A&A, 715, A15 (2026)
https://doi.org/10.1051/0004-6361/202660815
© The Authors 2026

Astronomy & Astrophysics

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

Received 6 May 2026 / Accepted 23 September 2026

ABSTRACT

We tested whether the so-called 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 ultraviolet (UV) luminosity function (LF) of broad-line active galactic nuclei (BLAGNs) at $z \gtrsim 4$ as the parent distribution, we forward-modeled the effects of the 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 the James Webb Space Telescope. The predicted LRD-to-BLAGN ratio is strongly luminosity dependent, rising from $\simeq$3% at $M_{1500} = -21$ mag 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 27% at $M_{4500} = -20$ mag and 39% at $M_{6500} = -21$. The fiducial solution adopted a baseline host extinction $A_V^{\rm host} = 0.5$ mag, a characteristic per-cloud nuclear extinction $\langle A_V^c \rangle = 3.0$ mag, and a gray $R_V = 5.5$ attenuation curve, and it yields a mean dust covering factor $\langle C_{\rm dust} \rangle \simeq 0.23$. These results can 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. In our model, the UV continuum is dominated by the active nucleus; LRDs and LBDs are drawn from the same observed BLAGN UV LF, but attenuation makes UV-selected LRDs intrinsically brighter and their inferred black holes more massive at fixed observed M_{1500}.

Key words. galaxies: active – galaxies: high-redshift – galaxies: nuclei – quasars: general – quasars: supermassive black holes

1. Introduction

The James Webb Space Telescope (JWST) has uncovered a large population of broad-line active galactic nuclei (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 ultraviolet (UV)-optical continua and is commonly classified as little red dots (LRDs), while the remainder show bluer continua and are often referred to as little blue dots (LBDs) (Brazzini et al. 2026). Recent analyses suggested that these categories might 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 BLAGN population when selection effects are taken into account, in particular because LRDs have often been preferentially targeted because of their unusual colors, whereas LBDs, whose colors are similar to those of 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 rapidly accreting BLAGN population. In this framework, the central engine is a geometrically thick radiation-supported accretion flow that emits anisotropically and is surrounded by an equatorially concentrated broad-line region (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 the intrinsic anisotropic continuum and the line-of-sight attenuation depend on viewing angle, the observed LRD-to-LBD ratio inferred from spectroscopic surveys is not a purely geometric quantity. It instead reflects the convolution of the 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 spectral energy distributions (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 favor a continuum dominated by an active galactic nucleus (AGN) in at least a substantial fraction of LRDs: the prevalence of broad Balmer lines and extreme Balmer equivalent widths (EWs) (Hviding et al. 2025; Yan et al. 2026), 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

⋆ Corresponding author: pmadau@ucsc.edu

Open Access article, published by EDP Sciences, under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.
This article is published in open access under the Subscribe to Open model. Subscribe to A&A to support open access publication.

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. 2025; Yan et al. 2026).

A key uncertainty concerns the origin of the rest-UV continuum that is used to construct the LRD luminosity function (LF) and the BLAGN LF. It has been argued that even if the rest-optical emission is AGN related, the UV spectrum might still be dominated by stellar light (e.g., de Graaff et al. 2026; 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. 2026; 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 this unification picture at the population level. We ask whether the observed LRD LF can be recovered from the 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 the 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 extreme-UV and soft-X-ray photons (Wang et al. 2014; Lupi et al. 2024; Madau 2026). Hard X-rays may be further weakened when the corona is efficiently Compton-cooled by the intense soft-photon field (Madau & Haardt 2024; Trinca et al. 2026). A dusty component with a 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 differences between LRDs and LBDs arise mainly from orientation and line-of-sight processing and not from distinct engines or evolutionary phases (e.g., Kido et al. 2025; Naidu et al. 2026; Pacucci et al. 2026).

We adopted 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. There is tentative evidence, at approximately 2.5σ, that the LRD-to-BLAGN ratio in the brightest bin, $M_{1500} = -21$, is lower than at $M_{1500} = -19$; the decrease in the $M_{1500} = -20$ bin is not statistically significant. In our framework, this offset reflects orientation-dependent dust interception: LRDs are the minority of BLAGNs viewed through an effective dusty obscurer, while less obscured sightlines appear as LBDs. This possible bright-end decrease motivated us to explore a luminosity-dependent dust covering factor, as may be expected if radiation pressure reduces circumnuclear dust in more luminous systems (e.g., Ricci et al. 2017, 2022).

