

Little red dots as obscured little blue dots: Relative abundances, luminosities, and black-hole masses

Piero Madau^{1,2} and Roberto Maiolino^{3,4,5}

¹ Dipartimento di Fisica “G. Occhialini,” Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

² Department of Astronomy & Astrophysics, University of California, 1156 High Street, Santa Cruz, CA 95064, USA

³ Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

⁴ Cavendish Laboratory, University of Cambridge, 19 JJ Thomson Avenue, Cambridge CB3 0HE, UK

⁵ Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

May 7, 2026

ABSTRACT

We test whether “little red dots” (LRDs) are the dust-reddened, high-inclination counterparts of bluer compact broad-line active galactic nuclei, here referred to as “little blue dots” (LBDs), by modeling their relative number densities and luminosities. Using the observed UV luminosity function (LF) of broad-line active galactic nuclei (BLAGNs) at $z \gtrsim 4$ as the parent distribution, we forward-model the effects of accretion rate, anisotropic emission, orientation, and dust obscuration within our super-Eddington unification framework. We show that a model with a geometrically thick accretion flow and a dusty circumnuclear cloud population reproduces the LRD LF over the luminosity range currently constrained by JWST. The predicted LRD/BLAGN fraction is strongly luminosity dependent, rising from $\simeq 3\%$ at $M_{1500} = -21$ to a peak value of $\sim 20\%$ near $M_{1500} \simeq -19$. The model also predicts a larger apparent LRD fraction at rest-frame optical wavelengths, reaching 26% at $M_{4500} = -20$ mag and 35% at $M_{6500} = -21$. The best-fitting solutions imply a characteristic per-cloud extinction $\langle A_V \rangle = 2.8^{+0.0}_{-0.4}$ mag and a mean dust covering factor $\langle C_{\text{dust}} \rangle = 0.23^{+0.27}_{-0.00}$ at 68% confidence, with the asymmetric uncertainties reflecting the degeneracy between cloud extinction and covering factor. These results may support an orientation-based unification of little dots and identify the LRD LF as a key demographic test of rapid accretion onto infant black holes at cosmic dawn. Within this same framework, UV-selected LRDs are predicted to host systematically more massive black holes than unobscured LBDs, not because they represent a distinct parent population, but because dust attenuation preferentially removes lower-mass obscured systems from the observed UV sample.

Key words. Accretion (14); Active galactic nuclei (16); James Webb Space Telescope (2291); Supermassive black holes (1663)

1. Introduction

The James Webb Space Telescope has uncovered a large population of compact BLAGNs at $z \gtrsim 4$, powered by accretion onto massive black holes with inferred masses of $\sim 10^6$ – $10^8 M_\odot$ (e.g., Harikane et al. 2023; Maiolino et al. 2024; Taylor et al. 2025; Juodžbalis et al. 2026). Within this emerging class, a subset of sources exhibits red, V-shaped UV-optical continua and is commonly classified as LRDs, while the remainder show bluer continua and are often referred to as LBDs (Brazzini et al. 2026). Recent analyses suggest that these categories may not represent two sharply distinct populations, but rather the opposite ends of a continuous sequence in compactness, continuum shape, and broad-line prominence (Billand et al. 2026). In current JWST spectroscopic samples, LRDs appear to constitute only $\sim 10\%$ – 30% of the compact BLAGN population once selection effects are taken into account – in particular the fact that LRDs have often been preferentially targeted because of their unusual colors, whereas LBDs, having colors similar to star-forming galaxies, are typically identified only serendipitously in spectroscopic surveys (Hainline et al. 2025; Taylor et al. 2025).

In Paper I (Madau & Maiolino 2026), we proposed that LRDs and LBDs are different viewing-angle manifestations of the same compact, rapidly accreting BLAGN population. In that framework, the central engine is a geometrically thick, radiation-supported accretion flow that emits anisotropically and is surrounded by an equatorially concentrated BLR and a dusty circumnuclear component of modest covering factor. LBDs correspond to lower-inclination, relatively unobscured sightlines, whereas LRDs are the same sources viewed through dust-intercepted lines of sight. Because both the intrinsic anisotropic continuum and the line-of-sight attenuation depend on viewing angle, the observed LRD/LBD ratio inferred from spectroscopic surveys is not a purely geometric quantity. It reflects instead the convolution of orientation-dependent luminosity, photometric selection, and dust obscuration.

The physical nature of LRDs remains debated. Purely stellar interpretations invoke compact, evolved stellar populations with strong nebular emission to explain the V-shaped SEDs (e.g., Baggen et al. 2024; Pérez-González et al. 2024; Leung et al. 2025; Labbe et al. 2025; Hainline et al. 2025). However, several lines of evidence increasingly favour

an AGN-dominated continuum in at least a substantial fraction of LRDs: the prevalence of broad Balmer lines and extreme Balmer EWs (Hviding et al. 2025; Yan et al. 2025), the frequent dominance of an unresolved rest-optical component (Killi et al. 2024; Hviding et al. 2025), and spectroscopic features around the Balmer break that are more naturally explained by absorption and radiative transfer in high-density gas than by an evolved stellar population (Juodžbalis et al. 2024; Ji et al. 2025; Inayoshi & Maiolino 2025; D’Eugenio et al. 2025b; Yan et al. 2025).

