causal con-




Figure 4. A spectacular example of an LRD with two validated close blue companions: PRIMER-COS-7236 ($z = 5.29$). *Left:* RGB composite image, with the red cross marking the LRD, the blue crosses marking the companions, and the dashed red circle marking the close-pair search radius. *Right:* Detection map of the same system, with contours at 2, 3, 5, 8, and $12\times$ the background RMS; a third validated blue companion (the structure to the lower left) lies just outside the window, at ≈ 2.3 kpc. Such close companions are of the kind required to irradiate the central source with LW photons.

nection between LRDs and blue companions requires a control.

We designed a controlled experiment, based on the photometric sample of LRDs in D. D. Kocevski et al. (2025), selecting 253 LRDs at $4.0 \leq z \leq 8.5$ across several JWST deep fields. We measured their blue companion statistics against > 7000 redshift-matched control galaxies, and against a more restrictive tier matched also in brightness and compactness. We performed every measurement on the same mosaics, with identical parameters. Our key findings are summarized here.

- The brightest quartile of the LRD sample hosts a blue companion within 2 kpc with a pair fraction of 27.7%. Identically measured, redshift-matched control galaxies pair at 11.6%. The excess is $\mathcal{R}_z = 2.39$, with a one-sided binomial probability of $p = 3.2 \times 10^{-4}$, significant at 3.3σ .
- The excess amplitude is essentially unchanged against the most restrictive control tier. Against controls matched in redshift, brightness, and compactness, the ratio is $\mathcal{R}_{z,m,c} = 2.17$, although at lower statistical significance due to the smaller number of controls.
- The phenomenon is restricted to the most luminous LRDs. The full LRD sample has a pair fraction of 11.5%, indistinguishable from the control rate.
- The phenomenon is equally restricted in separation. The pair excess declines from $\times 2.39$ within 2 kpc to $\times 1.79$ by 5 kpc. Additionally, the ratio declines to $\times 1.04$ in the 2 – 5 kpc annulus, if measured against the most restrictive control tier. In other words, beyond the close-pair window, the surroundings of bright LRDs are those of ordinary compact galaxies of the same brightness.
- The LRD companions are also remarkably different from companions of ordinary, coeval galaxies. Applying no color selection, the 26 companions of the brightest LRDs have a median UV slope of $\tilde{\beta} = -2.32$, against $\tilde{\beta} = -1.92$ for the 328 companions of the controls. A common parent distribution is rejected at $p = 6 \times 10^{-4}$, or 3.4σ . Hence, the companions of the brightest LRDs are significantly bluer than those of the controls, and overwhelmingly blue in their own right: 24 out of 26 (i.e., 92%) have $\beta < -1$. This population of sources, being young and weakly obscured, can produce a substantial LW flux.

The emerging picture is surprisingly specific, as this Letter has shown that three independent restrictions (i.e., separation, companion color, and luminosity) all point in the same direction. First, we have shown that the excess is confined entirely below 2 kpc. Since the LW flux received by a candidate DCBH at distance d scales as d^{-2} , only the closest companions can plausibly matter. Second, the companions we find are very significantly bluer (i.e., associated with young stellar populations) than those of ordinary galaxies at the same redshifts: exactly the population of sources maximizing the emission of LW dissociating radiation. Third, the excess appears at the bright end of the population, and nowhere else. This is expected if luminosity traces youth. In fact, the brightest LRDs are likely the most recently activated systems, observed while the triggering starburst is still UV-bright and inside the close-pair window. Then, as these systems age, the signature is erased, because the companion fades or merges.

This conclusion is entirely independent of *assuming* a priori the direct collapse interpretation. Two distinct measurements, i.e., the occurrence rate of close companions and their color distribution, reject the null hypothesis at $> 3\sigma$. In addition, this rejection occurs against very homogeneous control samples drawn from the same fields, observed to the same depth, and processed through the same pipeline.

Based on this evidence, we conclude that the association between luminous LRDs and their close blue companions must be physical. The two phenomena are *causally connected*: either one drives the other, or both emerge from the same synchronized formation event. This is a constraint that models for LRDs must confront. A successful theory explaining their nature can no longer describe an isolated object, but it must also answer a fundamental question: why are the most luminous of these sources found beside young, star-forming neighbors, at distances $\lesssim 2$ kpc?