Our implementation was a hybrid forward Monte Carlo model anchored to the empirical double power-law UV LF of 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 drew mock parent sources over the range $-23 \le M_{1500} \le -16.75$ mag using this best-fitting BLAGN LF, while the fit to the LRD LF itself was restricted to bins brighter than $M_{1500} = -17$ mag.

We parameterized the parent LF as

$$\Phi_{\rm BLAGN}(M) = \frac{\Phi_\star}{10^{0.4(\alpha+1)(M-M_\star)} + 10^{0.4(\beta+1)(M-M_\star)}}. \quad (1)$$

The weighted joint fit gives $\log_{10}(\Phi_\star/{\rm Mpc}^{-3}\,{\rm mag}^{-1}) = -4.72 \pm 0.21$, $M_\star = -21.75 \pm 0.22$, and $\alpha = -1.63 \pm 0.16$. The bright-end slope reaches the adopted lower bound and was fixed at $\beta = -8$; its precise value has a negligible effect over the luminosity range that dominates the forward model. The Taylor et al. point at $M_{1500} = -17$ is displayed but was excluded from this double power-law fit, since including its abrupt downturn would require an additional faint-end turnover. The shaded parent-LF intervals in Figure 1 propagate the covariance of the three fitted parameters.

Each mock source was 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 kept this mean fixed and treated the dispersion $\sigma_{\log \dot{m}}$ as a free parameter, in analogy to the width of the dust-attenuation distribution (see below). This choice anchored 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 assigned an isotropically distributed viewing angle and evaluated 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,

$$C_{\rm dust}(M_{1500}) = \frac{C_0}{1 + \exp\left[-k(M_{1500} - M_{\rm break})\right]}, \quad (2)$$

where C_0 is the faint-end plateau, $M_{\rm break}$ 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 fixed $k = 4$ and treated $(C_0, M_{\rm break})$ as free parameters.

It is important to distinguish this global covering factor from the inclination-dependent interception probability. At a given observed magnitude, $C_{\rm dust}(M_{1500})$ is the probability of dust interception averaged over an isotropic distribution of sightlines; it is not a second function evaluated at the inclination of an individual source. Equation (2) fixes the equatorial normalization $N_{0,d}(M_{1500})$ by requiring that the solid-angle average in Equation (A.3) equal $C_{\rm dust}(M_{1500})$. For a mock source whose inclination i has already been drawn, the probability of intercepting at least one cloud is then

$$1 - P_{{\rm esc},d}(i \mid M_{1500}) = 1 - \exp\{-N_{0,d}(M_{1500}) \times \exp\left[-\frac{\cos^2 i}{2\sigma_d^2}\right]\}. \quad (3)$$

Thus, the luminosity sets the population-averaged covering factor, while the inclination determines how that probability is

[1] We defined $\dot{m} \equiv 0.1\dot{M}c^2/L_{\rm Edd}$.

distributed over sightlines. The two dependences are complementary rather than degenerate definitions of the same quantity: at fixed M_{1500}, equatorial views have the highest interception probability, whereas the normalization averaged over all i is exactly $C_{\rm dust}(M_{1500})$.

Given $C_{\rm dust}$, we determined probabilistically whether the source intercepted the dusty obscurer using the same clumpy Poisson formalism as introduced in Paper I. The dusty cloud population was assumed to have an angular width $\sigma_d = 0.4$, corresponding to a geometrically thick equatorial distribution whose normalization is fixed by Equation (2). We adopted a gray AGN-type extinction curve, qualitatively similar to the flat attenuation laws inferred for reddened active nuclei (e.g., Gaskell et al. 2004). We fixed $R_V = 5.5$ and suppressed the 2175-Å feature ($B = 0$), giving $A_{1500}/A_V \simeq 1.30$. The extinction per obscuring cloud was drawn from a truncated normal distribution at $A_V^c > 0$ with fixed mean $\langle A_V^c \rangle = 3.0$ mag and free dispersion $\sigma_{A_V^c}$. We also included a uniform host-galaxy screen of $A_V^{\rm host} = 0.5$ mag applied to all sightlines. Thus, an unobscured LBD-like sightline received only the baseline host attenuation, whereas a one-cloud dust-intercepted sightline had a characteristic total extinction of $A_V \simeq 3.5$ mag. The total transmission along dust-intercepted sightlines was computed from the conditional Poisson cloud-crossing formalism described in the appendix; because more than one cloud might be intersected, their effective attenuation is not represented by a single fixed total A_V.

