



Central Massive Black Holes Are Not Ubiquitous in Local Low-mass Galaxies

Fan Zou¹, Elena Gallo¹, Anil C. Seth², Edmund Hodges-Kluck³, David Ohlson², Tommaso Treu⁴,
Vivienne F. Baldassare⁵, W. N. Brandt^{6,7,8}, Jenny E. Greene⁹, Piero Madau^{10,11}, Dieu D. Nguyen¹,
Richard M. Plotkin^{12,13}, Amy E. Reines¹⁴, Alberto Sesana^{11,15,16}, Jong-Hak Woo¹⁷, and Jianfeng Wu¹⁸

¹Department of Astronomy, University of Michigan, 1085 South University, Ann Arbor, MI 48109, USA; fanzou01@gmail.com

²Department of Physics and Astronomy, University of Utah, USA

³X-ray Astrophysics Laboratory, NASA/GSFC, Greenbelt, MD 20771, USA

⁴Physics and Astronomy Department, University of California, Los Angeles, CA 90095, USA

⁵Department of Physics and Astronomy, Washington State University, Pullman, WA 99164, USA

⁶Department of Astronomy and Astrophysics, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA

⁷Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

⁸Department of Physics, 104 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA

⁹Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

¹⁰Department of Astronomy & Astrophysics, University of California, 1156 High Street, Santa Cruz, CA 95064, USA

¹¹Dipartimento di Fisica “G. Occhialini,” Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

¹²Physics Department, University of Nevada, Reno, 1664 North Virginia Street, Reno, NV 89557, USA

¹³Nevada Center for Astrophysics, University of Nevada, Las Vegas, NV 89154, USA

¹⁴eXtreme Gravity Institute, Department of Physics, Montana State University, Bozeman, MT 59717, USA

¹⁵INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

¹⁶INAF—Osservatorio Astronomico di Brera, Via Brera 20, I-20121 Milano, Italy

¹⁷Astronomy Program, Department of Physics and Astronomy, Seoul National University, Seoul 08826, Republic of Korea

¹⁸Department of Astronomy, Xiamen University, Xiamen, Fujian 361005, People’s Republic of China

Received 2025 June 13; revised 2025 September 8; accepted 2025 September 11; published 2025 October 14

Abstract

The black hole occupation fraction (f_{occ}) defines the fraction of galaxies that harbor central massive black holes (MBHs), irrespective of their accretion activity level. While it is widely accepted that f_{occ} is nearly 100% in local massive galaxies with stellar masses $M_* \gtrsim 10^{10} M_\odot$, it is not yet clear whether MBHs are ubiquitous in less-massive galaxies. In this work, we present new constraints on f_{occ} based on over 20 yr of Chandra imaging data for 1606 galaxies within 50 Mpc. We employ a Bayesian model to simultaneously constrain f_{occ} and the specific accretion-rate distribution function, $p(\lambda)$, where the specific accretion rate is defined as $\lambda = L_X/M_*$, where L_X is the MBH accretion luminosity in the 2–10 keV range. Notably, we find that $p(\lambda)$ peaks around $10^{28} \text{ erg s}^{-1} M_\odot^{-1}$; above this value, $p(\lambda)$ decreases with increasing λ , following a power law that smoothly connects with the probability distribution of bona fide active galactic nuclei. We also find that the occupation fraction decreases dramatically with decreasing M_* : in high-mass galaxies ($M_* \approx 10^{11-12} M_\odot$), the occupation fraction is $>93\%$ (a 2σ lower limit), and then declines to $66\%_{-7\%}^{+8\%}$ (1σ errors) between $M_* \approx 10^{9-10} M_\odot$, and to $33\%_{-9\%}^{+13\%}$ in the dwarf galaxy regime between $M_* \approx 10^{8-9} M_\odot$. Our results have significant implications for the normalization of the MBH mass function over the mass range most relevant for tidal disruption events, extreme mass ratio inspirals, and MBH merger rates that upcoming facilities are poised to explore.

Unified Astronomy Thesaurus concepts: [Astrophysical black holes \(98\)](#); [Dwarf galaxies \(416\)](#); [Galaxy nuclei \(609\)](#)

Materials only available in the online version of record: machine-readable tables

1. Introduction

Massive black holes (MBHs) residing at the centers of nearby galaxies typically accrete material at rates far below their Eddington limits (e.g., L. C. Ho 2009; J. Kormendy & L. C. Ho 2013), making them comparatively difficult to identify using the diagnostic techniques conventionally applied to active galactic nuclei (AGNs). Dynamical methods—tracing the motion of stars and gas in the vicinity of the MBH—bypass this limitation, indicating that nearly all Local Volume galaxies as large as the Milky Way, or larger, host an MBH. One of the most thoroughly studied examples is the Galactic center MBH, Sgr A*, which has a dynamical black hole mass

of $M_{\text{BH}} \approx 4 \times 10^6 M_\odot$ (A. M. Ghez et al. 2008; S. Gillessen et al. 2009) and an Eddington ratio close to 10^{-9} (Event Horizon Telescope Collaboration et al. 2022).

It remains unclear whether central MBHs are ubiquitous in lower-mass galaxies, particularly dwarf galaxies with stellar masses of $M_* \lesssim 10^{9.5} M_\odot$ (J. E. Greene et al. 2020; A. E. Reines 2022 and references therein). This stellar mass range corresponds to black hole masses of $M_{\text{BH}} \lesssim 10^6 M_\odot$, based on the local $M_{\text{BH}}-M_*$ scaling relation established for massive galaxies (e.g., A. E. Reines & M. Volonteri 2015), albeit the scaling relation in the dwarf regime is poorly constrained (e.g., J. E. Greene et al. 2020). Due to the small sphere of influence, dynamical evidence for MBHs can only be acquired in a handful of such dwarf galaxies (D. D. Nguyen et al. 2018, 2019; J. E. Greene et al. 2020).

Quantifying the active fraction (f_a ; e.g., V. F. Baldassare et al. 2020; F. Zou et al. 2023; D. Ohlson et al. 2024; R. Pucha et al. 2025) and, more importantly, the true black hole



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

occupation fraction (f_{occ}) as a function of M_* is crucial to several high-impact astrophysical problems. For example, the occupation fraction in dwarf galaxies is thought to be highly sensitive to the black hole seeding mechanism at high redshift (M. Volonteri 2012; P. Natarajan 2014; B. Agarwal et al. 2016; R. Valiante et al. 2016; A. Ricarte & P. Natarajan 2018; T. E. Woods et al. 2019; K. Inayoshi et al. 2020). Models suggest that, by $z=0$, virtually all dwarf galaxies will contain an MBH if stellar-sized, Population III remnants ($M_* \approx 10^2 M_\odot$; e.g., C. L. Fryer et al. 2001) provide the dominant seeding mode, as opposed to global gas collapse on galaxy-size scales ($M_* \approx 10^{5-6} M_\odot$; e.g., M. Habouzit et al. 2016), which leads to extremely low occupation fractions in dwarfs; if seeding occurs preferentially via gravitational runaway in dense stellar clusters ($M_* \approx 10^{3-4} M_\odot$; e.g., Y. Sakurai et al. 2017), then about 50% of local dwarfs can be expected to host an MBH. The expected rates of extreme mass ratio inspirals (S. Babak et al. 2017) and tidal disruption events (N. C. Stone & B. D. Metzger 2016) that will be detectable by the Vera C. Rubin Observatory (Ž. Ivezić et al. 2019) and, later, the Laser Interferometer Space Antenna (LISA; P. Amaro-Seoane et al. 2017) are highly sensitive to the occupation fraction in dwarf galaxies. Additionally, if the occupation fraction of dwarf galaxies were close to 100%, then their MBHs (rather than supernovae) could be entirely responsible for quenching star formation in these systems (G. Dashyan et al. 2018; S. J. Penny et al. 2018; C. M. Dickey et al. 2019). Pursuing a local approach to the MBH seeding conundrum is particularly compelling because it is orthogonal and complementary to ongoing searches at very high redshifts with the James Webb Space Telescope. These observations appear to be more easily explained with heavy seeding (e.g., Á. Bogdán et al. 2024; J. E. Greene et al. 2024; R. Maiolino et al. 2024; P. Natarajan et al. 2024).

Hard X-rays ($\gtrsim 2$ keV) are arguably the best electromagnetic signpost of MBH activity within the Local Volume because the X-ray sky is dark, and even quiescent MBHs emit nonthermal X-rays due to persistent, low-Eddington-ratio accretion (S. Pellegrini 2010; B. Miller et al. 2012; R. She et al. 2017; S. Bi et al. 2020). Furthermore, hard X-rays are highly penetrating and significantly less diluted by galaxy starlight (e.g., W. N. Brandt & D. M. Alexander 2015; W. N. Brandt & G. Yang 2022) compared to both soft X-rays as well as optical photons. AGNs in dwarf galaxies can be identified through multiple diagnostics (e.g., A. E. Reines et al. 2013; L. F. Sartori et al. 2015; V. F. Baldassare et al. 2020; K. L. Birchall et al. 2020, 2022; C. J. Burke et al. 2022; A. E. Reines 2022; E. J. Wasleske & V. F. Baldassare 2024) and to cosmological distances (K. Pardo et al. 2016; M. Mezcuca et al. 2018, 2023; F. Zou et al. 2023), even though they account for only a small fraction of the underlying population. These studies generally find that f_a is a few percent or less. Such low fractions are in rough agreement with theoretical expectations (F. Pacucci et al. 2021).

Robustly constraining f_{occ} based on f_a requires leveraging a large, unbiased galaxy sample that spans as wide as possible a range of stellar masses down to the lowest possible Eddington ratios. However, there is an inherent limit to how faint or distant we can probe: due to contamination from X-ray binaries (XRBs) to the galaxy nuclear X-ray signal, the chosen X-ray instrument’s point-spread function effectively limits the distance within which this experiment can be efficiently

carried out. In practice, the known scaling of XRB luminosity with stellar mass and star formation rate (SFR; B. D. Lehmer et al. 2019) sets a resolution- and distance-dependent limiting X-ray sensitivity, below which XRBs dominate the signal (E. J. Hodges-Kluck et al. 2020). This limit is highly sensitive to spatial resolution, making the Chandra X-ray Observatory the only currently operating facility capable of making a constraining measurement of f_{occ} in the local Universe.¹⁹

Based on on-axis Chandra Advanced CCD Imaging Spectrometer (ACIS) observations of 326 nearby early-type galaxies, E. Gallo & A. Sesana (2019) employ Bayesian inference to constrain the MBH occupation fraction as a function of M_* , finding $f_{\text{occ}} > 47\%$ for host stellar masses in the range $10^{9-10} M_\odot$ (68% confidence level), with no upper bound. Despite the improvement compared to the initial sample of approximately 200 early types considered by B. P. Miller et al. (2015, hereafter M15), E. Gallo & A. Sesana (2019) have little or no constraining power in the bona fide dwarf galaxy regime.

In this work, we expand on the aforementioned results by (i) leveraging the full Chandra Source Catalog²⁰ Release 2.1 (hereafter CSC 2.1; I. N. Evans et al. 2024) and considering a total of about 1600 target galaxies within 50 Mpc, and (ii) foregoing the assumption of a log-linear relation with normal scatter between the MBH X-ray luminosity and the host stellar mass. C. J. Burke et al. (2025, hereafter B25) recently derived f_{occ} based on similar underlying data, and we will discuss our significant differences in Section 4.1.2.

2. Sample and Data Description

Our parent sample is drawn from the 50 Mpc Galaxy Catalog²¹ (50MGC; D. Ohlson et al. 2024), which contains consistent and homogenized stellar mass, distance, and morphological type measurements for 15,424 galaxies within ≈ 50 Mpc. The 50MGC stellar masses are estimated through color-dependent mass-to-light ratios and are in good consistency with other works (see Section 4 of D. Ohlson et al. 2024). The volume completeness is $>50\%$ above $10^9 M_\odot$ and $>20\%$ above $10^8 M_\odot$. To search for nuclear X-ray sources, we take advantage of the CSC 2.1, which includes measured properties for 407,806 unique (compact and extended) X-ray sources detected in ACIS or High Resolution Camera (HRC) observations that were released publicly prior to the end of 2021. We only consider X-ray sources covered by ACIS because of the low energy resolution of HRC. This yields 1606 unique galaxies with M_* measurements in the 50MGC within the ACIS footprint, which are shown in the plane of M_* versus distance in Figure 1. Since only 10% of the 50MGC sources are covered by ACIS, complex selection effects may arise. We address this issue in Appendix B, and argue that the selection effects associated with intentional Chandra observations of MBHs are likely smaller than the statistical uncertainties.