Spectroscopy of the companions is the natural next step, as they are thus far identified photometrically. While their separations are far too small for chance superposition to be a significant contaminant, only spectroscopic redshifts can definitively establish physical association and measure their star formation rates, providing an estimate of the emitted LW flux.

To summarize, our study suggests that the close blue companions of bright LRDs are not incidental neighbors: the evidence collected clearly points to their being active participants in the phenomenon itself. With the luminous, compact LRDs, we may be witnessing, for the first time, the formation of massive black hole seeds caught alongside the very sources of light that made it possible.

ACKNOWLEDGMENTS

F.P. and C.M.U. thank the organizers and participants of the Sesto/Sexten Center for Astrophysics workshop on “Deciphering the Growth of Supermassive Black Holes: Accretion, Mergers, and Cosmic Evolution” for stimulating discussions and a productive research environment. F.P. acknowledges discussions with Pierluigi Rinaldi and Dale Kocevski and support from NASA through Chandra Award No. GO3-24087A and JWST Award No. JWST-GO-03805.010-A. C.M.U. acknowledges support from the National Science Foundation under grant No. AST-2407751. Some of the data products presented herein were retrieved from the Dawn JWST Archive (DJA), an initiative of the Cosmic Dawn Center, funded under grant DNR140.

Software: `astropy` (Astropy Collaboration et al. 2022), `numpy` (C. R. Harris et al. 2020), `scipy` (P. Virtanen et al. 2020), `matplotlib` (J. D. Hunter 2007), `grizli` (G. Brammer 2023)

REFERENCES

- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)
- Baggen, J. F. W., Scoggins, M. T., van Dokkum, P., et al. 2026, *ApJL*, 1002, L4, doi: [10.3847/2041-8213/ae58a5](https://doi.org/10.3847/2041-8213/ae58a5)
- Baggen, J. F. W., van Dokkum, P., Labbé, I., et al. 2023, *ApJL*, 955, L12, doi: [10.3847/2041-8213/acf5ef](https://doi.org/10.3847/2041-8213/acf5ef)
- Bagley, M. B., Pirzkal, N., Finkelstein, S. L., et al. 2024, *ApJL*, 965, L6, doi: [10.3847/2041-8213/ad2f31](https://doi.org/10.3847/2041-8213/ad2f31)
- Barger, A. J., & Cowie, L. L. 2026, arXiv e-prints, arXiv:2605.27903. <https://arxiv.org/abs/2605.27903>
- Begelman, M. C., Volonteri, M., & Rees, M. J. 2006, *MNRAS*, 370, 289, doi: [10.1111/j.1365-2966.2006.10467.x](https://doi.org/10.1111/j.1365-2966.2006.10467.x)
- Bellovary, J. 2025, *ApJL*, 984, L55, doi: [10.3847/2041-8213/adce6c](https://doi.org/10.3847/2041-8213/adce6c)
- Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2010, *ApJL*, 708, L69, doi: [10.1088/2041-8205/708/2/L69](https://doi.org/10.1088/2041-8205/708/2/L69)
- Brammer, G. 2023, *grizli*, 1.9.11 Zenodo, doi: [10.5281/zenodo.8370018](https://doi.org/10.5281/zenodo.8370018)
- Bromm, V., & Loeb, A. 2003, *ApJ*, 596, 34, doi: [10.1086/377529](https://doi.org/10.1086/377529)