A key uncertainty concerns the origin of the rest-UV continuum used to construct the LRD and BLAGN luminosity functions. It has been argued that, even if the rest-optical emission is AGN-related, the UV spectrum may still be dominated by stellar light (e.g., de Graaff et al. 2025; Baggen et al. 2026). At the same time, detailed analyses of individual LRDs and related compact broad-line systems have found that the UV continuum is often consistent with being primarily AGN-powered, albeit with some host-galaxy contribution in certain cases (e.g., Labbe et al. 2024; D’Eugenio et al. 2025a; Ji et al. 2026; Tang et al. 2026; Pacucci et al. 2026; Torralba et al. 2026). In what follows, we adopt this AGN-dominated UV hypothesis as the working assumption of our LF modeling.

In this paper we test that unification picture at the population level. We ask whether the observed LRD LF can be recovered from the compact BLAGN parent population through orientation-dependent emission and dust interception alone. Using the observed BLAGN UV LF as input, we forward-model the effects of accretion rate, viewing angle, anisotropy, and extinction on the emergent rest-UV continuum. The central question is whether the same super-Eddington geometry that links LBDs and LRDs at the level of individual SEDs can also reproduce the observed abundance and luminosity dependence of LRDs with plausible dust covering factors and extinction distributions. Failure to do so would imply that orientation and obscuration alone are insufficient, and that additional physical distinctions between the two populations are required.

2. Hybrid forward modeling of the LRD LF

The logic of our model is straightforward. In the rapid accretion regime, the funnel geometry produces strongly anisotropic emission: polar sightlines receive a brighter, harder continuum, while equatorial views are self-shadowed and depleted in XUV/soft-X-ray photons (Wang et al. 2014; Lupi et al. 2024; Madau 2026). Hard X-rays may be further weakened if the corona is efficiently Compton-cooled by the intense soft-photon field (Madau & Haardt 2024; Trinca et al. 2026). A dusty component of modest covering factor can then redden high-inclination sightlines, producing the rest-optical slopes and V-shaped continua of LRDs, while lower-inclination views appear as LBDs. Thus, the observed LRD/LBD differences arise mainly from orientation and line-of-sight processing, rather than from distinct engines or evolutionary phases (e.g., Kido et al. 2025; Naidu et al. 2025; Pacucci et al. 2026).

We adopt the clumpy line-of-sight interception formalism of Paper I; details are given in the Appendix. As shown in Figure 1, the observed LRD LF lies below the parent BLAGN UV LF and falls more steeply at the bright end. In our framework, this offset reflects orientation-dependent dust interception: LRDs are the minority of com-

pact BLAGNs viewed through an effective dusty obscurer, while less obscured sightlines appear as LBDs. The declining LRD fraction toward bright UV magnitudes suggests a luminosity-dependent dust covering factor, as expected if radiation pressure reduces circumnuclear dust in more luminous systems (e.g., Ricci et al. 2017, 2022).

Our implementation is a hybrid forward Monte Carlo model anchored to the empirical double-power-law UV LF of compact BLAGNs obtained by jointly fitting the spectroscopic samples of Taylor et al. (2025) ($3.5 < z < 6$) and Juodžbalis et al. (2026) ($4 < z < 7$). We draw mock parent sources over the range $-23 \leq M_{1500} \leq -16.75$ mag using this best-fitting BLAGN LF, while the fit to the LRD LF itself is restricted to bins brighter than $M_{1500} = -17$ mag.

Each mock source is also assigned an accretion rate drawn from a Gaussian distribution in $\log \dot{m}$,¹ centered at $\langle \log \dot{m} \rangle = 1.3$, corresponding to the midpoint of the super-Eddington model grid explored here. We keep this mean fixed and treat the dispersion $\sigma_{\log \dot{m}}$ as a free parameter, in analogy with the width of the dust-attenuation distribution (see below). This choice anchors the luminosity inversion to the representative accretion-rate scale of our models while allowing for source-to-source variation within the rapidly accreting parent population.

For each mock source we then assign an isotropically distributed viewing angle and evaluate the corresponding inclination-dependent anisotropic SED predicted by our tabulated super-Eddington models (Madau 2026). The probability that the line of sight intersects circumnuclear dusty clouds is described by a luminosity-dependent covering factor,

$$C_{\text{dust}}(M_{1500}) = \frac{C_0}{1 + \exp[-k(M_{1500} - M_{\text{break}})]}, \quad (1)$$

where C_0 is the faint-end plateau, M_{break} is the characteristic transition magnitude, and k controls the sharpness of the decline toward the bright end. Since the present bright-end LRD counts only weakly constrain the transition width, we fix $k = 4$ and treat (C_0, M_{break}) as free parameters.