The CSC 2.1 Optical/IR Counterparts Catalog²² further provides the CSC 2.1 counterparts in Gaia Data Release 3, Legacy Survey DR10, Pan-STARRS-1, and the Two Micron All-Sky Survey. Since these optical–IR positions are generally

¹⁹ If selected, the Advanced X-ray Imaging Satellite (C. S. Reynolds et al. 2023) is expected to offer a higher throughput at arcsecond-level resolution, thereby enabling even more effective X-ray-based constraints (E. Gallo et al. 2023).

²⁰ <https://cxc.cfa.harvard.edu/csc/2.1/>

²¹ <https://github.com/davidohlson/50MGC>

²² https://cxc.cfa.harvard.edu/csc/csc_crossmatches.html




Figure 1. M_* vs. distance for the target sample of 1606 galaxies. Blue and orange points represent X-ray-detected and undetected galaxies, respectively. There are points visually “piling up” at a distance of 16.5 Mpc because they are from the Virgo Cluster (see Section 3.3 in D. Ohlson et al. 2024).

more accurate than X-ray ones, we start by matching the nominal positions of the 50MGC galaxies with the CSC 2.1 ACIS counterpart catalog using a radius of $1''.5$. This yields 477 X-ray sources. D. Ohlson et al. (2024) already conducted careful visual inspections (including nuclear coordinates) and removed contaminants when constructing the 50MGC. We also visually inspect the optical and X-ray images of our sources and find no apparent issues.

We use Equation (11) from I. N. Evans et al. (2024) to determine in which energy bands each source was detected. For inferring X-ray photometry, we select the detected band based on the following categories: hard band (2–7 keV; 71% of the sources), broad band (0.5–7 keV; 29%), medium band (1.5–2 keV; one source only), soft band (0.5–1.2 keV; none), and ultrasoft band (0.2–0.5 keV; none). We prioritize higher energy bands to minimize absorption effects and reduce contamination from extended hot gas. We further exclude all the extended sources identified by the CSC’s algorithm `extent_code`. This yields a total of 388 central, point-like X-ray nuclei. For those, we estimate the corresponding 2–10 keV flux (f_X) using aperture photometry within the 90% enclosed-energy radius (i.e., `flux_aper90` in the CSC) and assuming a photon index of 1.8 to convert the photometry in different energy bands to a common 2–10 keV range. The uncertainty from the assumed photon index would effectively be absorbed into the width of the specific accretion-rate distribution function (SARDF) later on, which is measured to be $\sigma = 1$ dex (see Section 3 for more details). However, varying the photon index between 1.4 and 2.2 could cause an at-most 0.15 dex difference, which is much smaller than $\sigma = 1$ dex and thus has a minimal impact. We note that different off-axis angles yield different values of the 90% enclosed-energy radii. The median aperture radius is $1''.1$, and the 25% and 75% percentiles are $1''.0$ and $1''.4$, respectively, albeit the radius can reach $20''$ in a handful of extreme cases.

We treat the remaining sources, including both undetected sources and extended ones, as upper limits, which are derived as follows. If there are X-ray sources detected within $5''$, we set the upper limit to $\max\{f_X + 3\sigma(f_X), 2f_X\}$, where f_X and $\sigma(f_X)$ are the X-ray flux and corresponding uncertainty of the X-ray flux, respectively. We set a minimum value of $2f_X$ to ensure that the limit is sufficiently above f_X even when $\sigma(f_X)$ is small.




Figure 2. Nuclear X-ray luminosity vs. M_* . Blue points represent sources with compact X-ray nuclei, with increased color transparency reflecting an increasing likelihood of XRB contamination to the signal. Most of the blue points are not transparent because most sources have low P_{XRB} . Downward orange triangles represent upper limits, which include both nondetections as well as (few) extended sources. The solid red and violet lines are nonparametric, Akritas–Theil–Sen (ATS) regression lines for massive galaxies (with $M_* > 10^{10} M_\odot$) and all the galaxies, respectively. The solid black line represents the best-fit $\log L_X = \log M_* + \log \lambda = \log M_* + \alpha + \beta(\log M_* - 10)$ relation (Equation (9)). The shaded region, dashed lines, and dotted lines represent the 1σ , 2σ , and 3σ uncertainties, respectively.

This conservative choice accommodates cases where (i) the nuclear source is extended and thus was removed from our fiducial sample; (ii) one of the nearby sources is the actual X-ray counterpart but was missed during matching (perhaps owing to an uncertain galaxy center determination); or (iii) nearby X-ray-bright sources may elevate the threshold sensitivity. In most other cases, which account for 88% of all instances, we adopt the sensitivity threshold quoted by CSC 2.1 as the actual upper limits.

Fluxes are converted to 2–10 keV luminosities (L_X) using the distances quoted in 50MGC, which use the best available redshift-independent and redshift-based distances. Figure 2 shows the 1606 target galaxies (out of which 388 have point-like nuclear X-ray sources) in the M_* – L_X plane. As expected, the X-ray-detected nuclei are much fainter than typical AGNs, whose luminosities usually exceed 10^{42} erg s^{-1} . Three of the X-ray-detected nuclei belong to galaxies with $M_* < 10^8 M_\odot$ and have luminosities exceeding $L_X \gtrsim 10^{39}$ erg s^{-1} . All were previously reported in the literature; more details on these systems are given in Appendix A.

The figure also presents some representative regression lines. The solid red and violet lines are based on the ATS estimator (M. G. Akritas et al. 1995; D. R. Helsel 2011; E. D. Feigelson & G. J. Babu 2012), which is a nonparametric linear regression method for censored data (i.e., including detection limits). The red line is for galaxies more massive than $M_* = 10^{10} M_\odot$, which generally have $f_{\text{occ}} = 1$, while the violet line is for all the sources. As we will discuss further in Section 3, adopting a linear relation to depict the relation is only sensible for galaxies actually hosting MBHs—whereas a large fraction of the upper limits at $M_* < 10^{10} M_\odot$ may be due to not hosting MBHs. Therefore, the red line is more scientifically relevant, although it relies on extrapolations down to $M_* < 10^{10} M_\odot$; the violet line effectively represents the extreme case of $f_{\text{occ}} = 1$ across the whole M_* range. These nonparametric lines only serve as approximate presentations,

whereas our parametric, scientifically motivated model will be presented in Section 3, whose corresponding relation is represented by the black line in Figure 2. The black line is roughly between the two ATS nonparametric lines and is thus reasonable at least at face value. More detailed validations of our model assumptions will be discussed in Sections 3 and 4.

2.1. X-Ray Binary Contamination Assessment

To estimate the sample’s active fraction, we must first assess, for each nuclear X-ray source, the likelihood that the measured X-ray luminosity originates from XRBs rather than from low-level accretion onto a central MBH. This issue was explored at length by A. Foord et al. (2017) and N. Lee et al. (2019), so we follow their approach here. The XRB population include long-lived low-mass XRBs with donor star masses $\lesssim 2 M_{\odot}$ and young high-mass XRBs with donor star masses $\gtrsim 8 M_{\odot}$, and XRBs with intermediate-mass donors quickly evolve into low-mass XRBs due to a strong accretion instability. For the low- and high-mass XRB populations, we adopt the X-ray luminosity functions derived by B. D. Lehmer et al. (2019, 2021). The former depends on the host galaxy’s stellar mass, while the latter depends on the host’s SFR and metallicity.

The 50MGC provides host-galaxy M_{\star} ; to measure SFRs, we proceed as follows. The Local Volume Galaxy catalog²³ (I. D. Karachentsev et al. 2013) provides SFRs for 50MGC galaxies within ≈ 11 Mpc, measured from H α or far-UV fluxes (I. D. Karachentsev & E. I. Kaisina 2013), and we adopt their results for these systems. For others, we infer SFRs by fitting their spectral energy distributions (SEDs). Photometry of nearby galaxies through typical source-detection pipelines is often less reliable due to oversegmentation into multiple components. Therefore, we rely on dedicated catalogs and adopt the far-UV to mid-infrared photometry compiled in the NASA–Sloan Atlas (NSA)²⁴ and the Siena Galaxy Atlas 2020 (SGA-2020; J. Moustakas et al. 2023).²⁵

The 50MGC already provides matched identifiers in the NSA. For SGA-2020, most of our sources with central compact X-ray detections (304 out of 388) can be matched using a 2" matching radius, while unmatched sources primarily reside outside the SGA-2020 footprint. We fit the corresponding SEDs with Code Investigating Galaxy Emission v2025.0 (M. Boquien et al. 2019; G. Yang et al. 2020), following the parameter settings in Table 4 of F. Zou et al. (2022), except that we do not include the X-ray module. Briefly, we adopt the stellar templates from G. Bruzual & S. Charlot (2003), the Chabrier initial mass function (G. Chabrier 2003), a delayed star formation history, the dust attenuation model from D. Calzetti et al. (2000), and the dust emission templates from D. A. Dale et al. (2014). We refer readers to Section 3 of F. Zou et al. (2022) for further discussion of the SED fitting methodology. The fitting returns both M_{\star} and SFR. We have verified that the returned M_{\star} values are in good agreement with the 50MGC values, with a median offset of 0.1 dex and a scatter of 0.2 dex. However, we will retain the 50MGC M_{\star} for the remaining analyses to maintain consistency. In total, 86% of our galaxies have SFR measurements.

Next, we turn to high-resolution photometric data (available for 60% of the X-ray-detected galaxies) to estimate the nuclear M_{\star} and SFR. Specifically, for those galaxies with Hubble Space Telescope (HST) data, we measure the fraction of light by fitting isophotes to the HST imaging. We use observations from either the Advanced Camera for Surveys Wide Field Camera or the Wide Field Camera 3 UVIS/IR. In order to capture the underlying stellar population and mitigate contamination from star-forming regions, we restrict our measurements to data in wide, relatively red HST filters. Filters used included F555W, F606W, F814W, F850LP, F110W, and F160W. We calculate the enclosed light fraction, γ , within an aperture radius equal to the 90% enclosed-energy radius of the compact X-ray source. The nuclear M_{\star} and SFR are then obtained by multiplying the total M_{\star} and SFR by γ , thereby assuming a constant mass-to-light ratio.

We follow N. Lee et al. (2019) to calculate the probability P_{XRB} that the expected nuclear XRB emission exceeds a given L_{X} . The expected number of XRBs is $\mathbb{E}(N) = \int \frac{dN}{dL} dL$, where $\frac{dN}{dL}$ is the (M_{\star} - and SFR-dependent) XRB X-ray luminosity function. This number is a Poisson random variable, and the total XRB X-ray luminosity is given by $L_{\text{X}}^{\text{XRB}} = \sum_{i=1}^N L_{\text{X},i}^{\text{XRB}}$, where $L_{\text{X},i}^{\text{XRB}}$ is the individual XRB luminosity drawn from the XRB X-ray luminosity function in B. D. Lehmer et al. (2019, 2021). We note that this approach is notably different from, and arguably superior to, discarding as “contaminated” all nuclear X-ray sources with luminosity below 2 or 3 times the expected luminosity arising from XRBs (e.g., M. Mezcuca et al. 2018; K. L. Birchall et al. 2020; S. D. Bykov et al. 2024) for two reasons. First, the latter adopts a normal approximation that cannot accurately depict the tail distribution of $L_{\text{X}}^{\text{XRB}}$, to which P_{XRB} is sensitive. Second, considerable information is lost when simply discarding contaminated sources.