- Chisholm, J., Berg, D. A., Boylan-Kolchin, M., et al. 2026, ApJL, 1004, L4, doi: [10.3847/2041-8213/ae6dae](https://doi.org/10.3847/2041-8213/ae6dae)
- Claeysens, A., Adamo, A., Richard, J., et al. 2023, MNRAS, 520, 2180, doi: [10.1093/mnras/stac3791](https://doi.org/10.1093/mnras/stac3791)
- Dijkstra, M., Haiman, Z., Mesinger, A., & Wyithe, J. S. B. 2008, MNRAS, 391, 1961, doi: [10.1111/j.1365-2966.2008.14031.x](https://doi.org/10.1111/j.1365-2966.2008.14031.x)
- Draine, B. T., & Bertoldi, F. 1996, ApJ, 468, 269, doi: [10.1086/177689](https://doi.org/10.1086/177689)
- Dunlop, J. S., Abraham, R. G., Ashby, M. L. N., et al. 2021, PRIMER: Public Release IMAGING for Extragalactic Research,, JWST Proposal. Cycle 1, ID. #1837
- Eisenstein, D. J., Willott, C., Alberts, S., et al. 2026, ApJS, 283, 6, doi: [10.3847/1538-4365/ae3163](https://doi.org/10.3847/1538-4365/ae3163)
- Fan, X., Narayanan, V. K., Lupton, R. H., et al. 2001a, AJ, 122, 2833, doi: [10.1086/324111](https://doi.org/10.1086/324111)
- Ferrara, A., Salvadori, S., Yue, B., & Schleicher, D. 2014, MNRAS, 443, 2410, doi: [10.1093/mnras/stu1280](https://doi.org/10.1093/mnras/stu1280)
- Finkelstein, S. L., Bagley, M. B., Arrabal Haro, P., et al. 2025, ApJL, 983, L4, doi: [10.3847/2041-8213/adbbd3](https://doi.org/10.3847/2041-8213/adbbd3)
- Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, ApJ, 964, 39, doi: [10.3847/1538-4357/ad1e5f](https://doi.org/10.3847/1538-4357/ad1e5f)
- Gupta, A. R., Taylor, A., Curtis-Lake, E., et al. 2026, arXiv e-prints, arXiv:2605.30414, doi: [10.48550/arXiv.2605.30414](https://doi.org/10.48550/arXiv.2605.30414)
- Habouzit, M., Volonteri, M., Latif, M., Dubois, Y., & Peirani, S. 2016, MNRAS, 463, 529, doi: [10.1093/mnras/stw1924](https://doi.org/10.1093/mnras/stw1924)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Hazlett, R., Visbal, E., Bryan, G. L., & Kulkarni, M. 2026, arXiv e-prints, arXiv:2608.13656, <https://arxiv.org/abs/2608.13656>
- Hunter, J. D. 2007, Computing in Science and Engineering, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Jones, B. L., Kocevski, D. D., Pacucci, F., et al. 2025, arXiv e-prints, arXiv:2510.07376, doi: [10.48550/arXiv.2510.07376](https://doi.org/10.48550/arXiv.2510.07376)
- Juodžbalis, I., Ji, X., Maiolino, R., et al. 2024, MNRAS, 535, 853, doi: [10.1093/mnras/stae2367](https://doi.org/10.1093/mnras/stae2367)
- Kocevski, D. D., Onoue, M., Inayoshi, K., et al. 2023, ApJL, 954, L4, doi: [10.3847/2041-8213/ace5a0](https://doi.org/10.3847/2041-8213/ace5a0)
- Kocevski, D. D., Finkelstein, S. L., Barro, G., et al. 2025, ApJ, 986, 126, doi: [10.3847/1538-4357/adbc7d](https://doi.org/10.3847/1538-4357/adbc7d)
- Kokorev, V., Caputi, K. I., Greene, J. E., et al. 2024, ApJ, 968, 38, doi: [10.3847/1538-4357/ad4265](https://doi.org/10.3847/1538-4357/ad4265)
- Labbé, I., van Dokkum, P., Nelson, E., et al. 2023, Nature, 616, 266, doi: [10.1038/s41586-023-05786-2](https://doi.org/10.1038/s41586-023-05786-2)
- Latif, M. A., Bovino, S., Van Borm, C., et al. 2014, MNRAS, 443, 1979, doi: [10.1093/mnras/stu1230](https://doi.org/10.1093/mnras/stu1230)