Given C_{dust} , we determine probabilistically whether the source intercepts the dusty obscurer using the same clumpy Poisson formalism introduced in Paper I. The dusty cloud population is assumed to have angular width $\sigma_d = 0.4$, corresponding to a geometrically thick equatorial distribution whose normalization is fixed by equation (1). We adopt a gray AGN-type extinction curve, qualitatively similar to the flat attenuation laws inferred for reddened active nuclei (e.g., Gaskell et al. 2004), corresponding in our implementation to an effective far-UV attenuation of $A_{1500}/A_V \simeq 1.36$. The extinction per obscuring cloud is drawn from a truncated normal distribution at $A_V^c > 0$ with fixed mean $\langle A_V^c \rangle = 2.8$ mag and free dispersion $\sigma_{A_V^c}$. We also include an additional uniform host-galaxy screen of $A_V^{\text{host}} = 0.7$ mag applied to all sightlines. The total transmission along dust-intercepted lines of sight is then computed from the conditional Poisson cloud-crossing formalism described in the Appendix.

At this stage, each mock source has a specified target observed magnitude M_{1500} , accretion rate $\dot{m}$, viewing angle, and total line-of-sight attenuation. We then determine the black hole mass iteratively by requiring that the final attenuated SED reproduce exactly the assigned observed UV

¹ Here, we define $\dot{m} \equiv 0.1 \dot{M} c^2 / L_{\text{Edd}}$.




Fig. 1. Observed and model UV LFs of LRDs and their parent BLAGN population. The upper panel shows the fixed double-power-law parent BLAGN UV LF obtained by jointly fitting the spectroscopic BLAGN samples of Taylor et al. (2025) ($3.5 < z < 6$) and Juodžbalis et al. (2026) ($4 < z < 7$), together with the individual BLAGN LF measurements. It also shows the photometric LRD UV LFs from Kocevski et al. (2025) and Kokorev et al. (2024) over $4.5 < z < 6.5$. The green dashed curve shows the best-fitting LRD LF obtained from our hybrid forward model, in which mock BLAGNs drawn from the parent LF are processed through angle-dependent super-Eddington SEDs, luminosity-dependent dust covering, stochastic cloud obscuration, and photometric LRD selection. The lower panel shows the resulting apparent ratio $\Phi_{\text{LRD}}/\Phi_{\text{BLAGN}}$ as a function of observed M_{1500} , together with the 95% confidence band obtained from the family of $\Delta\chi^2$ -accepted forward-model solutions around the best fit. The best-fitting solution has $\chi^2/\text{dof} = 1.3$ for 10 data points and 4 free parameters, with $\sigma_{\log \dot{m}} = 0.1$, $\sigma_{A_V^c} = 0.06$, $C_0 = 0.24$, and $M_{\text{break}} = -20.23$; the remaining parameters are fixed to $\langle \log \dot{m} \rangle = 1.3$, $\langle A_V^c \rangle = 2.8$ mag, $\sigma_d = 0.4$, $A_V^{\text{host}} = 0.7$ mag, and $k = 4$. The implied population-averaged effective dust covering factor is modest, $\langle C_{\text{dust}} \rangle = 0.22$, while $\simeq 15\%$ of all compact BLAGNs satisfy the adopted $(\beta_{\text{opt}}, \beta_{\text{UV}})$ photometric LRD selection.

magnitude at 1500 Å. In this way, the empirical BLAGN UV LF is preserved by construction, while dust-reddened sightlines are naturally mapped onto intrinsically more luminous and typically more massive systems than unobscured ones at the same observed M_{1500} .

From the final attenuated continuum we measure the rest-frame UV slope β_{UV} using windows centered at 1500, 2000, and 2500 Å, and the optical slope β_{opt} using windows at 4500, 6000, and 8000 Å. A mock source is classified as a photometric LRD if

$$\beta_{\text{opt}} > 0, \quad -2.8 < \beta_{\text{UV}} < -0.37, \quad (2)$$

following Kocevski et al. (2025). The predicted LRD LF is then obtained from the fraction of mock BLAGNs in each magnitude bin that satisfy these criteria, multiplied by the fixed parent BLAGN LF.

The free parameters of the fiducial model are therefore $\sigma_{\log \dot{m}}$, $\sigma_{A_V^c}$, and the two sigmoid parameters (C_0, M_{break}) .

The mean accretion rate, $\langle \log \dot{m} \rangle = 1.3$, the mean per-cloud extinction, $\langle A_V^c \rangle = 2.8$ mag, the dust angular width, $\sigma_d = 0.4$, the host reddening screen, $A_V^{\text{host}} = 0.7$ mag, and the sigmoid sharpness, $k = 4$, are held fixed. These parameters are constrained by minimizing the χ^2 difference between the predicted and observed LRD LFs from Kocevski et al. (2025) and Kokorev et al. (2024). In this way, the model tests whether the observed abundance and luminosity dependence of LRDs can be reproduced as the dust-intercepted tail of the compact BLAGN population within a single super-Eddington framework.