At fixed values of the remaining model parameters, a sensitivity scan favors $R_V \simeq 5.5$; steeper curves with $R_V = 4.6$–5.1 tend to overproduce the color-selected LRDs. Re-optimizing the LF parameters partially compensates for this effect through a smaller fitted covering factor, while still grayer curves with $R_V \gtrsim 6$ provide statistically comparable fits only by requiring substantially larger covering factors. The LRD LF therefore does not constrain R_V independently of the dust geometry. We retained $R_V = 5.5$ as a representative gray AGN attenuation curve and treated the dependence on R_V as a systematic uncertainty rather than introducing it as an additional free parameter.

At this stage, each mock source had a specified target observed magnitude M_{1500}, accretion rate $\dot{m}$, viewing angle, and total line-of-sight attenuation. We then determined the black hole mass iteratively by requiring that the final attenuated SED reproduce exactly the assigned observed UV magnitude at 1500 Å. In this way, the empirical BLAGN UV LF was preserved by construction, while dust-reddened sightlines were 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 measured the rest-frame UV slope $\beta_{\rm UV}$ using windows centered at 1500, 2000, and 2500 Å, and the optical slope $\beta_{\rm opt}$ using windows at 4500, 6000, and 8000 Å. A mock source was classified as a photometric LRD if

$$\beta_{\rm opt} > 0, \qquad -2.8 < \beta_{\rm UV} < -0.37, \tag{4}$$

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

Figure 2 makes the orientation dependence explicit. The left panel shows the synthetic continuum slopes of the mock catalog, color-coded by inclination. The LRD selection preferentially isolates high-inclination objects: for the best-fitting realization,

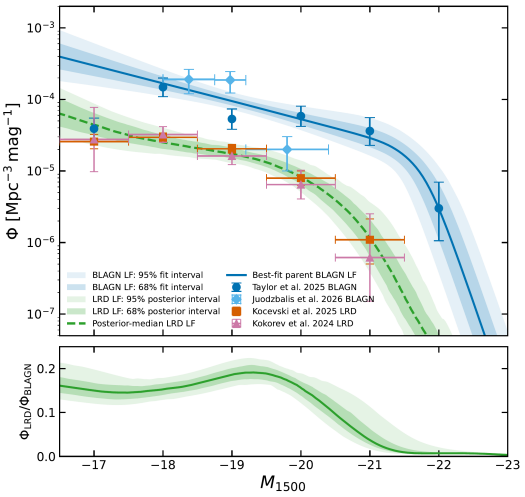


Fig. 1. Observed and modeled UV LFs of LRDs and their parent BLAGN population. The upper panel shows the fitted double power-law parent BLAGN LF, the measured LRD LFs, and the posterior-median model prediction. Dark and light bands enclose the 68% and 95% intervals. The lower panel shows $\Phi_{\rm LRD}/\Phi_{\rm BLAGN}$. The best fit has a reduced $\chi^2 = 1.01$ for 10 data points and four free parameters; the fitted and fixed parameters are listed in Table 1. The implied mean covering factor is $\langle C_{\rm dust} \rangle = 0.23$, and 15% of the modeled BLAGNs satisfy the photometric LRD selection.

the median inclination is 59.8° for the full isotropic mock, 73.6° for dust-intercepted sightlines, and 78.9° for photometrically selected LRDs. The right panel shows that $C_{\rm dust}(M_{1500})$ controls the solid-angle normalization, while the interception probability at every luminosity rises strongly toward edge-on sightlines.

The free parameters of the fiducial model are therefore $\sigma_{\log \dot{m}}$, $\sigma_{A_V^c}$, and the two sigmoid parameters $(C_0, M_{\rm break})$. The mean accretion rate, $\langle \log \dot{m} \rangle = 1.3$, the mean per-cloud extinction, $\langle A_V^c \rangle = 3.0$ mag, the dust angular width, $\sigma_d = 0.4$, the host reddening screen, $A_V^{\rm host} = 0.5$ mag, the attenuation-curve parameters, $R_V = 5.5$ and $B = 0$, and the sigmoid sharpness, $k = 4$, were held fixed. We sampled the four-parameter posterior using bounded uniform priors over the ranges listed in Table 1. The resulting marginalized constraints are also listed in the table and shown in Figure 3; the two dispersion parameters remain weakly constrained toward their lower prior bounds. All final model predictions and confidence bands were evaluated using 40 000 mock sources.