We generate 10^4 realizations of $L_{\text{X}}^{\text{XRB}}$ and derive P_{XRB} with the following expression: $P_{\text{XRB}} = P(L_{\text{X}}^{\text{XRB}} \geq L_{\text{X}})$. P_{XRB} is a function of the nuclear L_{X} and the aperture radius r (expressed in physical units), plus other source-specific quantities, i.e., M_{\star} , SFR, and the surface brightness profile. We further define $\lambda = L_{\text{X}}/M_{\star}$ (in units of $\text{erg s}^{-1} M_{\odot}^{-1}$) to normalize L_{X} over M_{\star} (see Section 3.1 for more discussions of λ). We take the $P_{\text{XRB}}(\lambda, r; M_{\star}, \text{SFR}, \gamma)$ functions derived above for the galaxies with HST data and calculate their median $P_{\text{XRB}}(\lambda, r)$ separately for early and late types (the latter tend to have higher high-mass XRB contributions). We use these median functions to estimate P_{XRB} for the galaxies with no available HST data, which are assumed to follow a Sérsic profile with index = 4 for early types and index = 1 for late types. The resulting median P_{XRB} of our sample is 5%, and the 25% and 75% percentiles are 1% and 18%, respectively.

To verify the applicability of these median $P_{\text{XRB}}(\lambda, r)$ functions to the sources lacking HST coverage, we apply them to predict P_{XRB} for the HST-covered galaxies, which have their own P_{XRB} measurements available. We find that the predicted and the actual P_{XRB} are reasonably consistent across the M_{\star} range—the median values of the actual and predicted P_{XRB} for HST-covered galaxies are 2% and 5%, respectively.

We note that the adopted X-ray luminosity function, and the ensuing P_{XRB} , already account for the possible presence of ultraluminous X-ray sources (ULXs; see discussion in B. D. Lehmer et al. 2021). In comparison, the scaling relation with M_{\star} and SFR obtained by K. Kovalakas et al. (2020) yields

²³ <https://relay.sao.ru/lv/lvgdb/>

²⁴ <https://www.sdss4.org/dr17/manga/manga-target-selection/nsa/>

²⁵ <https://www.legacysurvey.org/sga/sga2020/>

Table 1
Source Properties of the Sample

Name	R.A. (deg)	Decl. (deg)	Distance (Mpc)	$\log L_X$ (erg s^{-1})	$\log M_*$ (M_\odot)	$\log \text{SFR}$ ($M_\odot \text{ yr}^{-1}$)	P_{XRB}
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2MASXJ15150097+5525555	228.753971	55.432026	49.67	39.22 ± 0.16	9.66	-2.79	2.1%
ESO005-004	91.423788	-86.631870	23.66	40.50 ± 0.04	10.17	NaN	0.6%
ESO097-013	213.291254	-65.339152	4.20	38.75 ± 0.02	10.06	NaN	0.6%
ESO121-006	91.874124	-61.807596	14.14	40.37 ± 0.01	9.47	-0.16	0.0%
ESO138-001	252.834282	-59.234564	29.37	41.25 ± 0.01	9.81	NaN	0.0%
...

Note. Column (1): galaxy names as in the 50MGC. Columns (2) and (3): J2000 coordinates. Column (4): distances. Column (5): nuclear 2–10 keV luminosities. Column (6): host-galaxy stellar masses. Column (7): host SFR. Column (8): the probability that the detected nuclear X-ray emission is from XRBs.

(This table is available in its entirety in machine-readable form in the [online article](#).)

an average rate of 0.02 ULXs with $L_X > 10^{39} \text{ erg s}^{-1}$ per galactic nucleus in our sample, which is comparable to our derived P_{XRB} . Although there are inevitable uncertainties when deriving P_{XRB} —for example, the Sérsic index of low-mass early-type galaxies may deviate from the nominal value of four (J. Kormendy et al. 2009)—the exact values of P_{XRB} are generally not critical. This is because they are sufficiently small, and the detected X-ray signals primarily originate from MBHs.

We summarize the properties of all the 1606 sources in our sample in Table 1.

3. Bayesian Inference

It is apparent from Figures 1 and 2 that the fraction of detected X-ray nuclei decreases dramatically with decreasing M_* : 54% of sources with $M_* > 10^{10} M_\odot$ have detected compact X-ray nuclei, while this fraction drops to 12% for $M_* < 10^{10} M_\odot$. Qualitatively, this could imply that MBHs are ubiquitous in low-mass galaxies but are significantly more underluminous than the detection limits, that the MBH occupation fraction decreases at lower masses, or a combination of the two. In this section, we quantitatively constrain f_{occ} .

3.1. Modeling the Occupation Fraction and Specific Accretion-rate Distribution

Our methodology updates the framework presented in both M15 and B25. We first define the specific accretion rate, $\lambda = L_X/M_*$, as an empirical proxy for the Eddington ratio. This choice aligns with the definition used in several recent works on AGN accretion rates (e.g., A. Bongiorno et al. 2016; G. Yang et al. 2018; F. Zou et al. 2024b; P. Guetzoyan et al. 2025) and aims to bypass the uncertainties introduced by using the actual Eddington ratio $\lambda_{\text{Edd}} = k_{\text{bol}} L_X / L_{\text{Edd}}$. These uncertainties include the choice of the scaling relation between M_* and M_{BH} (and its scatter) as well as the X-ray to bolometric correction k_{bol} . For reference, $\lambda = 10^{34} \text{ erg s}^{-1} M_\odot^{-1}$ corresponds to an Eddington ratio of $\lambda_{\text{Edd}} \approx 1$ when assuming $k_{\text{bol}} = 25$ and a typical scaling relation of $M_{\text{BH}} = 0.002 M_*$ (as discussed in Section 4.4, k_{bol} likely depends on both L_X and λ_{Edd} ; however, because k_{bol} is not included in the definition of λ , its exact value does not affect our results).

We define the SARDF as the conditional probability density per unit $\log \lambda$ for a galaxy with a central MBH to have a specific accretion rate of λ . The SARDF is represented by $p(\lambda) = p(\log \lambda | M_*, I_{\text{MBH}} = 1)$, where $I_{\text{MBH}} = 1$ indicates the

presence of a central MBH. We first consider the simple case where there is no XRB contamination (i.e., $P_{\text{XRB}} = 0$) to help illustrate our methodology. As will be shown in Section 3.3, the analysis for a nonzero P_{XRB} can be generalized from this case. Our data are described by the following distribution:

$$\log \lambda \sim \text{Mixture} \{ [p(\lambda), \delta(-\infty)], [f_{\text{occ}}, 1 - f_{\text{occ}}] \}, \quad (1)$$

where ‘‘Mixture’’ stands for a mixture model in statistics. This means a galaxy has a probability of f_{occ} of hosting a central MBH, in which case its $\log \lambda$ follows the distribution $p(\lambda)$. Conversely, the galaxy has a probability of $(1 - f_{\text{occ}})$ of not hosting a central MBH, resulting in $\log \lambda = -\infty$. Therefore, we need to simultaneously constrain both $p(\lambda)$ and f_{occ} .

Since $p(\lambda)$ and f_{occ} are coupled in Equation (1), the constraints on f_{occ} depend on the choice of $p(\lambda)$. Previous work by M15 assumes a log-linear relation with lognormal scatter between the MBH accretion luminosity in X-rays, L_X , and the host-galaxy stellar mass, M_* . The more recent, multiwavelength investigation by B25 assumes that the luminosity scatter at a given mass is described by the Eddington ratio distribution function (ERDF). The probability distribution of the ERDF is modeled as a mass-independent, smoothly broken power law, implying that galaxies have an increasingly high likelihood of hosting underluminous MBHs. This choice is based on A. K. Weigel et al. (2017), who conclude that the inferred ERDF for (luminous) AGNs follows a broken power law and is roughly mass independent in the local Universe.²⁶ The ERDF chosen by B25 introduces a key difference compared to M15’s approach, as B25’s ERDF is inconsistent with lognormal scatter in the L_X – M_* relation. In this work, we will adopt a flexible distribution family for $p(\lambda)$ to encompass both the possibilities presented by M15 and B25, and we will revisit the viability of each functional form in Section 4.2.

3.2. Parameterization

To parameterize $f_{\text{occ}}(M_*)$ and $p(\lambda)$, we aim to adopt flexible functional forms that can empirically represent as broad a range of behaviors as possible. For f_{occ} , we follow B25 and assume that the expression parameterizing the MBH occupation fraction as a function of M_* is a generalized logistic function. This family of functions is sufficiently flexible to

²⁶ The mass independence is not necessarily correct in the distant Universe. As shown by J. Aird et al. (2018) and F. Zou et al. (2024b), the ERDF becomes increasingly mass dependent as the redshift increases up to $z \sim 2$.




Figure 3. Visual illustration of the chosen parameterization for f_{occ} (left; Equation (2)) and $p(\lambda)$ (right; Equation (7)). Left: in this plot, $\log M_0 = 9.5$. The black horizontal line represents the limiting f_{occ} as $\delta \rightarrow 0$ for any θ . The other black curves correspond to $\delta \rightarrow +\infty$ at different θ values, reaching $f_{\text{occ}} = 1$ at $\log M_{\star} = \log M_0$. The colored curves show how f_{occ} evolves with δ when $\theta = 1$. As δ increases, the curve becomes less flat and approaches the corresponding limiting black curve for $\theta = 1$. Right: the plot illustrates $p(\lambda)$ for different (μ, σ, ξ) values, as indicated in the legend. When $\xi = 1$, $p(\lambda)$ is a normal distribution, shown by the black curve. The blue and orange curves have different ξ values, resulting in heavier tails: on the left when $\xi > 1$ and on the right when $\xi < 1$. The green curve, compared to the orange curve, demonstrates that σ controls the width of $p(\lambda)$. The red curve shows that our $p(\lambda)$ can mimic a power-law distribution over the plotted $\log \lambda$ range when both μ and ξ are small.

model typical growth patterns (F. J. Richards 1959). It has the form

$$f_{\text{occ}}(M_{\star}) = [1 + \delta e^{-\theta \delta (\log M_{\star} - \log M_0)}]^{-\frac{1}{\delta}}, \quad (2)$$

where $\log M_0$, δ , and θ are free parameters. A minor difference between our parameterization and that of B25 is the inclusion of an additional δ factor in the exponentiation term, which reduces the correlation between θ and δ . In our definition, θ controls the limiting behavior of f_{occ} with respect to δ , effectively decoupling from δ at large values. Mathematically, it can be shown that

$$\lim_{\delta \rightarrow 0} f_{\text{occ}} = \frac{1}{e}, \quad (3)$$

$$\lim_{\delta \rightarrow +\infty} f_{\text{occ}} = \min\{e^{\theta (\log M_{\star} - \log M_0)}, 1\}. \quad (4)$$

The above equation indicates that θ (when positive) directly affects f_{occ} in low-mass galaxies, particularly when $\delta \gg 1$. The left panel of Figure 3 illustrates how these parameters influence the shape of f_{occ} . In brief, $\log M_0$ shifts the function along the $\log M_{\star}$ axis. Increasing θ suppresses f_{occ} when $\delta \gg 1$, and increasing δ makes the curve less flat and closer to the limiting $f_{\text{occ}}(\delta \rightarrow +\infty)$ curve.

For the parameterization of $p(\lambda)$, we assume that it follows a Box–Cox distribution for a given M_{\star} , and that the distribution can shift linearly with $\log M_{\star}$. This family of distributions includes commonly adopted models for the Eddington ratio or specific accretion-rate distribution, such as the normal distribution (M15) and the power-law distribution (B25). The Box–Cox distribution, also known as the power-normal distribution (e.g., J. Freeman & R. Modarres 2006), originates from the Box–Cox transformation (G. E. P. Box & D. R. Cox 1964):

$$Y = \frac{X^{\xi} - 1}{\xi}, \quad X > 0, \quad (5)$$

which reduces to $Y = \ln X$ for $\xi = 0$. The Box–Cox transformation is a statistical technique used to stabilize variance and make data more closely approximate a normal distribution. It is particularly useful for transforming skewed data X , often with sizable tails, into a more symmetrical bell-shaped distribution for Y . In our model, we will use $X = \log \lambda$ and set ξ as a free parameter. If Y follows a (truncated)²⁷ normal distribution, then $X = \log \lambda$ follows a Box–Cox distribution. Note that $Y = X - 1$ if $\xi = 1$, making a truncated normal distribution a special case of the Box–Cox distribution. Additionally, for certain values of ξ , a Box–Cox distribution can model a power-law-like tail over a large range of X .