3. Results and discussion

Before presenting the LF comparison, we briefly summarize the logic of the population synthesis:

1. We adopt the observed rest-frame UV LF of compact BLAGNs at $z \gtrsim 4$ as the parent distribution of the little-dot population. In this framework, the Taylor et al. and Juodžbalis et al. BLAGN LF provides the underlying UV number density from which both unobscured LBDs and dust-reddened LRDs are generated through orientation and circumnuclear obscuration.
2. Each Monte Carlo source is assigned a target observed UV magnitude M_{1500} drawn from this empirical BLAGN LF, together with an accretion rate $\dot{m}$ drawn from a lognormal distribution with fixed center $\langle \log \dot{m} \rangle = 1.3$ and free dispersion $\sigma_{\log \dot{m}}$.
3. We then assign each source a random isotropic viewing angle and evaluate the corresponding inclination-dependent anisotropic SED predicted by the accretion-flow model. The viewing angle changes both the observed continuum normalization and the UV-to-optical spectral shape, as expected for a geometrically thick super-Eddington emitter.
4. We probabilistically determine whether the line of sight intercepts circumnuclear dusty clouds. The interception probability is described by a luminosity-dependent covering factor $C_{\text{dust}}(M_{1500})$, modeled as a sigmoid that decreases toward the bright end. This luminosity dependence is required because LRDs are rare among the most UV-luminous compact BLAGNs; a luminosity-independent covering factor would overproduce bright LRDs.
5. For dust-intercepted sightlines, the inclination-dependent SED is attenuated with a gray extinction law. The per-cloud extinction is drawn from a distribution centered on $\langle A_V^c \rangle = 2.8$ mag, and a modest foreground host-galaxy screen is applied to all sightlines. Dust-free sightlines therefore retain only the foreground attenuation, whereas obscured sightlines acquire substantially larger UV suppression.
6. Given the total line-of-sight attenuation and the chosen inclination, we determine the black-hole mass iteratively by requiring that the final attenuated SED reproduce exactly the target observed UV magnitude M_{1500} . In this way, the mock catalog preserves by construction the observed BLAGN UV LF while allowing dust-reddened sightlines to correspond to intrinsically more luminous and typically more massive systems than unobscured ones at the same observed M_{1500} .
7. From each final mock SED we measure the synthetic UV and optical continuum slopes and classify as LRDs only

those objects satisfying the adopted photometric color-selection criteria. Counting these selected mock LRDs in each magnitude bin yields the predicted UV LRD LF.

8. The same mock SEDs are also used to compute rest-frame optical magnitudes at 4500 Å and 6500 Å. Because dust-reddened, high-inclination sources have much redder optical-to-UV colors than unobscured BLAGNs, objects that are only moderately bright in observed M_{1500} can populate substantially brighter optical-magnitude bins. This color migration increases the apparent LRD/BLAGN ratio in the optical relative to the UV.

3.1. A low covering solution

Figure 1 compares the resulting forward-model predictions with the observed UV LF measurements. The parent BLAGN LF is described by a fixed double power law (blue solid curve), while the predicted LRD LF (green dashed curve) is obtained by applying orientation-dependent dust interception and the LRD color selection to this parent population. The model reproduces the observed normalization and overall shape of the LRD LF over $-21 \lesssim M_{1500} \lesssim -17$. As shown in the lower panel, it also predicts a strongly luminosity-dependent apparent LRD fraction, $\Phi_{\text{LRD}}/\Phi_{\text{BLAGN}}$, rising from below 1% at $M_{1500} \sim -22$ to a maximum of $\sim 20\%$ near $M_{1500} \sim -19$, with the subsequent faint-end behavior depending on the extrapolated BLAGN LF and the magnitude range sampled in the Monte Carlo draw. This behavior reflects the combined effects of a dust covering factor that declines toward high luminosities and a color-selection efficiency that depends on luminosity, inclination, and attenuation. The same mock catalog predicts a substantially larger apparent LRD/BLAGN ratio when the population is selected in the rest-frame optical rather than in the UV; this additional wavelength dependence is discussed below.

The preferred solution corresponds to an effective dust covering factor of $\langle C_{\text{dust}} \rangle \simeq 0.22$, implying that about one fifth of all BLAGN sightlines are dust-intercepted over the modeled luminosity range. In the mock population, $\simeq 15\%$ of BLAGNs satisfy the photometric LRD criterion, while an additional $\simeq 7\%$ are dust-intercepted but fall outside the adopted LRD selection box. Thus about two-thirds of the obscured sightlines are classified as LRDs, indicating that the adopted gray dust prescription is effective at producing LRD-like continua for a substantial fraction of dust-intercepted sources. The fit favors both a relatively narrow accretion-rate distribution, with $\sigma_{\log \dot{m}} \simeq 0.10$, and a narrow per-cloud extinction distribution, with $\sigma_{A_V^c} \simeq 0.06$. This indicates that the observed LRD LF and colors are reproduced most naturally by a fairly homogeneous rapidly accreting parent population, together with dust-intercepted sightlines that cluster around a characteristic per-cloud extinction.