The posterior constraints are conditional on the fixed characteristic cloud extinction and attenuation curve. To quantify this systematic dependence, we performed a conditional profile-likelihood scan in the $(\langle A_V^c \rangle, C_0)$ plane, minimizing over $\sigma_{A_V^c}$ at each grid point while holding $\sigma_{\log \dot{m}}$ and $M_{\rm break}$ at their maximum-likelihood values. Figure 4 shows a pronounced degeneracy: a lower characteristic extinction per cloud can be compensated for by a larger dust covering factor, which both increases the fraction of intercepted sightlines and, in the adopted Poisson clumpy model, increases the mean number of clouds conditional on interception. The fiducial four-parameter maximum-likelihood solution lies within the 68% confidence region. Thus, the LRD LF does not determine the extinction and covering factor independently, and the narrow posterior interval on C_0 should be interpreted as conditional on the adopted dust prescription.

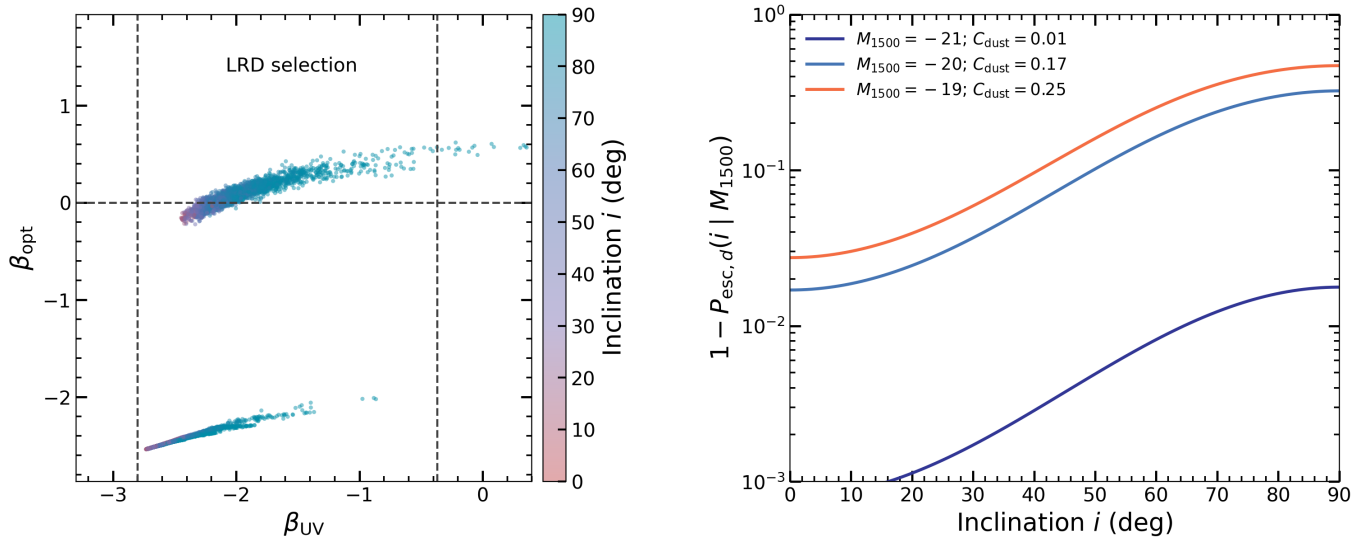


Fig. 2. Orientation dependence of the best-fitting forward model. Left: Synthetic rest-frame continuum slopes, color-coded by inclination; dashed lines delimit the photometric LRD selection. Right: Probability of intercepting at least one dusty cloud versus inclination at three observed UV magnitudes. Each listed $C_{\rm dust}$ is the solid-angle average of the corresponding curve, weighted by $\sin i$.