We assume the following truncated normal distribution for Y :

$$Y \sim \begin{cases} \mathcal{TN}\left(\frac{\mu^{\xi} - 1}{\xi}, \mu^{2\xi - 2}\sigma^2; -\frac{1}{\xi}, +\infty\right), & \xi > 0, \\ \mathcal{TN}\left(\frac{\mu^{\xi} - 1}{\xi}, \mu^{2\xi - 2}\sigma^2; -\infty, -\frac{1}{\xi}\right), & \xi < 0, \end{cases} \quad (6)$$

where μ and σ are parameters that approximate, but are not identical to, the mean and standard deviation of $\log \lambda$, and $Y \sim \mathcal{TN}(a_1, a_2; a_3, a_4)$ represents a normal distribution with a mean of a_1 and a variance of a_2 , constrained within the interval $a_3 < Y < a_4$.

Inverting the transformation in Equation (5), we obtain the following expression for the SARDF:

$$p(\lambda) = \frac{1}{\sqrt{2\pi}\sigma A} \left(\frac{\log \lambda}{\mu}\right)^{\xi - 1} e^{-\frac{1}{2}\left(\frac{\mu}{\xi\sigma}\right)^2 \left[\left(\frac{\log \lambda}{\mu}\right)^{\xi} - 1\right]^2}, \quad (7)$$

²⁷ Y is truncated by definition because X is assumed to be positive. Specifically, $Y > -\frac{1}{\xi}$ for $\xi > 0$, and $Y < -\frac{1}{\xi}$ for $\xi < 0$. Truncation is almost negligible in our modeling because $X = \log \lambda \gg 0$.

$$A = \Phi\left(\frac{\mu}{|\xi|\sigma}\right), \quad (8)$$

where $\text{sgn}(x)$ is the sign function, and $\Phi(x) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{x}{\sqrt{2}}\right)\right]$ is the cumulative distribution function of the standard normal distribution.

We illustrate how these parameters influence $p(\lambda)$ in the right panel of Figure 3. Changing μ shifts the curve, while σ controls the width of $p(\lambda)$. For a given set of (μ, σ) , ξ determines the overall shape of $p(\lambda)$. When $\xi = 1$, $p(\lambda)$ is a normal distribution, as shown by the black curve in the figure. As ξ increases (or decreases), $p(\lambda)$ develops a heavier left (or right) tail, as illustrated by the blue and orange curves. A notable case occurs when $p(\lambda)$ mimics a power-law distribution, represented by the red curve; in this case, both μ and ξ are small, and the distribution approximates a power law across the plotted range of $\log \lambda$. To account for potential changes in $p(\lambda)$ with M_* , we further assume that μ is linearly correlated with $\log M_*$:

$$\mu = \alpha + \beta(\log M_* - 10), \quad (9)$$

where α and β are free coefficients. This allows $p(\lambda)$ to shift linearly with $\log M_*$ while maintaining roughly the same overall shape for different M_* values.

3.3. Likelihood

We explicitly derive the likelihood from Equation (1) below. We first present the likelihood for the case when the data are noiseless, and then we derive the full likelihood with a rigorous treatment of the measurement uncertainties. The former is relatively simpler and helps to illustrate the model, while the latter, although much more complex, can intuitively be regarded as a perturbative correction to the noiseless likelihood. Hereafter, we use the subscript or superscript k to refer to a given galaxy with the identifier k in our sample.

3.3.1. Noiseless Likelihood

For an uncensored (i.e., compact X-ray detected) source k , the likelihood is given by

$$\ln \mathcal{L}_k^{\text{det}} = \ln p(\lambda_k) + \ln f_{\text{occ}}(M_k). \quad (10)$$

For censored sources (i.e., with L_X upper limits), Equation (1) reduces to

$$(I_k = 0) \sim \text{Bernoulli}(p_k), \quad (11)$$

$$p_k = f_{\text{occ}}(M_k) S_\lambda(\log \lambda_{\text{lim},k}), \quad (12)$$

where I_k is set to 0 (1) if $\lambda_k < \lambda_{\text{lim},k}$ ($\lambda_k \geq \lambda_{\text{lim},k}$); λ_{lim} is the upper limit of λ ; p_k is the probability that a source k has $\lambda_k \geq \lambda_{\text{lim},k}$; and $S_\lambda(\log \lambda_{\text{lim}}) = P(\{\log \lambda \sim p(\lambda)\} \geq \log \lambda_{\text{lim}})$ is the survival function corresponding to $p(\lambda)$.

$S_\lambda(\log \lambda_{\text{lim}})$ can be derived from Equations (5) and (6):

$$S_\lambda(\log \lambda_{\text{lim}}) = \begin{cases} \frac{1}{A} \left[1 - \Phi\left(\frac{\mu}{\xi\sigma} \left[\left(\frac{\log \lambda_{\text{lim}}}{\mu}\right)^\xi - 1 \right] \right) \right], & \xi > 0, \\ 1 - \frac{1}{A} \Phi\left(\frac{\mu}{\xi\sigma} \left[\left(\frac{\log \lambda_{\text{lim}}}{\mu}\right)^\xi - 1 \right] \right), & \xi < 0. \end{cases} \quad (13)$$

The likelihood for a censored source k is given by

$$\ln \mathcal{L}_k^{\text{cens}} = \ln(1 - p_k). \quad (14)$$

Note that the above equations hold when $P_{\text{XRB}} = 0$, as assumed in Equation (1). If $P_{\text{XRB}} \neq 0$, the log-likelihood function of an uncensored source is the weighted mean of $\ln \mathcal{L}^{\text{det}}$ and $\ln \mathcal{L}^{\text{cens}}$:

$$\ln \mathcal{L}_k = (1 - P_{\text{XRB},k}) \ln \mathcal{L}_k^{\text{det}} + P_{\text{XRB},k} \ln \mathcal{L}_k^{\text{cens}}, \quad (15)$$

where the first term refers to sources whose X-ray emission originates solely from MBH accretion, and the second term addresses sources whose X-ray emission is contaminated by XRBs. Consequently, the overall log-likelihood function is

$$\ln \mathcal{L} = \sum_{k \in \{\text{uncensored}\}} [(1 - P_{\text{XRB},k}) \ln \mathcal{L}_k^{\text{det}} + P_{\text{XRB},k} \ln \mathcal{L}_k^{\text{cens}}] + \sum_{k \in \{\text{censored}\}} \ln \mathcal{L}_k^{\text{cens}}. \quad (16)$$

3.3.2. Likelihood with Measurement Uncertainties

We now extend the derivation by incorporating measurement uncertainties into the likelihood function. This allows for a more realistic and rigorous statistical treatment of the data. Our data include uncertainties in both axes ($\log M_*$, $\log \lambda$), and it is generally a nontrivial task to account for uncertainties in both variables, especially for the abscissa. We derive the corresponding likelihood following the Bayesian approach described in B. C. Kelly (2007). The key idea is to treat the true values of $(\log M_*, \log \lambda)$ —which would be available in the case of perfectly precise measurements—as missing data, and then introduce and marginalize over these latent parameters for each data point. We denote (χ, ζ) as the unknown true values of $(\log M_*, \log \lambda)$. The likelihood for each source k is

$$\mathcal{L}_k = \iint p(\log M_k, \log \lambda_k | \chi_k, \zeta_k) p(\zeta_k | \chi_k) p(\chi_k) d\chi_k d\zeta_k. \quad (17)$$

Throughout this subsection, we denote $\mathcal{N}(\mathbf{m}, \Sigma)$ as either the univariate or multidimensional normal distribution with a mean vector \mathbf{m} and covariance matrix Σ . We also denote $\mathcal{N}(\mathbf{x} | \mathbf{m}, \Sigma)$ as the corresponding probability distribution function evaluated at \mathbf{x} .

Following the same approach as in Section 3.3.1, we first consider uncensored sources under the $p_{\text{XRB}} = 0$ condition, i.e., corresponding to $\mathcal{L}_k^{\text{det}}$ in Equation (10). The first term of the integrand in Equation (17), $p(\log M_k, \log \lambda_k | \chi_k, \zeta_k)$,

represents measurement uncertainties, such that

$$(\log M_k, \log \lambda_k) | (\chi_k, \zeta_k) \sim \mathcal{N}((\chi_k, \zeta_k), \Sigma_k), \quad (18)$$

$$\begin{aligned} \Sigma_k &= \begin{pmatrix} \Sigma_{11,k} & \Sigma_{12,k} \\ \Sigma_{12,k} & \Sigma_{22,k} \end{pmatrix} \\ &= \begin{pmatrix} \sigma^2 \{\log M_k\} & -\sigma^2 \{\log M_k | D_k\} \\ -\sigma^2 \{\log M_k | D_k\} & \sigma^2 \{\log f_{X,k}\} + \sigma^2 \{\log M_k | D_k\} \end{pmatrix}, \end{aligned} \quad (19)$$

where $\sigma \{\log M_k\}$ and $\sigma \{\log f_{X,k}\}$ are the uncertainties in $\log M_k$ and $\log f_X$, respectively, and D is the source distance. The $\sigma \{\log M_k\}$ values quoted by D. Ohlson et al. (2024) already include the (generally subdominant) distance uncertainties; the uncertainty contributed by the remaining, distant-independent factors is estimated to be $\sigma \{\log M_k | D_k\} = 0.136$ (see Section 4 in D. Ohlson et al. 2024). $\log \lambda$ is not affected by distance uncertainties because the distance dependence of both L_X and M_* canceled out. Therefore, only $\sigma \{\log M_k | D_k\}$ contributes to $\Sigma_{12,k} = \text{Cov}\{\log \lambda_k, \log M_k | \chi_k, \zeta_k\}$ and $\Sigma_{22,k} = \text{Var}\{\log \lambda_k | \chi_k, \zeta_k\}$.

The second term of the integrand in Equation (17), $p(\zeta_k | \chi_k)$, is set by Equation (1):

$$\begin{aligned} \zeta_k | \chi_k &\sim \text{Mixture} \{ [p_\lambda(\zeta_k | \chi_k), \delta(-\infty)], \\ & [f_{\text{occ}}(\chi_k), 1 - f_{\text{occ}}(\chi_k)] \}, \end{aligned} \quad (20)$$

where $p_\lambda(\zeta_k | \chi_k)$ is the same as in Equation (7) but with (χ_k, ζ_k) replacing $(\log M_*, \log \lambda)$. Therefore, similarly to Equation (10), we have

$$p(\zeta_k | \chi_k) = p_\lambda(\zeta_k | \chi_k) f_{\text{occ}}(\chi_k). \quad (21)$$

The third term of the integrand in Equation (17), $p(\chi_k)$, represents the underlying distribution of χ_k . Following B. C. Kelly (2007), we model it with a mixture of normal distributions:

$$p(\chi_k) = \sum_j w_j \mathcal{N}(\chi_k | t_j, s_j^2), \quad (22)$$

$$\sum_j w_j = 1. \quad (23)$$

In principle, the (w_j, t_j, s_j) parameters could be incorporated as free parameters within the occupation fraction model. However, this would introduce a substantial number of additional parameters (e.g., eight more for a three-component mixture), significantly increasing the model's dimensionality and computational cost. Since our sample size is very large, the $\log M_*$ distribution is already well constrained, so fixing (w_j, t_j, s_j) has a negligible impact on the inference and greatly improves computational efficiency. Since Equation (18) gives $\log M_k | \chi_k \sim \mathcal{N}(\chi_k, \Sigma_{11,k})$, we have that

$$\begin{aligned} p(\log M_k) &= \int p(\log M_k | \chi_k) p(\chi_k) d\chi_k \\ &= \sum_j \frac{w_j}{\sqrt{2\pi(\Sigma_{11,k} + s_j^2)}} e^{-\frac{(t_j - \log M_k)^2}{2(\Sigma_{11,k} + s_j^2)}}. \end{aligned} \quad (24)$$

By maximizing $\prod_k p(\log M_k)$ to fit the inferred $\log M_*$ distribution for our sample (including both censored and uncensored data), we can constrain (w_j, t_j, s_j) . We find that a

mixture of three normal distributions models the $\log M_*$ distribution well, with the corresponding parameters (w_j, t_j, s_j) given by (0.70, 8.75, 0.97), (0.17, 10.06, 0.35), and (0.13, 10.64, 0.25).