3.2. Parameter degeneracies in the dust-obscuration fit

To illustrate the non-uniqueness of the dust-obscuration solution, we relax the assumption of a fixed $\langle A_V^c \rangle = 2.8$ mag and perform a 3D grid scan over $(\langle A_V^c \rangle, C_0, \sigma_{A_V^c})$, projecting the minimum χ^2 over $\sigma_{A_V^c}$ at each grid point. The fiducial best-fit parameters quoted above therefore correspond to

one particular slice through this broader parameter space. Allowing $\langle A_V^c \rangle$ to vary together with C_0 and $\sigma_{A_V^c}$ shifts the global minimum only modestly, confirming that the fixed- $\langle A_V^c \rangle$ solution already lies close to the broader optimum.

Figure 2 shows the resulting projected $\Delta\chi^2$ confidence contours in the $(\langle C_{\text{dust}} \rangle, \langle A_V^c \rangle)$ plane. Rather than a sharply localized minimum, the allowed region forms an elongated degeneracy valley extending from the fiducial solution toward larger covering factors and slightly lower characteristic cloud extinctions. This means that the observed LRD number counts can be reproduced by a family of dust-obscuration models with different combinations of effective obscured fraction and per-cloud attenuation.

The physical origin of this degeneracy is straightforward. If individual dusty clouds are highly opaque in the far UV, only a relatively modest fraction of BLAGNs needs to be intercepted in order to shift enough sources into the photometric LRD selection region. Conversely, if the characteristic extinction per cloud is weaker, simply increasing the number of obscured sightlines is not sufficient by itself: many such sightlines would remain only mildly reddened and would fail the Kocevski LRD criteria. The fit therefore compensates in two linked ways. First, it increases $\langle C_{\text{dust}} \rangle$, so that a larger fraction of BLAGNs is processed by dust. Second, it favors a broader cloud-extinction distribution, so that even when the mean $\langle A_V^c \rangle$ is lower, a non-negligible high- A_V^c tail remains available to produce the red optical slopes and V-shaped continua required for LRD selection.




Fig. 2. Projected $\Delta\chi^2 \equiv \chi^2 - \chi_{\text{min}}^2$ confidence contours in the $(\langle C_{\text{dust}} \rangle, \langle A_V^c \rangle)$ plane, obtained from a 3D grid scan over $(\langle A_V^c \rangle, C_0, \sigma_{A_V^c})$ with M_{break} held fixed at its fiducial best-fit value, and projected by minimizing over $\sigma_{A_V^c}$ at each fixed $(\langle A_V^c \rangle, C_0)$. Here $\langle C_{\text{dust}} \rangle$ is the number-density-weighted mean dust covering factor. Solid, dashed, and dotted contours denote the 68.3%, 95.4%, and 99.7% confidence regions, respectively. The elongated confidence valley illustrates the trade-off between cloud extinction and obscured fraction: lower per-cloud extinction requires a larger dust covering factor to reproduce the observed LRD counts. The red star marks the global minimum of the 3D scan projected onto this plane.

Along this valley, the preferred width of the cloud-extinction distribution also changes systematically. At the high- $\langle A_V^c \rangle$ end, the best solutions correspond to nearly

Table 1. Predicted apparent LRD/BLAGN ratios at representative UV and optical magnitudes.

Magnitude	1500 Å	4500 Å	6500 Å
−21	0.03	0.26	0.35
−20	0.15	0.26	0.30
−19	0.20	0.22	0.17

single-valued per-cloud extinction. Toward lower $\langle A_V^c \rangle$, however, the preferred dispersion rises to $\sigma_{A_V^c} \sim 0.3\text{--}0.4$. In this regime, a broader cloud-extinction distribution is needed to populate a non-negligible high- A_V^c tail, ensuring that at least some sightlines remain sufficiently reddened to satisfy the LRD color criteria even though the characteristic cloud is less opaque.

These degenerate solutions are not observationally identical. While they produce similar LRD number counts within the adopted $(\beta_{\text{opt}}, \beta_{\text{UV}})$ selection box, they predict different abundances of mildly reddened compact BLAGNs lying just outside that box. Relaxing the current photometric selection criteria, or explicitly targeting sources with intermediate UV-to-optical slopes, therefore offers a possible observational route to breaking this degeneracy.

3.3. Predicted wavelength dependence of the apparent LRD fraction

Table 1 compares the best-fitting apparent ratio $\Phi_{\text{LRD}}/\Phi_{\text{BLAGN}}$ when the compact BLAGN population is binned by observed rest-frame magnitude at 1500, 4500, and 6500 Å. The UV-selected ratio rises from 3% at $M_{1500} = -21$ to 20% at $M_{1500} = -19$, consistent with the obscured minority fraction inferred from the fiducial dust covering model. In contrast, the corresponding optical-selected ratios are systematically larger: the 4500 Å ratio reaches 26% at $M_{4500} = -20$, while the 6500 Å ratio is larger still, 35% at $M_{6500} = -20$.