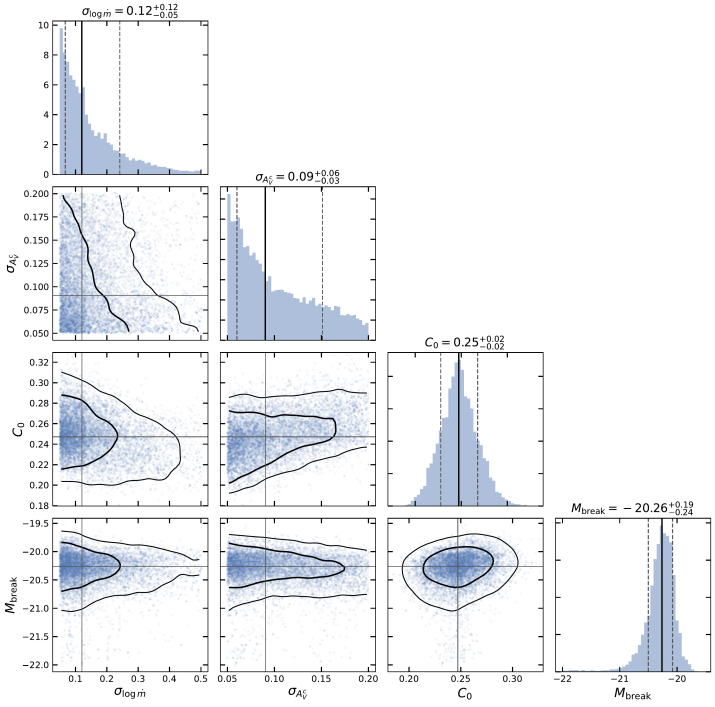


Fig. 3. Posterior distributions for the four free parameters. The lines mark the medians and 16th–84th percentile intervals. The contours enclose the 68% and 95% probability. The best fit has $\chi^2 = 6.08$ for six degrees of freedom (reduced $\chi^2 = 1.01$). The two dispersion parameters remain weakly constrained toward their lower prior limits.

3. Results and discussion

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

1. We adopted the observed rest-frame UV LF of 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 unobscured LBDs and dust-reddened LRDs are generated through orientation and circumnuclear obscuration.
2. Each Monte Carlo source was 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 assigned each source a random isotropic viewing angle and evaluated the corresponding inclination-dependent anisotropic SED predicted by the accretion-flow model. The viewing angle changes the observed continuum normalization and the UV-to-optical spectral shape, as expected for a geometrically thick super-Eddington emitter.
4. We probabilistically determined whether the line of sight intercepted circumnuclear dusty clouds. The interception probability is described by a luminosity-dependent covering factor $C_{\rm dust}(M_{1500})$, modeled as a sigmoid that decreases toward the bright end. This luminosity dependence is required because LRDs are rare among the most UV-luminous BLAGNs; a luminosity-independent covering factor would overproduce bright LRDs.

Table 1. Model parameters.

Parameter	Status	Prior range	Result
$\sigma_{\log \dot{m}}$	Free	[0.05, 0.5]	$0.12^{+0.12}_{-0.05}$
$\sigma_{A_V^c}$	Free	[0.05, 0.2]	$0.09^{+0.06}_{-0.03}$
C_0	Free	[0.05, 0.95]	$0.25^{+0.02}_{-0.02}$
$M_{\rm break}$	Free	[−22, −16]	$-20.26^{+0.19}_{-0.24}$
$\langle \log \dot{m} \rangle$	Fixed	–	1.30
$A_V^{\rm host}$	Fixed	–	0.5 mag
$\langle A_V^c \rangle$	Fixed	–	3.0 mag
R_V	Fixed	–	5.5
B	Fixed	–	0
σ_d	Fixed	–	0.40
k	Fixed	–	4.0

Notes. The uncertainties give the 16th–84th posterior percentiles.

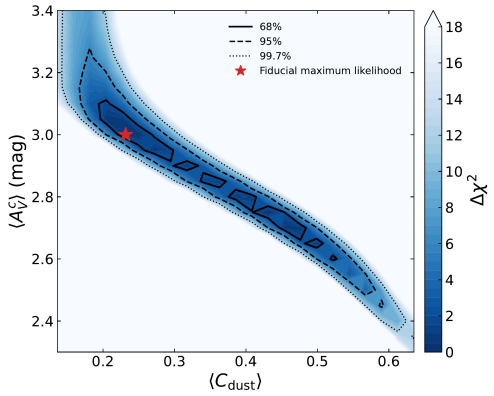


Fig. 4. Conditional profile likelihood in the $(\langle C_{\rm dust} \rangle, \langle A_V^c \rangle)$ plane. Solid, dashed, and dotted contours enclose the 68%, 95%, and 99.7% confidence regions. The red star marks the fiducial maximum-likelihood solution: $C_0 = 0.255$, $\langle C_{\rm dust} \rangle = 0.232$, and $\langle A_V^c \rangle = 3.0$ mag. The remaining profiling assumptions are described in the text.