With these three terms, Equation (17) becomes

$$\begin{aligned} \mathcal{L}_k^{\text{det}} &= \sum_j w_j \iint \mathcal{N}((\log M_k, \log \lambda_k) | (\chi_k, \zeta_k), \Sigma_k) \\ &\quad \times \mathcal{N}(\chi_k | t_j, s_j^2) p_\lambda(\zeta_k | \chi_k) f_{\text{occ}}(\chi_k) d\chi_k d\zeta_k. \end{aligned} \quad (25)$$

After solving for the first two terms of the integrand above, $\mathcal{L}_k^{\text{det}}$ reduces to

$$\begin{aligned} \mathcal{L}_k^{\text{det}} &= \sum_j \frac{w_j}{\sqrt{2\pi(\Sigma_{11,k} + s_j^2)}} e^{-\frac{(t_j - \log M_k)^2}{2(\Sigma_{11,k} + s_j^2)}} \\ &\quad \times \iint \mathcal{N}((x_1, x_2) | (c_1^{kj}, c_2^{kj}), \eta^{kj}) g^k(x_1, x_2) dx_1 dx_2, \end{aligned} \quad (26)$$

$$c_1^{kj} = \frac{\Sigma_{11,k}}{\Sigma_{11,k} + s_j^2} (t_j - \log M_k), \quad (27)$$

$$c_2^{kj} = \frac{\Sigma_{12,k}}{\Sigma_{11,k} + s_j^2} (t_j - \log M_k), \quad (28)$$

$$\begin{aligned} \eta^{kj} &= \begin{pmatrix} \eta_{11}^{kj} & \eta_{12}^{kj} \\ \eta_{12}^{kj} & \eta_{22}^{kj} \end{pmatrix} \\ &= \begin{pmatrix} \frac{\Sigma_{11,k} s_j^2}{\Sigma_{11,k} + s_j^2} & \frac{\Sigma_{12,k} s_j^2}{\Sigma_{11,k} + s_j^2} \\ \frac{\Sigma_{12,k} s_j^2}{\Sigma_{11,k} + s_j^2} & \frac{\Sigma_{11,k} \Sigma_{22,k} - \Sigma_{12,k}^2 + \Sigma_{22,k} s_j^2}{\Sigma_{11,k} + s_j^2} \end{pmatrix}, \end{aligned} \quad (29)$$

$$\begin{aligned} g^k(x_1, x_2) &= p_\lambda(\log \lambda_k + x_2 | \log M_k + x_1) \\ &\quad \times f_{\text{occ}}(\log M_k + x_1). \end{aligned} \quad (30)$$

There is an important practical difference between our fitting procedure and typical regressions that assume Gaussian scatter. In the latter case, $g^k(x_1, x_2)$ is a Gaussian function, so the integration in Equation (26) can be solved analytically in closed form; previous treatments of measurement uncertainties usually rely implicitly on this property (e.g., B. C. Kelly 2007; D. W. Hogg et al. 2010). However, the integration in our $\mathcal{L}_k^{\text{det}}$ cannot be solved analytically. Accurate calculation would require a two-dimensional numerical integration for every data point during each evaluation of $\mathcal{L}_k^{\text{det}}$, which is not practically feasible for sampling, as numerical integrations are extremely slow. We therefore seek an approximate method to avoid numerical integration.

We note that the typical measurement uncertainty (median $\{\sqrt{\Sigma_{11,k}}\} = 0.14$ and median $\{\sqrt{\Sigma_{22,k}}\} = 0.19$) is small compared to the overall span of our data range (several dex). Therefore, in Equation (17), the main contribution to the integrand comes from the region around $(\log M_k, \log \lambda_k)$, due to the first term, $p(\log M_k, \log \lambda_k | \chi_k, \zeta_k)$. The remaining part of the integrand, $p(\zeta_k | \chi_k) p(\chi_k)$, varies much more slowly on a scale comparable to the total data range. Thus, it is a good approximation to perform a Taylor expansion of $p(\zeta_k | \chi_k)$

around $(\log M_k, \log \lambda_k)$, i.e., to expand $g^k(x_1, x_2)$ around $(0, 0)$:

$$g^k(x_1, x_2) \approx \sum_{i_1} \sum_{i_2} g_{i_1 i_2}^k x_1^{i_1} x_2^{i_2}, \quad (31)$$

where the coefficients $g_{i_1 i_2}^k$ are calculated with `TaylorSeries.jl` (L. Benet & D. P. Sanders 2019). The integration in Equation (26) reduces to evaluating moments of the normal distribution $\mathcal{N}((c_1^{kj}, c_2^{kj}), \eta^{kj})$, which are given by Isserlis's theorem (L. Isserlis 1918). We expand $g^k(x_1, x_2)$ up to the fourth order, and Equation (26) then becomes

$$\begin{aligned} \mathcal{L}_k^{\text{det}} \approx & \sum_j \frac{w_j}{\sqrt{2\pi(\Sigma_{11,k} + s_j^2)}} e^{-\frac{(t_j - \log M_k)^2}{2(\Sigma_{11,k} + s_j^2)}} \\ & \times \{g_{00}^k + g_{10}^k c_1^{kj} + g_{01}^k c_2^{kj} + g_{20}^k [(c_1^{kj})^2 + \eta_{11}^{kj}] \\ & + g_{11}^k [c_1^{kj} c_2^{kj} + \eta_{12}^{kj}] + g_{02}^k [(c_2^{kj})^2 + \eta_{22}^{kj}] \\ & + g_{30}^k [(c_1^{kj})^3 + 3c_1^{kj} \eta_{11}^{kj}] \\ & + g_{21}^k [(c_1^{kj})^2 c_2^{kj} + 2c_1^{kj} \eta_{12}^{kj} + c_2^{kj} \eta_{11}^{kj}] \\ & + g_{12}^k [c_1^{kj} (c_2^{kj})^2 + 2c_2^{kj} \eta_{12}^{kj} + c_1^{kj} \eta_{22}^{kj}] \\ & + g_{03}^k [(c_2^{kj})^3 + 3c_2^{kj} \eta_{22}^{kj}] \\ & + g_{40}^k [(c_1^{kj})^4 + 6(c_1^{kj})^2 \eta_{11}^{kj} + 3(\eta_{11}^{kj})^2] \\ & + g_{31}^k [(c_1^{kj})^3 c_2^{kj} + 3(c_1^{kj})^2 \eta_{12}^{kj} + 3c_1^{kj} c_2^{kj} \eta_{11}^{kj} + 3\eta_{11}^{kj} \eta_{12}^{kj}] \\ & + g_{22}^k [(c_1^{kj})^2 (c_2^{kj})^2 + (c_1^{kj})^2 \eta_{22}^{kj} + (c_2^{kj})^2 \eta_{11}^{kj}] \\ & + 4c_1^{kj} c_2^{kj} \eta_{12}^{kj} + \eta_{11}^{kj} \eta_{22}^{kj} + 2(\eta_{12}^{kj})^2] \\ & + g_{13}^k [c_1^{kj} (c_2^{kj})^3 + 3(c_2^{kj})^2 \eta_{12}^{kj} + 3c_1^{kj} c_2^{kj} \eta_{22}^{kj} + 3\eta_{12}^{kj} \eta_{22}^{kj}] \\ & + g_{04}^k [(c_2^{kj})^4 + 6(c_2^{kj})^2 \eta_{22}^{kj} + 3(\eta_{22}^{kj})^2]\}. \quad (32) \end{aligned}$$

For censored sources, the likelihood is

$$\begin{aligned} \mathcal{L}_k^{\text{cens}} = & p(I_k = 0, \log M_k) = p(\log M_k) \\ & - p(I_k = 1, \log M_k). \quad (33) \end{aligned}$$

$p(\log M_k)$ is given in Equation (24), and the second term is

$$\begin{aligned} p(I_k = 1, \log M_k) &= \iint p(I_k = 1 | \zeta_k) p(\log M_k | \chi_k) p(\zeta_k | \chi_k) p(\chi_k) d\chi_k d\zeta_k \\ &= \int p(\log M_k | \chi_k) p(\chi_k) S_\lambda(\log \lambda_{\text{lim},k} | \chi_k) f_{\text{occ}}(\chi_k) d\chi_k \\ &= \sum_j w_j \int \mathcal{N}(0; x, \Sigma_{11}) \mathcal{N}(x; t_j - \log M_k, s_j^2) h(x) dx \\ &= \sum_j \frac{w_j}{\sqrt{2\pi(\Sigma_{11,k} + s_j^2)}} e^{-\frac{(t_j - \log M_k)^2}{2(\Sigma_{11,k} + s_j^2)}} \\ & \times \int \mathcal{N}(x; c_1^{kj}, \eta_{11}^{kj}) h^k(x) dx, \quad (34) \end{aligned}$$

$$h^k(x) = S_\lambda(\log \lambda_{\text{lim},k} | \log M_k + x) f_{\text{occ}}(\log M_k + x). \quad (35)$$

We expand $h^k(x)$ up to the fourth order:

$$h^k(x) \approx \sum_{i=1}^4 h_i^k x^i. \quad (36)$$

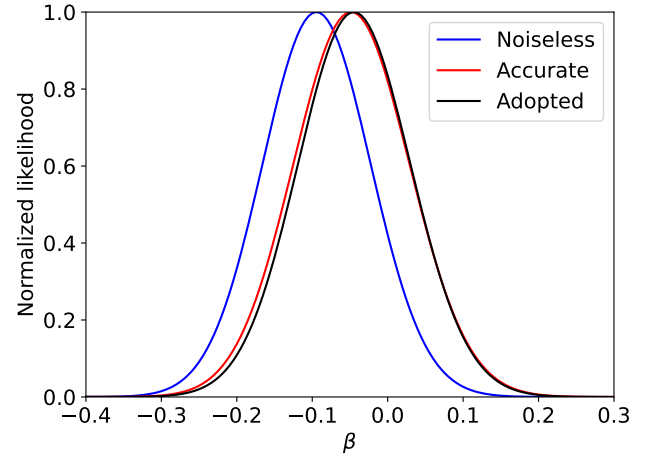


Figure 4. Example comparison among the likelihood ignoring measurement uncertainties (blue), the likelihood accurately accounting for measurement uncertainties (red), and our adopted approximation (black). All likelihoods are normalized to a maximum value of one. Here, β is varied while all other parameters are fixed at their median posterior values. The near-complete overlap of the black and red curves indicates that our adopted likelihood closely approximates the fully accurate likelihood.

Equation (33) then reduces to

$$\begin{aligned} \mathcal{L}_k^{\text{cens}} \approx & \sum_j \frac{w_j}{\sqrt{2\pi(\Sigma_{11,k} + s_j^2)}} e^{-\frac{(t_j - \log M_k)^2}{2(\Sigma_{11,k} + s_j^2)}} \\ & \times \{1 - h_0^k - h_1^k c_1^{kj} - h_2^k [(c_1^{kj})^2 + \eta_{11}^{kj}] \\ & - h_3^k [(c_1^{kj})^3 + 3c_1^{kj} \eta_{11}^{kj}] \\ & - h_4^k [(c_1^{kj})^4 + 6(c_1^{kj})^2 \eta_{11}^{kj} + 3(\eta_{11}^{kj})^2]\}. \quad (37) \end{aligned}$$

The final likelihood is calculated by substituting Equations (32) and (37) into Equation (16). As an example, Figure 4 compares the noiseless likelihood described in Section 3.3.1, the precise likelihood calculated via numerical integration of Equations (26) and (34), and our adopted approximate likelihood. We show how each likelihood evolves with respect to β , while holding other parameters fixed at their median posterior values from Section 3.4. The figure demonstrates that measurement uncertainties induce a slight shift in the likelihood, and that our approximation closely matches the true likelihood, accurately accounting for measurement uncertainties. By comparing these likelihoods, we further find that β is the parameter mostly affected by measurement uncertainties, while the f_{occ} parameters, $(\log M_0, \delta, \theta)$, generally remain much more robust against these uncertainties.