- Lodato, G., & Natarajan, P. 2006, MNRAS, 371, 1813, doi: [10.1111/j.1365-2966.2006.10801.x](https://doi.org/10.1111/j.1365-2966.2006.10801.x)
- Loeb, A., & Rasio, F. A. 1994, ApJ, 432, 52, doi: [10.1086/174548](https://doi.org/10.1086/174548)
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, A&A, 691, A145, doi: [10.1051/0004-6361/202347640](https://doi.org/10.1051/0004-6361/202347640)
- Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, ApJ, 963, 129, doi: [10.3847/1538-4357/ad2345](https://doi.org/10.3847/1538-4357/ad2345)
- Modelevsky, E., Nitzan, O., Sari, R., & Quataert, E. 2026, arXiv e-prints, arXiv:2608.05283, <https://arxiv.org/abs/2608.05283>
- Mortlock, D. J., Warren, S. J., Venemans, B. P., et al. 2011, Nature, 474, 616, doi: [10.1038/nature10159](https://doi.org/10.1038/nature10159)
- Naidu, R. P., Matthee, J., Katz, H., et al. 2025, arXiv e-prints, arXiv:2503.16596, doi: [10.48550/arXiv.2503.16596](https://doi.org/10.48550/arXiv.2503.16596)
- Nandal, D., & Loeb, A. 2026, ApJ, 998, 124, doi: [10.3847/1538-4357/ae32f3](https://doi.org/10.3847/1538-4357/ae32f3)
- Oh, S. P., & Haiman, Z. 2002, ApJ, 569, 558, doi: [10.1086/339393](https://doi.org/10.1086/339393)
- Omukai, K. 2001, ApJ, 546, 635, doi: [10.1086/318296](https://doi.org/10.1086/318296)
- Pacucci, F., Ferrara, A., & Kocevski, D. D. 2026, arXiv e-prints, arXiv:2601.14368, doi: [10.48550/arXiv.2601.14368](https://doi.org/10.48550/arXiv.2601.14368)
- Pacucci, F., Nguyen, B., Carniani, S., Maiolino, R., & Fan, X. 2023, ApJL, 957, L3, doi: [10.3847/2041-8213/ad0158](https://doi.org/10.3847/2041-8213/ad0158)
- Pacucci, F., Pallottini, A., Ferrara, A., & Gallerani, S. 2017, MNRAS, 468, L77, doi: [10.1093/mnrasl/slx029](https://doi.org/10.1093/mnrasl/slx029)
- Rieke, M. J., Robertson, B., Tacchella, S., et al. 2023, ApJS, 269, 16, doi: [10.3847/1538-4365/acf44d](https://doi.org/10.3847/1538-4365/acf44d)
- Rusakov, V., Watson, D., Nikopoulos, G. P., et al. 2026, Nature, 649, 574, doi: [10.1038/s41586-025-09900-4](https://doi.org/10.1038/s41586-025-09900-4)
- Schaerer, D. 2003, A&A, 397, 527, doi: [10.1051/0004-6361:20021525](https://doi.org/10.1051/0004-6361:20021525)
- Setton, D. J., Greene, J. E., de Graaff, A., et al. 2025, ApJ, 995, 118, doi: [10.3847/1538-4357/ae1500](https://doi.org/10.3847/1538-4357/ae1500)
- Taylor, A. J., Finkelstein, S. L., Kocevski, D. D., et al. 2025, ApJ, 986, 165, doi: [10.3847/1538-4357/add15b](https://doi.org/10.3847/1538-4357/add15b)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Visbal, E., Haiman, Z., & Bryan, G. L. 2014, MNRAS, 445, 1056, doi: [10.1093/mnras/stu1794](https://doi.org/10.1093/mnras/stu1794)
- Volonteri, M., & Rees, M. J. 2005, ApJ, 633, 624, doi: [10.1086/466521](https://doi.org/10.1086/466521)
- Wang, F., Yang, J., Fan, X., et al. 2021, ApJL, 907, L1, doi: [10.3847/2041-8213/abd8c6](https://doi.org/10.3847/2041-8213/abd8c6)
- Weaver, J. R., Cutler, S. E., Pan, R., et al. 2024, ApJS, 270, 7, doi: [10.3847/1538-4365/ad07e0](https://doi.org/10.3847/1538-4365/ad07e0)
- Weibel, A., Oesch, P. A., Barrufet, L., et al. 2024, MNRAS, 533, 1808, doi: [10.1093/mnras/stae1891](https://doi.org/10.1093/mnras/stae1891)