5. For dust-intercepted sightlines, the inclination-dependent SED was attenuated with a gray extinction law. The per-cloud extinction was drawn from a distribution centered on $\langle A_V^c \rangle = 3.0$ mag, and a foreground host-galaxy screen with $A_V^{\rm host} = 0.5$ mag was applied to all sightlines. The dust-free sightlines therefore retained only the foreground attenuation, whereas obscured sightlines acquired substantially larger UV suppression.
6. Given the total line-of-sight attenuation and the chosen inclination, we determined the black-hole mass iteratively by requiring that the final attenuated SED reproduced exactly the target observed UV magnitude M_{1500}. In this way, the mock catalog preserved the observed BLAGN UV LF by construction 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 measured the synthetic UV and optical continuum slopes and classified as LRDs only those objects satisfying the adopted photometric color-selection criteria. Counting these selected mock LRDs in each magnitude bin yielded the predicted UV LRD LF.
8. The same mock SEDs were 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-to-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 was described by a fixed double power law (solid blue curve), while the predicted LRD LF (dashed green curve) was 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, $\Phi_{\rm LRD}/\Phi_{\rm BLAGN}$, rising from below 1% at $M_{1500} \sim -22$ to a maximum of ∼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 higher apparent LRD-to-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 fiducial solution corresponds to a number-density-weighted mean dust-covering factor of $\langle C_{\rm dust} \rangle \simeq 0.23$, implying that about one quarter of BLAGN sightlines are dust-intercepted over the modeled luminosity range. In the mock population, ≃15% of BLAGNs satisfy the photometric LRD criterion, while an additional ≃8% 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 maximum-likelihood solution favors narrow distributions, with $\sigma_{\log \dot{m}} = 0.06$ and $\sigma_{A_V^c} = 0.05$; after marginalizing over the covering-factor parameters, the posterior medians are 0.12 and 0.09, respectively, with both dispersions remaining weakly constrained toward their lower prior bounds.

3.2. Predicted wavelength dependence of the apparent LRD fraction

Table 2 compares the best-fitting apparent ratio $\Phi_{\rm LRD}/\Phi_{\rm BLAGN}$ when the modeled BLAGN population was 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 higher: the 4500 Å ratio reaches 27% at $M_{4500} = -20$, while the 6500 Å ratio is higher still, 37% 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

Table 2. Predicted apparent LRD-to-BLAGN ratios at representative UV and optical magnitudes.

Magnitude	1500 Å	4500 Å	6500 Å
−21	0.03	0.17	0.39
−20	0.15	0.27	0.37
−19	0.20	0.29	0.33

model, the true fraction of dust-intercepted sightlines remained fixed by the adopted circumnuclear cloud covering factor and was ∼20%. The higher 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 BLAGN population was sampled from the observed UV LF, and the emergent SED was then computed for its assigned inclination and dust interception. The dust attenuation shifts the same object by an amount of

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

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-to-BLAGN ratio at optical wavelengths.

The consequence is a color-driven migration in all 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 $\Phi_{\mathrm{LRD}}/\Phi_{\mathrm{BLAGN}}$ 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 still less shifted at 6500 Å, causing the observed LRD counts to be increasingly recovered relative to the blue BLAGN denominator. The systematic rise in 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 higher 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 BLAGN population, rest-frame optical surveys should recover a substantially larger apparent LRD fraction than UV-selected samples over comparable observed luminosity ranges.

3.3. Black hole mass demographics

Figure 5 summarizes the black hole mass demographics implied by the best-fitting forward model. The top panel compares the

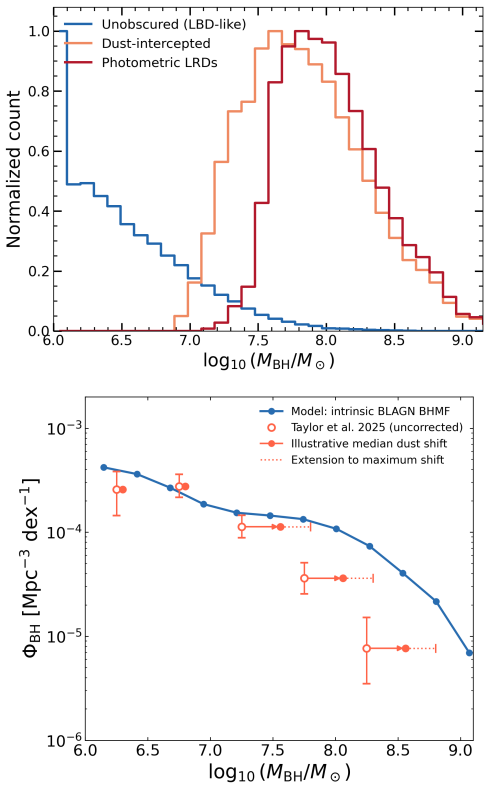


Fig. 5. Black hole mass demographics implied by the best-fitting model. Top: Normalized mass distributions of unobscured sightlines, all dust-intercepted sources, and photometric LRDs. Bottom: Intrinsic BLAGN black hole mass function (BHMF) and the uncorrected measurements of Taylor et al. (2025) (open red points). Filled points apply their median mass corrections; dotted extensions show the largest correction. Number densities are unchanged, so these shifts are not a recomputed extinction-corrected BHMF.