3.4. Sampling Results

We assume a flat prior for all our parameters, with bounds established to ensure the propriety of the prior (H. Tak et al. 2018): $\log M_0 \sim \mathcal{U}[5, 12]$, $\delta \sim \mathcal{U}[0, 50]$, $\theta \sim \mathcal{U}[0, 10]$, $\alpha \sim \mathcal{U}[15, 40]$, $\beta \sim \mathcal{U}[-3, 3]$, $\xi \sim \mathcal{U}[-15, 15]$, and $\sigma \sim \mathcal{U}[0, 10]$. We utilize `DynamicHMC.jl`,²⁸ which implements the Hamiltonian Monte Carlo algorithm (M. Betancourt 2017), to sample the posterior distribution. The sampling results are displayed in Figure 5, where the median values and

²⁸ <https://www.tamasapp.eu/DynamicHMC.jl/stable/>

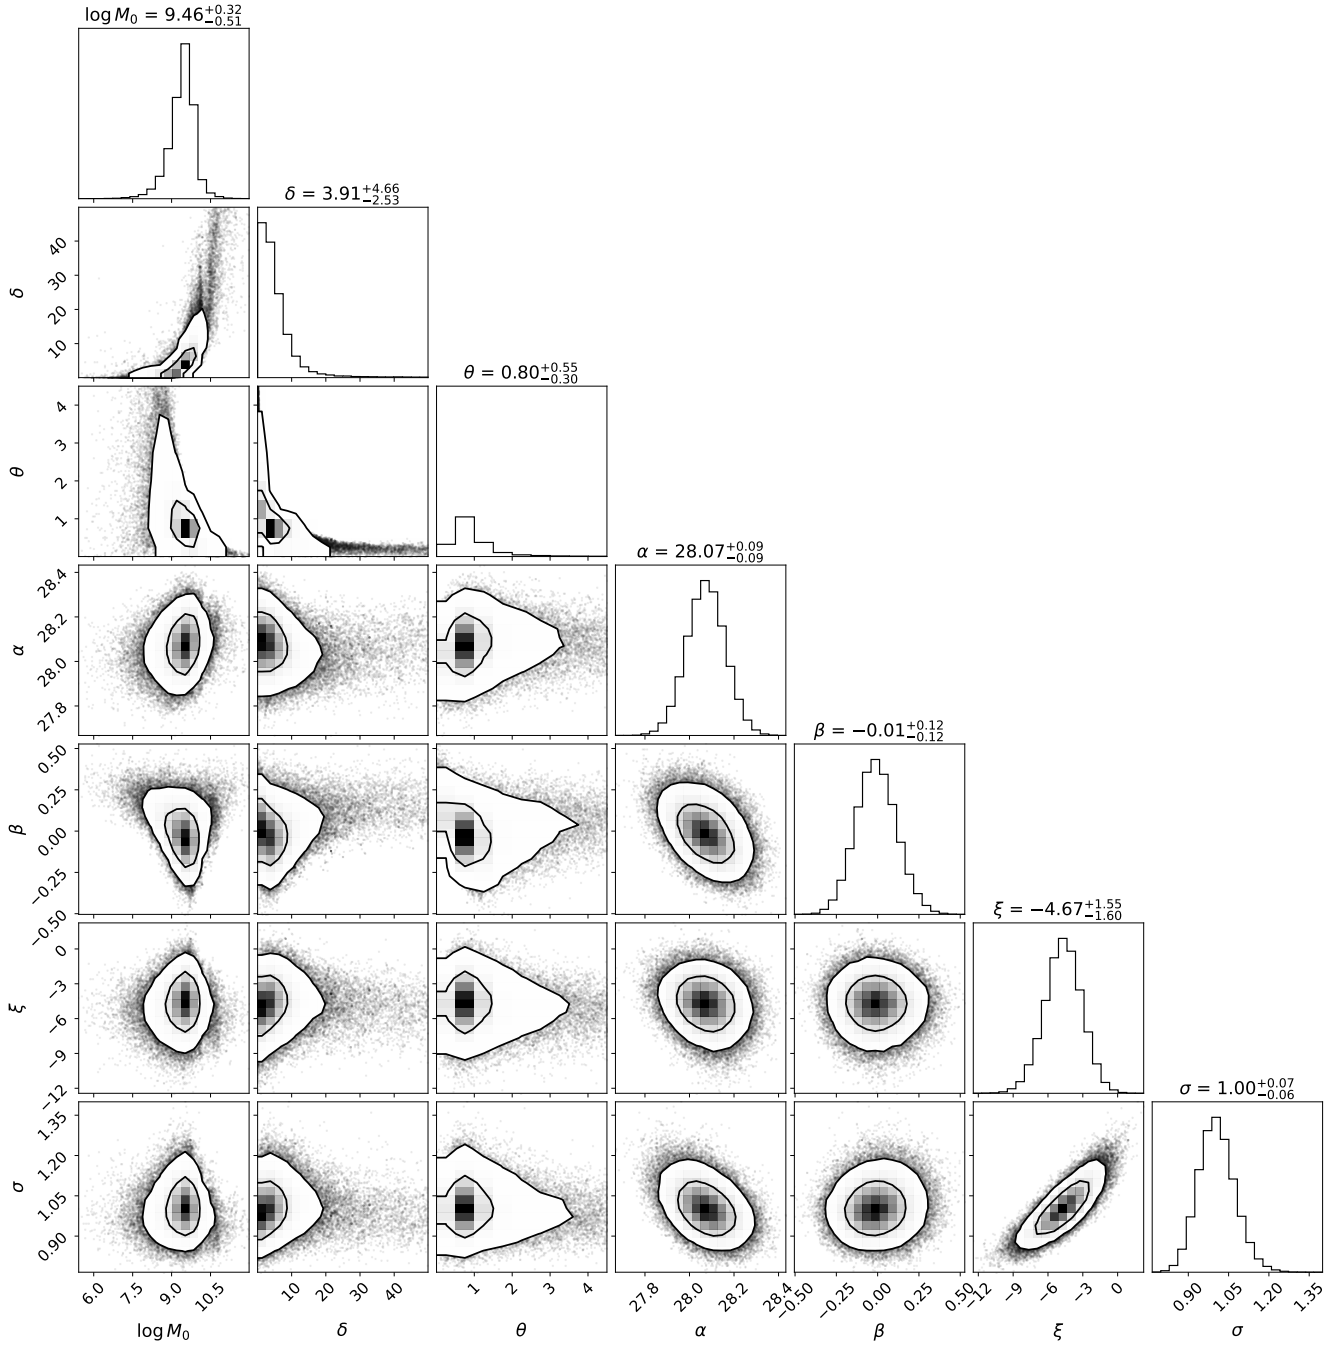


Figure 5. Sampling results. Contours represent 68% and 95% levels; grayscale pixels represent probabilities, and the rasterized points outside the 95% contours are individual sampling points. The first three parameters, ($\log M_0$, δ , θ), represent the dependence of f_{occ} on M_* (Equation (2)), as illustrated in the left panel of Figure 3. α represents the typical $\log \lambda$ of galaxies with $M_* = 10^{10} M_\odot$, and β is the slope of the mass dependence of the SARDF (Equation (9)). ξ and σ control the shape of $p(\lambda)$ (Equation (7)), as shown in the right panel of Figure 3. The fitted β is consistent with zero, indicating that $p(\lambda)$ has little dependence on M_* . The fitted ξ is $\ll 1$, suggesting that $p(\lambda)$ is highly right-skewed compared to normal distributions.

1σ uncertainties for all parameters are listed at the top of each column. The fitted posterior distributions are much narrower than the prior distributions for all these parameters, indicating that the data are sufficiently informative and the inference is insensitive to the choice of prior.

For replicability purposes, the resulting $p(\lambda)$ can be obtained by substituting these median values of our sampled parameters into Equation (7).²⁹ The results for the MBH occupation

fraction are presented in Table 2 and Figure 6, where the median, along with the 1σ , 2σ , and 3σ ranges of the posterior distribution are plotted as a function of M_* .

Due to the fivefold increase in sample size, the findings of this work are significantly more precise than, though still consistent with, previous f_{occ} constraints by M15 and E. Gallo & A. Sesana (2019), both of whom did not find an upper limit for f_{occ} in the dwarf galaxy regime. As a consistency check, in Appendix C, we verify that applying the model presented above to fit the same data used in M15 does not yield an upper bound on f_{occ} .

²⁹ Strictly speaking, the median $p(\lambda)$ is not exactly the same as this calculation, but we have verified that these two are nearly identical.




Figure 6. Local black hole occupation fraction constraints from Chandra ACIS imaging observations of 1606 galaxies within 50 Mpc. The black solid curve is the median of our sampling results; the dashed, dashed–dotted, dotted curves, and their corresponding shaded regions enclose the 1σ , 2σ , and 3σ ranges, respectively. The 1σ measurements in **M15** and **B25** are also shown for comparison. Our f_{occ} quickly drops down when M_* decreases below $10^{10} M_\odot$ and is below unity at a statistically significant level for low-mass galaxies.

Table 2

Black Hole Occupation Fraction as a Function of Host-galaxy Stellar Mass

$\log M_*$	0.13%	2.3%	16%	50%	84%	97.7%	99.87%
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
7.0	0.002	0.014	0.046	0.106	0.218	0.425	0.695
7.1	0.003	0.017	0.052	0.115	0.230	0.436	0.701
7.2	0.005	0.021	0.059	0.125	0.242	0.447	0.705
7.3	0.007	0.026	0.066	0.136	0.255	0.459	0.712
7.4	0.010	0.032	0.075	0.147	0.269	0.471	0.718
...

Note. Column (1): host-galaxy stellar mass. Columns (2)–(8) refer to the (2) 3σ lower bound; (3) 2σ lower bound; (4) 1σ lower bound; (5) median; (6) 1σ upper bound; (7) 2σ upper bound; and (8) 3σ upper bound of the sampled f_{occ} . (This table is available in its entirety in machine-readable form in the [online article](#).)

The key result of this study is that f_{occ} decreases sharply with decreasing M_* . The M_* value at which f_{occ} decreases to 50% is determined to be $\log M_* = 9.0_{-0.4}^{+0.2}$.

By convolving the inferred f_{occ} with the galaxy stellar mass function derived by S. P. Driver et al. (2022), we obtain the following values (1σ) or lower limits (2σ):

$$\begin{aligned}
 f_{\text{occ}} &> 93\% && \text{for } 10^{11} M_\odot \leq M_* \leq 10^{12} M_\odot, \\
 f_{\text{occ}} &= 66\%_{-7\%}^{+8\%} && \text{for } 10^9 M_\odot \leq M_* \leq 10^{10} M_\odot, \\
 f_{\text{occ}} &= 33\%_{-9\%}^{+13\%} && \text{for } 10^8 M_\odot \leq M_* \leq 10^9 M_\odot.
 \end{aligned}$$

The f_{occ} uncertainty is asymmetric, and moderate occupation (i.e., $f_{\text{occ}} \approx 50\%$) is still allowed down to $M_* = 10^8 M_\odot$ at the 2σ level.