This behavior does not imply that a larger fraction of lines of sight intrinsically produce LRDs at longer wavelength. In the model, the true fraction of dust-intercepted sightlines remains fixed by the adopted circumnuclear cloud covering factor and is of order $\sim 20\%$. The larger optical ratios arise because the quantity plotted is not the intrinsic sightline fraction, but the ratio of observed LFs evaluated at fixed observed magnitude in each band. For each mock source, the parent compact BLAGN population is sampled from the observed UV LF, and the emergent SED is then computed for its assigned inclination and dust interception. Dust attenuation shifts the same object by an amount

$$M_{\lambda,\text{obs}} = M_{\lambda,\text{int}} + A_{\lambda},$$

with a wavelength dependence that is highly non-uniform. At fixed observed M_{1500} , a dust-intercepted source must be intrinsically more luminous than an unobscured source in order to compensate for the UV attenuation. Although the adopted extinction curve is gray compared with a standard Milky-Way law, the attenuation still decreases from 1500 to 4500 and 6500 Å. The same reddened objects therefore appear relatively brighter in the rest-frame optical than in the UV, shifting them into brighter optical-magnitude bins and increasing the apparent LRD/BLAGN ratio at optical wavelengths.

The consequence is a color-driven migration across magnitude bins. An intrinsically luminous dusty source that appears extremely faint in the UV can still populate a moderate optical luminosity bin, whereas the unobscured BLAGN population remains concentrated closer to its intrinsic UV luminosity ranking. Rest-frame optical bins are therefore populated efficiently by red, high-inclination objects that would be strongly underrepresented in a UV-selected comparison. The apparent ratio $\Phi_{\text{LRD}}/\Phi_{\text{BLAGN}}$ at fixed optical magnitude can thus substantially exceed the underlying geometric obscured fraction, even though the number of LRD-producing lines of sight is unchanged.

This effect becomes progressively stronger toward longer wavelength because A_{λ} decreases monotonically from the UV to the optical. The same dust-reddened subset that is heavily suppressed at 1500 Å is only moderately shifted at 4500 and less shifted still at 6500 Å, causing the observed LRD counts to be increasingly recovered relative to the blue BLAGN denominator. The systematic rise of the ratio from UV to optical wavelengths is therefore a direct consequence of wavelength-dependent extinction combined with magnitude binning in the observed frame.

The larger ratios predicted at 4500 Å and 6500 Å provide a direct observational test of the orientation-based unification scenario. If LRDs are indeed dust-intercepted members of the same compact BLAGN population, rest-frame optical surveys should recover a substantially larger apparent LRD fraction than UV-selected samples over comparable observed luminosity ranges.

3.4. Black-hole mass demographics

Figure 3 summarizes the black-hole mass demographics implied by the best-fitting forward model. The top panel compares the normalized inferred mass distributions of unobscured sightlines (LBD analogs), all dust-intercepted sources, and the subset of dust-intercepted sources that additionally satisfy the photometric LRD color criteria.

The inferred black-hole mass distributions of unobscured and dust-reddened mock sources differ markedly. Non-dusty sightlines are concentrated toward the low-mass end of the allowed SED grid, since only modest intrinsic UV luminosities are required to reproduce the assigned observed M_{1500} in the absence of strong attenuation. By contrast, dust-intercepted sightlines are shifted to substantially larger inferred masses, peaking near $\log(M_{\text{BH}}/M_{\odot}) \sim 7.5$ and extending to beyond $10^{8.5} M_{\odot}$, while the subset that additionally satisfies the photometric LRD color criteria is displaced even further toward the high-mass tail.

This pronounced mass segregation should not be interpreted as evidence that LRDs and LBDs arise from intrinsically distinct parent populations. In the model, both classes are drawn from the same compact BLAGN UV LF and differ only by orientation and line-of-sight obscuration. The separation in inferred M_{BH} is instead a direct selection effect produced by the iterative luminosity inversion. Once the target observed M_{1500} is held fixed, a dust-reddened source must correspond to a substantially brighter intrinsic UV continuum than an unobscured source in order to remain detectable within the same observed UV magnitude bin. This forces the obscured solutions toward larger black-hole masses. Conversely, dust-intercepted BLAGNs with intrinsically lower masses would be attenuated below the UV




Fig. 3. Black-hole mass demographics implied by the best-fitting forward model. *Top panel:* normalized inferred mass distributions of unobscured (LBD-like) sightlines, all dust-intercepted sources, and the subset classified as photometric LRDs. *Bottom panel:* black-hole mass function implied by the same mock BLAGN population after the iterative inversion from the observed compact BLAGN UV LF. Red points show the BLAGN BHMf inferred by Taylor et al. (2025) from single-epoch virial mass estimates based on broad H α emission.

luminosity range sampled by the present data and would therefore not enter the observed LRD population.