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 because 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 higher inferred masses, peaking near $\log(M_{\mathrm{BH}}/M_\odot) \sim 7.5$ 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 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. When 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 higher black-hole masses. Conversely, dust-intercepted BLAGNs with intrinsically lower masses would be attenuated below the UV 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 BLAGN parent population. Over the range $10^6 \lesssim M_{\rm BH}/M_\odot \lesssim 10^9$, the inferred BHMF declines monotonically, with a progressively steeper fall toward the highest masses. This behavior reflects the bright-end cutoff of the parent 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 BLAGNs selected through the observed UV LF under the best-fitting accretion-rate distribution.

For comparison, Figure 5 also shows the BLAGN BHMF inferred by Taylor et al. (2025) using uncorrected broad-Hα virial masses. The authors estimated median extinction corrections of +0.05 dex below $10^7\,M_\odot$ and +0.31 dex above it, with a maximum correction of +0.55 dex for their most reddened source. We illustrate these corrections as horizontal shifts at fixed number density; because the individual mass posteriors and completeness weights are not available, the shifted points should not be interpreted as a recomputed extinction-corrected BHMF. The comparison therefore shows only the direction and approximate magnitude of the mass correction: extinction shifts the Taylor et al. estimates toward higher masses, while the corresponding change in number density cannot be inferred from the published binned measurements.

To characterize the high-mass tail directly, we examined the mock sources with $\log(M_{\rm BH}/M_\odot) \geq 8$. Of these, 94% were viewed through the dusty cloud distribution and 90% satisfied the adopted LRD selection. Their median observed and intrinsic UV magnitudes were $M_{1500} = -18.8$ and -23.5, respectively; for the LRD subset, the corresponding values were -18.7 and -23.6. The high-mass excess is therefore associated primarily with objects that are present at moderate observed UV luminosities but are shifted to much brighter intrinsic luminosities, and hence to higher inferred masses, by the dust correction.

Because the inversion from UV luminosity to black hole mass remains partially degenerate with the adopted form of $p(\log \dot{m})$, we repeated the fit at fixed $\langle \log \dot{m} \rangle = 0.67$, 0.85, 1.09, 1.30, and 1.50, holding $\sigma_{\log \dot{m}} = 0.06$ fixed and reoptimizing $\langle A_V^c \rangle$, $\sigma_{A_V^c}$, C_0, and $M_{\rm break}$. The refitted mean cloud extinction remained in the narrow range $\langle A_V^c \rangle = 2.8$–$3.0$ mag, while the number-density-weighted mean covering factor spanned only $\langle C_{\rm dust} \rangle = 0.25$–$0.29$. The solutions through $\langle \log \dot{m} \rangle = 1.30$ retained acceptable fits to the LRD LF, with reduced $\chi^2 = 0.9$–1.3, while the 1.50 case was less satisfactory (reduced $\chi^2 = 1.9$). As expected, decreasing the mean accretion rate shifted the inferred mass distribution and BHMF toward higher masses, whereas the 1.50 model shifted them downward. None of these conditional fits, however, removed the high-mass offset relative to the Taylor et al. determination. The high-mass tail remained sensitive to the assumed accretion-rate distribution, the dust correction, and 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 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 and low-mass ends of the BHMF.

The preceding results assume that the rest-frame UV continuum is AGN dominated. Under this assumption, the same nuclear component must reproduce the blue UV slope and the red optical–near-infrared (near-IR) continuum of LRDs. Strong obscuration by a standard extinction curve would suppress and redden the AGN far-UV emission too strongly; an AGN-dominated interpretation therefore requires a comparatively gray attenuation law, as adopted in our fiducial model. This assumption also ties the observed M_{1500} directly to the attenuated AGN luminosity and hence underlies the luminosity inversion used to infer the black-hole mass.

If the observed UV continuum is instead dominated by an unobscured stellar component, the blue UV slope no longer constrains the extinction law toward the nucleus, and the AGN can be obscured by a more conventional, non-gray dust screen. A nuclear extinction of a few magnitudes can then strongly suppress the AGN in the far-UV, leaving the unobscured stellar component to produce the observed blue UV spectrum, while the obscured AGN re-emerges at optical and near-IR wavelengths and powers the broad Hα emission. Our modeling of LRDs showed that after we added an unobscured stellar component that dominated at 1500 Å, the optical–near-IR SED still required $A_V^c \simeq 3$ mag toward the AGN (Madau & Maiolino, in prep.).