The median of our sampling results for the probability distribution $p(\lambda)$ of the SARDF is shown as the solid red curve in Figure 7 for galaxies with $M_* = 10^{10} M_\odot$. The figure also compares the inferred $p(\lambda)$ with results from previous works,

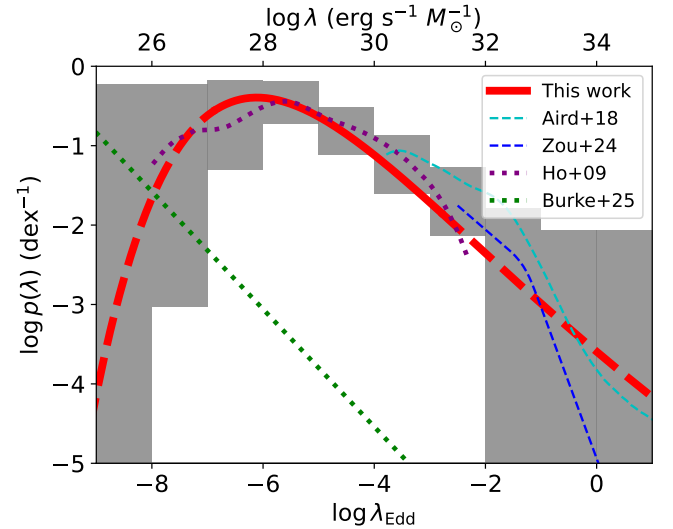


Figure 7. The probability distribution, $p(\lambda)$, of the SARDF for galaxies with $M_* = 10^{10} M_\odot$ is shown as a function of $\log \lambda$ (top axis) and $\log \lambda_{\text{Edd}} = \log \lambda - 34$ (bottom axis). Here, λ is converted to λ_{Edd} by assuming $k_{\text{bol}} = 25$ and $M_{\text{BH}} = 0.002 M_*$ (see Section 4.4 for a more in-depth discussion). The gray shaded regions bracket the range of allowed $p(\lambda)$ as inferred directly from the data (see Section 4.2), based on a subsample of 457 nearby objects with $M_* > 10^{10} M_\odot$. The red curve represents our inferred parametric $p(\lambda)$, which is formally valid between $27 \lesssim \log \lambda \lesssim 31$ (shown as a solid line). The probability distributions derived by J. Aird et al. (2018) and F. Zou et al. (2024b) for AGNs are shown as dashed curves for comparison. The probability distributions obtained by L. C. Ho (2009) and B25 based on different samples of nearby galaxies are shown as dotted lines.

which will be discussed in Section 4.2. The main properties of $p(\lambda)$ can be summarized as follows: (i) the distribution exhibits a peak near $\log \lambda \approx 28$; (ii) above this value, $p(\lambda)$ is well approximated by a declining power-law function of $\log \lambda$, up to approximately 31.5, beyond which the bona fide AGN regime begins; and (iii) because the best-fit value of β in

Equation (9) is consistent with zero ($\beta = -0.01_{-0.12}^{+0.12}$), $p(\lambda)$ is nearly independent of M_* .

A peak in the distribution at $\log \lambda \approx 28$ is entirely consistent with the data. Among all the massive galaxies with $M_* > 10^{10} M_\odot$ in our sample, 42% have detected X-ray nuclei with $\log \lambda > 28$ (the actual fraction may be even higher when accounting for undetected sources), implying that $\log \lambda = 28$ approximately corresponds to the median λ of massive galaxies.

The main results of this work can be intuitively understood as follows: the SARDF and its possible dependence on M_* are primarily determined by the massive galaxies, which have significantly higher X-ray detection rates compared to low-mass galaxies. The dependence of the SARDF on M_* is encapsulated by Equation (9), which assumes a linear relationship between μ (a proxy for $\log \lambda$) and $\log M_*$. This behavior is assumed to hold even in the low-mass regime. Our sampling yields $\beta \approx 0$, implying that the SARDF of low-mass galaxies is similar to that of massive galaxies. It follows that if most or all low-mass galaxies have central MBHs (i.e., $f_{\text{occ}} \approx 1$), the model would predict a significantly higher X-ray detection fraction than is measured, even after accounting for the fact that, at a fixed Eddington ratio, low-mass galaxies are less luminous. Our model reflects these processes self-consistently. In Appendix D, we demonstrate that our model can recover f_{occ} in both simulated low- and high- f_{occ} scenarios.

We emphasize that the inferred $p(\lambda)$ is only formally valid within the range $27 \lesssim \log \lambda \lesssim 31$, and extrapolations beyond this range may not be reliable. There is evidence against extrapolating $p(\lambda)$ too far in either direction: (i) at the high- λ end, it is known that the AGN $p(\lambda)$ follows a broken power law (e.g., A. K. Weigel et al. 2017; J. Aird et al. 2018; F. Zou et al. 2024b), and the break is not included in our model; (ii) at the low- λ end, our $p(\lambda)$ is in tension with observations of the Milky Way MBH Sgr A*, whose X-ray luminosity is $L_X = 2.4 \times 10^{33} \text{ erg s}^{-1}$ (F. K. Baganoff et al. 2003). For a host stellar mass of $M_* = 6.08 \times 10^{10} M_\odot$ (T. C. Licquia & J. A. Newman 2015), this yields $\log \lambda = 22.6$, which is over 5 orders of magnitude below the inferred peak in $p(\lambda)$. Including the Sgr A* data point in our sample significantly elevates $p(\lambda)$ (compared to the distribution shown in Figure 7) below $\log \lambda \lesssim 26$, although the range between $27 \lesssim \log \lambda \lesssim 31$ remains virtually unchanged and still exhibits a peak at $\log \lambda \approx 28$. However, we find that with the inclusion of Sgr A*, assuming that $p(\lambda)$ follows a Box-Cox distribution would be falsified (see Equation (38) and Section 4.1.1). This implies that the Box-Cox distribution family is not sufficiently flexible to model the full $p(\lambda)$ distribution (all the way down to $\log \lambda \lesssim 23$ and up to $\log \lambda \gtrsim 33$). We further note that this limitation does not undermine our f_{occ} inference, as it is only sensitive to $p(\lambda)$ in the range $28 \lesssim \log \lambda \lesssim 31$.

We anticipate that the occupation fraction of low-mass galaxies would be higher if the SARDF of low-mass galaxies had a peak at lower λ values. To precisely determine where the peak in the SARDF would need to be to achieve full occupation in the dwarf galaxy regime, we consider two mass bins: $\log M_* = 8 \pm 0.5$ and 9 ± 0.5 . We fix the SARDF shape parameters (σ, ξ) to the values shown in Figure 5 and assume $f_{\text{occ}} = 1$. Maximizing the likelihood over μ (which is a proxy for the distribution peak value) returns $\mu = 26.0$ and

27.4 for the $\log M_* = 8 \pm 0.5$ and 9 ± 0.5 bins, respectively. This means that, to reproduce the low X-ray detection rates in the case of full occupation, the peaks of the SARDFs of low-mass galaxies would need to shift toward lower values (specifically by 2.1 and 0.7 dex for the $\log M_* = 8 \pm 0.5$ and 9 ± 0.5 bins, respectively). However, we emphasize that there is currently no evidence to suggest that MBHs in low-mass galaxies are more likely to have lower λ compared to those in massive galaxies. This work provides an observational benchmark for the SARDF of massive galaxies in the low- λ regime (see Section 4.2). High-resolution simulations of individual galaxies yield predictions that are broadly consistent with our estimates (e.g., F. Yuan et al. 2018). However, observational constraints of the SARDF for low-mass galaxies in the low- λ regime are currently unavailable; additional uncertainties driven by stellar feedback, seed mass, and nuclear star clusters (e.g., C. Partmann et al. 2025; J. Petersson et al. 2025) make theoretical modeling at the low-mass end particularly challenging.

Indirect clues about the accretion properties of low-mass galaxies come from AGNs. Although simulations still face challenges and lack consensus (e.g., R. S. Sharma et al. 2022; R. S. Beckmann et al. 2023), observational studies show that the AGN SARDF and AGN fraction of low-mass galaxies with $10^8 M_\odot \lesssim M_* \lesssim 10^{10} M_\odot$ and $z \lesssim 0.6$ are similar to those of massive galaxies—if AGNs are identified based on λ rather than L_X (K. L. Birchall et al. 2020, 2022; F. Zou et al. 2023). While we acknowledge that any statistical inference model involves assumptions, we argue that the model presented in this work is both sufficiently flexible and physically reasonable.

Our results are in stark contrast to those obtained by B25, who report a very high f_{occ} (close to 100% down to the lowest-mass bin). We will examine the reasons for this disagreement in detail in Section 4.1.2, and provide independent evidence for a plateau in the SARDF (ERDF) at low λ (λ_{Edd}) in Section 4.2 below.

4. Discussion

This work presents new constraints on the local black hole occupation fraction, using Chandra imaging data for approximately 1600 galaxies within a distance of 50 Mpc. We employ a Bayesian model to simultaneously constrain f_{occ} and the probability distribution $p(\lambda)$ of the SARDF. We find that $p(\lambda)$ shows little to no evolution with M_* and exhibits a peak around $10^{28} \text{ erg s}^{-1} M_\odot^{-1}$ (Figure 7). Simultaneously, f_{occ} drops below 50% for host-galaxy stellar masses $M_* \lesssim 10^9 M_\odot$ (Figure 6).

A key assumption throughout this work is that $p(\lambda)$ follows a Box-Cox distribution (Figure 3). This distribution is, in principle, more flexible than the normal distribution assumed by M15 or the broken power-law distribution (based on A. K. Weigel et al. 2017) assumed by B25, as it can reasonably approximate both functional shapes. In the following sections, we aim to derive model-independent constraints for the SARDF and assess the viability of the aforementioned assumptions, as well as their implications for the inferred occupation fraction, in light of these model-independent constraints. It is important to clarify that $p(\lambda)$ should not be confused with the distribution of bona fide AGNs (i.e., systems with $\log \lambda \gtrsim 32$, which corresponds to $\lambda_{\text{Edd}} \gtrsim 0.01$), which is indeed well described by a double

power law (e.g., A. Bongiorno et al. 2016; J. Aird et al. 2018; G. Yang et al. 2018; F. Zou et al. 2024b). Instead, our focus is on constraining the low-luminosity (or quiescent) range below $\log \lambda \lesssim 31$. These systems may exhibit different accretion physics from AGNs—as the accretion flow likely transitions from a geometrically thin disk to an optically thin, radiatively inefficient flow (e.g., F. Yuan & R. Narayan 2014 and references therein). We will further discuss whether the inferred $p(\lambda)$ for low-luminosity galactic nuclei accurately extrapolates into the AGN regime.

4.1. The Accretion-rate Distribution of Local Massive Black Holes

4.1.1. Box–Cox versus Normal Distribution

First, we verify whether the data are statistically consistent with a Box–Cox distribution. By applying the inverse transformation of Equation (5), our assumption becomes equivalent to

$$\frac{\mu}{\xi\sigma} \left[\left(\frac{\log \lambda}{\mu} \right)^\xi - 1 \right] \sim \mathcal{N}(0, 1). \quad (38)$$

We then perform a normality test for censored data (S. P. Millard 2013) and obtain a high p -value of 0.2, suggesting that a Box–Cox distribution accurately represents $p(\lambda)$.

Next, we evaluate whether a normally distributed $p(\lambda)$, as assumed by M15, is consistent with the data. The Box–Cox distribution simplifies to a normal distribution when $\xi = 1$. However, our best-fit value is $\xi = -4.67_{-1.60}^{+1.55}$, which differs from one at a highly statistically significant level. To confirm this discrepancy, we fix ξ to one and refit our model. The normality test in Equation (38) yields a low p -value of 9×10^{-4} , thus refuting the normal distribution assumption. This outcome arises because the data distribution is right-skewed with a heavy tail toward high λ , characteristics that normal distributions cannot adequately capture. Moreover, if we adopt the noiseless likelihood in Section 3.3.1, the inferred $\log M_* - \log \lambda$ slope under the normal distribution assumption would be $-0.16_{-0.12}^{+0.12}$, corresponding to a $\log M_* - \log L_X$ slope of $+0.84$. This value is in good agreement with the slope reported by B. Miller et al. (2012) and M15, who argued for downsizing in black hole accretion—a scenario where lower-mass galaxies exhibit relatively higher λ values compared to more massive ones. However, after we account for the measurement uncertainties (Section 3.3.2), the $\log M_* - \log \lambda$ slope becomes $-0.04_{-0.13}^{+0.12}$ under the normal distribution assumption and $-0.01_{-0.12}^{+0.12}$ under our Box–Cox distribution assumption. As Figure 4 shows, the slope would be slightly underestimated if not accounting for the measurement uncertainties. Therefore, our analysis suggests that previous downsizing results likely no longer hold.