The lower panel shows the black-hole mass function (BHMf) implied by the same best-fitting mock compact BLAGN parent population. Over the range $10^6 \lesssim M_{\text{BH}}/M_{\odot} \lesssim 10^9$, the inferred BHMf declines monotonically, with a progressively steeper fall toward the highest masses. This behavior reflects the bright-end cutoff of the parent compact BLAGN UV LF, together with the adopted mapping between UV luminosity, accretion rate, and black-hole mass. It should therefore not be interpreted as a direct measurement of the intrinsic cosmic AGN BHMf, but rather as the mass distribution of actively accreting compact BLAGNs selected through the observed UV LF under the best-fitting accretion-rate distribution.

For comparison, Figure 3 also shows the BLAGN BHMf inferred by Taylor et al. (2025) using standard virial mass estimates. Our recovered BHMf is systematically flatter and remains elevated toward the high-mass end relative to the Taylor et al. determination. In the present framework, this offset arises naturally because dust-reddened compact BLAGNs that remain observable after attenuation must correspond to intrinsically brighter UV sources. In the iterative luminosity inversion, these sources are therefore assigned larger black-hole masses. The recovered BHMf thus represents the full compact BLAGN parent population, including UV-suppressed obscured systems, rather than only the unobscured subset whose masses can be estimated directly from continuum-based virial methods.

Because the inversion from UV luminosity to black-hole mass remains partially degenerate with the adopted form of $p(\log \dot{m})$, the inferred absolute mass scale is conditional on this choice: broader or differently parameterized accretion-rate distributions would shift and broaden the recovered BHMf. The high-mass tail is especially sensitive to these assumptions, to the dust-obscuration correction, and to the limited dynamic range of the present BLAGN LF, and should therefore be regarded as indicative rather than definitive. We finally reiterate that the inferred BHMf is for actively accreting compact BLAGNs that are individually detectable in current spectroscopic samples. Fainter accreting black holes, such as those inferred from stacked analyses (e.g., Geris et al. 2026), would preferentially add low-mass systems and therefore steepen the faint-/low-mass end of the BHMf.

4. Summary

We have presented a forward model in which LRDs are the dust-reddened, high-inclination subset of the same compact BLAGN population that appears as LBDs along cleaner sightlines. Starting from the observed compact BLAGN UV LF, the model combines a super-Eddington, anisotropic continuum, an inclination-dependent obscuring cloud population, and the same photometric color selection used in current LRD samples.

The model reproduces the observed LRD LF using a modest luminosity-dependent dust covering factor, whose average value is of order $\langle C_{\text{dust}} \rangle \simeq 0.2$, together with characteristic cloud extinctions of a few magnitudes. These parameters are not independently fixed: the fit shows an elongated degeneracy in which lower cloud extinctions can be compensated by larger dust covering factors, while larger extinctions require a smaller intercepted fraction. The predicted apparent LRD/BLAGN fraction is strongly luminosity dependent, rising from a few percent at the bright end to $\sim 20\%$ near $M_{1500} \simeq -19$.

A key prediction is that the apparent LRD/BLAGN ratio should be larger when the population is binned by rest-frame optical magnitude than by UV magnitude. This does not require a larger intrinsic obscured fraction at longer wavelengths, but follows from color-driven migration across observed magnitude bins: dust-reddened sources are strongly displaced in the UV while remaining comparatively bright in the optical.

Finally, the same best-fitting mock population predicts a strong segregation in inferred black-hole mass between unobscured and dust-reddened sightlines. This segregation

is a consequence of the luminosity inversion rather than evidence for distinct parent populations. The resulting BHMF should therefore be interpreted as a model-dependent mass distribution for actively accreting compact BLAGNs, conditional on the assumed accretion-rate distribution and dust correction.