In this star-dominated limit, the observed M_{1500} primarily traces the host and not the attenuated nuclear luminosity. Nuclear obscuration therefore produces less migration between UV-magnitude bins, and the mapping from the parent BLAGN LF to the LRD LF depends on the joint stellar–AGN luminosity distribution, the AGN-to-stellar contrast, and the dust-covering factor. This does not imply that the preference for massive LRDs necessarily disappears: the intrinsic AGN luminosity and black hole mass must instead be constrained from the extinction-corrected optical continuum and broad Hα. This distinction is also relevant to the comparison with Taylor et al. (2025), who retained uncorrected broad-Hα luminosities in their fiducial virial masses. They estimated that an internal-extinction correction would increase the inferred masses by a median of 0.19 dex for the full sample and 0.31 dex above $10^7\,M_\odot$, reaching 0.55 dex for the most reddened source. Such corrections would move their BHMF toward the high-mass tail predicted here. A self-consistent treatment of the stellar-dominated case requires a two-component population model and is deferred to future work.

4. Summary

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

The model reproduces the observed LRD LF using a modest luminosity-dependent dust covering factor, whose average value is about $\langle C_{\rm dust} \rangle \simeq 0.2$, 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 for by larger dust covering factors, while higher extinctions require a smaller intercepted fraction. The predicted apparent LRD-to-BLAGN ratio is strongly luminosity dependent, rising from a few percent at the bright end to ∼20% near $M_{1500} \simeq -19$.

A key prediction is that the apparent LRD-to-BLAGN ratio should be higher when the population is binned by rest-frame optical magnitude instead of by UV magnitude. This does not require a larger intrinsic obscured fraction at longer wavelengths,

but follows from color-driven migration in the 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 and is not evidence of distinct parent populations. The resulting BHMF should therefore be interpreted as a model-dependent mass distribution for actively accreting BLAGNs, conditional on the assumed accretion-rate distribution and dust correction.

References

Baggen, J. F. W., van Dokkum, P., Brammer, G., et al. 2024, ApJ, 977, L13
Baggen, J. F. W., Scoggins, M. T., van Dokkum, P., et al. 2026, ApJ, 1002, L4
Billand, J. B., Elbaz, D., Franco, M., et al. 2026, A&A, 714, A327
Brazzini, M., D'Eugenio, F., Maiolino, R., et al. 2026, A&A, 713, A332
de Graaff, A., Hviding, R. E., Naidu, R. P., et al. 2026, MNRAS, 551, stag1567
D'Eugenio, F., Nelson, E., Ji, X., et al. 2025, ApJ, submitted [arXiv:2510.00101]
D'Eugenio, F., Maiolino, R., Perna, M., et al. 2026, MNRAS, 547, stag401
Elitzur, M. 2008, New Astron. Rev., 52, 274
Elitzur, M., & Shlosman, I. 2006, ApJ, 648, L101
Gaskell, C. M., Goosmann, R. W., Antonucci, R. R. J., & Whysong, 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
Labbe, I., Greene, J. E., Matthee, J., et al. 2024, ApJ, submitted [arXiv:2412.04557]
Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, ApJ, 978, 92
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, A&A, 714, A276
Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, A&A, 691, A145
Naidu, R. P., Matthee, J., Katz, H., et al. 2026, Nature, 656, 329
Nenkova, M., Sirocky, M. M., Ivezić, Ž., & Elitzur, M. 2008, ApJ, 685, 147
Pacucci, F., Ferrara, A., & Kocevski, D. D. 2026, ApJ, submitted [arXiv:2601.14368]
Pérez-González, P. G., Barro, G., Rieke, G. H., et al. 2024, ApJ, 968, 4
Ricci, C., Trakhtenbrot, B., Koss, M. J., et al. 2017, Nature, 549, 488
Ricci, C., Ananna, T. T., Temple, M. J., et al. 2022, ApJ, 938, 67
Tang, M., Stark, D. P., Mason, C. A., et al. 2026, ApJ, submitted [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, A&A, 710, A289
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. 2026, ApJ, 1002, 159

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

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

(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

$$P_{\mathrm{esc},d}(i) = \exp[-N_d(i)]. \tag{A.2}$$

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,

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

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

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

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).