References

- Baggen, J. F. W., Scoggins, M. T., van Dokkum, P., et al. 2026, arXiv e-prints, arXiv:2602.02702
- Baggen, J. F. W., van Dokkum, P., Brammer, G., et al. 2024, ApJ, 977, L13
- Billand, J.-B., Elbaz, D., Franco, M., et al. 2026, arXiv e-prints, arXiv:2604.11677
- Brazzini, M., D'Eugenio, F., Maiolino, R., et al. 2026, arXiv e-prints, arXiv:2601.22214
- de Graaff, A., Hviding, R. E., Naidu, R. P., et al. 2025, arXiv e-prints, arXiv:2511.21820
- D'Eugenio, F., Maiolino, R., Perna, M., et al. 2025a, arXiv e-prints, arXiv:2503.11752
- D'Eugenio, F., Nelson, E., Ji, X., et al. 2025b, arXiv e-prints, arXiv:2510.00101
- Elitzur, M. 2008, New A Rev., 52, 274
- Elitzur, M. & Shlosman, I. 2006, ApJ, 648, L101
- Gaskell, C. M., Goosmann, R. W., Antonucci, R. R. J., & Whyson, D. H. 2004, ApJ, 616, 147
- Geris, S., Maiolino, R., Isobe, Y., et al. 2026, MNRAS, 545, staf1979
- Hainline, K. N., Maiolino, R., Juodžbalis, I., et al. 2025, ApJ, 979, 138
- Harikane, Y., Zhang, Y., Nakajima, K., et al. 2023, ApJ, 959, 39
- Hviding, R. E., de Graaff, A., Miller, T. B., et al. 2025, A&A, 702, A57
- Inayoshi, K. & Maiolino, R. 2025, ApJ, 980, L27
- Ji, X., Maiolino, R., Übler, H., et al. 2025, MNRAS, 544, 3900
- Ji, X., Pezzulli, G., D'Eugenio, F., et al. 2026, arXiv e-prints, arXiv:2604.03370
- Juodžbalis, I., Ji, X., Maiolino, R., et al. 2024, MNRAS, 535, 853
- Juodžbalis, I., Maiolino, R., Baker, W. M., et al. 2026, MNRAS, 546, stag086
- Kido, D., Ioka, K., Hotokezaka, K., Inayoshi, K., & Irwin, C. M. 2025, MNRAS, 544, 3407
- Killi, M., Watson, D., Brammer, G., et al. 2024, A&A, 691, A52
- Kocevski, D. D., Finkelstein, S. L., Barro, G., et al. 2025, ApJ, 986, 126
- Kokorev, V., Caputi, K. I., Greene, J. E., et al. 2024, ApJ, 968, 38
- Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, ApJ, 978, 92
- Labbe, I., Greene, J. E., Matthee, J., et al. 2024, arXiv e-prints, arXiv:2412.04557
- Leung, G. C. K., Finkelstein, S. L., Pérez-González, P. G., et al. 2025, ApJ, 992, 26
- Lupi, A., Trinca, A., Volonteri, M., Dotti, M., & Mazzucchelli, C. 2024, A&A, 689, A128
- Madau, P. 2026, A&A, 708, A116
- Madau, P. & Haardt, F. 2024, ApJ, 976, L24
- Madau, P. & Maiolino, R. 2026, arXiv e-prints, arXiv:2602.22386
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, A&A, 691, A145
- Naidu, R. P., Matthee, J., Katz, H., et al. 2025, arXiv e-prints, arXiv:2503.16596
- Nenkova, M., Sirocky, M. M., Ivezić, Ž., & Elitzur, M. 2008, ApJ, 685, 147
- Pacucci, F., Ferrara, A., & Kocevski, D. D. 2026, arXiv e-prints, arXiv:2601.14368
- Pérez-González, P. G., Barro, G., Rieke, G. H., et al. 2024, ApJ, 968, 4
- Ricci, C., Ananna, T. T., Temple, M. J., et al. 2022, ApJ, 938, 67
- Ricci, C., Trakhtenbrot, B., Koss, M. J., et al. 2017, Nature, 549, 488
- Tang, M., Stark, D. P., Mason, C. A., et al. 2026, arXiv e-prints, arXiv:2604.03563
- Taylor, A. J., Finkelstein, S. L., Kocevski, D. D., et al. 2025, ApJ, 986, 165
- Torralba, A., Matthee, J., Pezzulli, G., et al. 2026, A&A, 707, A75
- Trinca, A., Lupi, A., Haardt, F., & Madau, P. 2026, arXiv e-prints, arXiv:2602.22305
- Wada, K. 2012, ApJ, 758, 66
- Wang, J.-M., Qiu, J., Du, P., & Ho, L. C. 2014, ApJ, 797, 65
- Yan, Z., Inayoshi, K., Chen, K., & Guo, J. 2025, arXiv e-prints, arXiv:2512.11050

Appendix A: Clumpy dust obscuration model

Following Paper I, we describe the dusty obscurer as a clumpy, equatorially concentrated cloud distribution by specifying the mean number of clouds intersected along a line of sight at inclination i as

$$N_d(i) = N_{0,d} \exp \left[-\frac{\cos^2 i}{2\sigma_d^2} \right], \quad (\text{A.1})$$

(see, e.g., the formalism of [Nenkova et al. 2008](#)), where σ_d controls the angular thickness of the cloud distribution about the equatorial plane. The corresponding escape probability for direct disk photons is

$$P_{\text{esc},d}(i) = \exp \left[-N_d(i) \right]. \quad (\text{A.2})$$

The normalization $N_{0,d}$ is fixed by requiring that the solid-angle-averaged probability of intercepting at least one dusty cloud equals the global dust covering factor,

$$C_{\text{dust}} = \int_0^{\pi/2} [1 - P_{\text{esc},d}(i)] \sin i \, di. \quad (\text{A.3})$$

Here, C_{dust} is the solid-angle covering factor of the effective obscurer, i.e. the cloud population whose interception reddens the continuum sufficiently to produce a photometric LRD. The corresponding conditional inclination distribution of dust-intercepting sightlines is

$$p(i | d) = \frac{[1 - P_{\text{esc},d}(i)] \sin i}{C_{\text{dust}}}. \quad (\text{A.4})$$

Physically, a broader angular distribution for the dusty component than for the BLR is expected because the obscurer lies at larger radii, beyond the dust sublimation front, where radiation pressure on dust and/or disk winds can produce a larger scale height than in the dust-free BLR (e.g., [Elitzur & Shlosman 2006](#); [Nenkova et al. 2008](#); [Elitzur 2008](#); [Wada 2012](#)).