4.1.2. Box–Cox versus Power-law Distribution

We specifically address the apparent disagreement with the recent findings of B25, who report a very high occupation fraction across the host stellar mass spectrum. It is important to note that B25 use a consistent statistical model to explore various diagnostics of MBH activity, such as nuclear X-ray and radio emissions, as well as optical variability, consistently finding a high occupation fraction across the mass range

($f_{\text{occ}} > 90\%$ at $M_* \gtrsim 10^8 M_\odot$, at the 2σ level). In this discussion, we concentrate on the X-ray sample and identify their assumed ERDF as the primary source of the discrepancy.

Although both our study and B25 begin with the 50MGC to draw galaxy samples, B25 utilizes several X-ray catalogs, including ROSAT, Chandra, XMM-Newton, and eROSITA. In contrast, we restrict our selection to CSC 2.1 sources with ACIS data. An additional difference is the chosen matching radius between CSC 2.1 (or other catalogs) and 50MGC, as well as the aperture over which XRB contamination is evaluated. We emphasize that these differences do not influence the conclusions presented below. In fact, we verified that when we apply a similar Bayesian inference method as B25 to the Chandra sample described in Section 2, we recover very high occupation fractions down to low galaxy masses, in agreement with B25.

In this work, we assume the same³⁰ functional expression for the MBH occupation fraction probability as B25, but make different assumptions for the underlying $p(\lambda)$. For clarity, we recall that we make use of the SARDF, which is a function of $\lambda = L_X/M_*$, while B25 define the ERDF as a function of $\lambda_{\text{Edd}} = k_{\text{bol}} L_X/L_{\text{Edd}}$. A rough conversion between the SARDF and the ERDF is given by $\log \lambda \approx \log \lambda_{\text{Edd}} + 34$ (this is formally correct for an X-ray bolometric correction $k_{\text{bol}} = 25$ and black hole to host-galaxy stellar mass ratio $M_{\text{BH}}/M_* = 0.002$). The ERDF assumed by B25 is described by a smoothly broken power-law model, whereby the probability of hosting a low- λ_{Edd} MBH increases monotonically with decreasing λ_{Edd} , down to a minimum value $\log \lambda_{\text{Edd, min}}$, which is inferred to be ≈ -8.7 (see their Figure 19). This assumption originates from the fact that the probability distribution of bona fide AGNs is thought to follow a double power law (e.g., A. K. Weigel et al. 2017; J. Aird et al. 2018; F. Zou et al. 2024b). Strong evidence for this comes from the measured AGN luminosity function (Y. Ueda et al. 2014). However, there is no empirical evidence to suggest that the double power law can be extended to quiescent MBHs, i.e., down to $\log \lambda \approx 25$, or $\log \lambda_{\text{Edd}} \approx -9$.

The assumption of a Box–Cox distribution for the SARDF can indeed model a power-law-like tail across a wide range of λ values (Figure 3), although it implicitly assumes the presence of a peak in the distribution. Our prior ($\alpha \sim \mathcal{U}[15, 40]$) allows the peak of $p(\lambda)$ to be as low as approximately 15 (in $\log \lambda$), which is more than 10 orders of magnitude below the inferred peak value of $\log \lambda \approx 28$. This effectively permits an ERDF that is qualitatively similar to that assumed by B25. Appendix D directly demonstrates that if the underlying $p(\lambda)$ is indeed a broken power law, our model should be able to recover it well.

It is also noteworthy that the ERDF in B25 is assumed to be independent of black hole mass. This assumption implies that all galaxies must share the same distribution of $\log L_X - 1.61 \log M_*$ (or equivalently, that $\lambda \propto M_*^{0.61}$), where the 1.61 slope is derived from the chosen $M_{\text{BH}}-M_*$ scaling relation (Equation (3) in B25). In contrast, we do not fix the scaling of $\log M_*$ with $\log \lambda$: in our model, the slope is a free parameter with a fitted value of $-0.01_{-0.12}^{+0.12}$, compared to the 0.61 assumed by B25. A steeper slope implies that lower-mass galaxies tend to have lower L_X and are thus less likely to be detected, which in turn favors higher occupation.

³⁰ Modulo the additional δ factor in the exponentiation term of Equation (2).

We wish to emphasize that, up to this point, we have not provided compelling reasons to prefer one distribution function over the other. In the following section, we derive model-independent constraints on the SARDF over a range of λ (or λ_{Edd}) that is significantly below the levels typically probed by AGN observations. We also demonstrate the presence of a plateau in the SARDF.

4.2. Bracketing the Specific Accretion-rate Distribution Function of Quiescent Massive Black Holes

Robust upper and lower bounds to $p(\lambda)$ can be derived directly from the Chandra data without relying on Bayesian inference. To do so, we only consider galaxies with $M_* > 10^{10} M_\odot$ because their occupation fraction is widely accepted to be 100%, and, by definition, $p(\lambda)$ can only be defined for sources with central MBHs. There are 457 sources in this mass bin, and 245 (54%) have compact X-ray nuclei. For these galaxies, a meaningful range of $p(\lambda)$ values can be determined by examining two extreme scenarios. Consider a specific accretion ratio bin $\log \lambda$, which includes values between $\log \lambda_1$ and $\log \lambda_2$. The lowest possible $p(\lambda)$ for that bin occurs when none of the sources with upper limits falls within the range between $\log \lambda_1$ and $\log \lambda_2$. Conversely, the highest possible $p(\lambda)$ occurs when all sources with upper limits higher than $\log \lambda_1$ actually have $\log \lambda$ values falling between $\log \lambda_1$ and $\log \lambda_2$.

Mathematically, we proceed by labeling (1, 2, ...) bins in ascending order of $\log \lambda$ and denote with N_i^{cens} the number of censored $\log \lambda$ points within the i th bin. The minimum number of sources in the i th bin is given by

$$N_{\min,i} = \sum_k (1 - P_{\text{XRB},k}) \mathcal{I}(\log \lambda_k \in \{\text{ith bin}\}), \quad (39)$$

where k runs over all the sources with $M_* > 10^{10} M_\odot$, and $\mathcal{I}(\log \lambda_k \in \{\text{ith bin}\}) = 1$ (0) if $\log \lambda_k$ falls within (outside of) the i th bin. The maximum number of sources is

$$N_{\max,i} = N_{\min,i} + \sum_{j \geq i} [N_j^{\text{cens}} + P_{\text{XRB},k} \mathcal{I}(\log \lambda_k \in \{\text{jth bin}\})]. \quad (40)$$

The lower (upper) bound of $p(\lambda) \Delta \log \lambda$ is set to the 2σ lower (upper) bound of a binomial proportion confidence interval (E. B. Wilson 1927) with the total number of trials set equal to the number of sources with $M_* > 10^{10} M_\odot$, and the number of successes being $N_{\min,i}$ ($N_{\max,i}$).

The shaded gray areas in Figure 7 illustrate the range of allowed SARDFs. Generalizing the above reasoning, $N_{\max,i} = 0$ if the bin covers a range that exceeds the highest $\log \lambda$ measured in the sample. $N_{\max,i}$ equals the total number of censored data points in the sample if the bin covers a range lower than the smallest $\log \lambda$ in the sample. Note that there are almost no data points, whether censored or uncensored, below $\log \lambda < 26$ and above $\log \lambda > 33$. In these cases, $N_{\min,i} = 0$. Furthermore, not all $p(\lambda)$ curves within the viable gray range are permissible, as they must integrate to unity by definition. Critically, this analysis demonstrates that the SARDF ceases increasing with decreasing $\log \lambda$ below $\log \lambda \lesssim 28$, or $\log \lambda_{\text{Edd}} \lesssim -6$. The inferred SARDF, shown by the red curve in Figure 7, is comfortably within the range of allowed values.

Figure 7 also includes a set of published SARDF/ERDF probability distributions, including the ERDF derived by B25

(dotted green curve), as well as a representative ERDF based on the results³¹ of L. C. Ho (2009, dotted purple curve). The B25 and L. C. Ho (2009) ERDFs are shifted by +33.4 and +33.8, respectively, to allow for a uniform representation in units of λ .³² For comparison, we also include the SARDFs inferred for AGN using different samples (J. Aird et al. 2018; F. Zou et al. 2024b), which are in qualitative agreement with the broken power law derived by A. K. Weigel et al. (2017).

A visual comparison of the different probability distributions shows that the SARDF derived in this work is in good qualitative agreement with both the results of L. C. Ho (2009, based on a sample with comparable luminosities) and the AGN findings from various groups, whereas it is inconsistent with the probability distribution assumed by B25. Aside from the chosen functional shape, the parameterization of B25 implies a very low normalization for the ERDF probability distribution above $\log \lambda_{\text{Edd}} > \log \lambda_{\text{Edd},\min} = -8.7$ (equivalent to $\log \lambda > 24.7$). Above this threshold, the normalization is set by the mass-independent active fraction $f_a(\lambda_{\text{Edd}} > \lambda_{\text{Edd},\min}) = \int_{\log \lambda_{\text{Edd},\min}}^{+\infty} p(\lambda_{\text{Edd}}) d \log \lambda_{\text{Edd}} = 0.14$. The resulting ERDF probability distribution, however, is inconsistent with the observed active fraction of massive galaxies. For instance, the bracketing method described in Equations (39) and (40) yields an active fraction for $\log \lambda > 29$ between 13% and 28% for $M_* > 10^{10} M_\odot$, whereas integrating B25's ERDF over the same range yields less than 1%. In summary, our comparison with the findings of B25 suggests that their chosen ERDF parameterization, combined with the implicit assumption of a steep L_X - M_* relation, leads to the inference of a large population of highly underluminous MBHs.

4.3. Dependence on Galactic Morphology and Star-forming Properties

We further derive f_{occ} and $p(\lambda)$ separately for different types of galaxies. Using the 50MGC morphological classifications (D. Ohlson et al. 2024), we divide our sample into 618 early-type galaxies (with 186 compact X-ray detections) and 988 late-type galaxies (with 202 compact X-ray detections). We repeat the analysis as described in Section 3 for both galaxy types. The results, presented in the top panels of Figure 8, are qualitatively similar to our original findings, with both early- and late-type galaxies showing a decrease in f_{occ} toward lower M_* . Late-type galaxies tentatively exhibit higher f_{occ} at $M_* \lesssim 10^9 M_\odot$ along with elevated λ , indicating that the corresponding $p(\lambda)$ is higher at $\lambda \gtrsim 10^{29} \text{ erg s}^{-1} M_\odot^{-1}$. However, within 2σ uncertainties, early- and late-type galaxies remain statistically consistent.

³¹ L. C. Ho (2009) assessed the Eddington ratio, λ_{Edd} , for a sample of 175 nearby galaxies with Chandra data but did not explicitly calculate the sample's ERDF. We derive the ERDF based on their data set. In summary, we first employ the Nelson-Aalen estimator (O. Aalen 1978) to compute the cumulative hazard function, $H(\log \lambda_{\text{Edd}})$, for their λ_{Edd} distribution. Given that the resultant $H(\log \lambda_{\text{Edd}})$ is nondifferentiable, we apply a spline fit to smooth it and subsequently convert it to $p(\log \lambda_{\text{Edd}})$ using the relation $p(\log \lambda_{\text{Edd}}) = e^{-H} \frac{dH}{d \log \lambda_{\text{Edd}}}$. This approach involves significant uncertainties, particularly in the $\log \lambda_{\text{Edd}}$ intervals that are inadequately represented by the data. Consequently, we consider it only a coarse approximation of their ERDF.

³² The different shifts are applied in order to ensure self-consistency between the definition of λ_{Edd} in the original works and the definition of λ adopted in this work. Specifically, B25 adopt $k_{\text{bol}} = 10$ and $M_{\text{BH}}/M_* = 8 \times 10^{-5}$ at $M_* = 10^{10} M_\odot$, whereas L. C. Ho (2009) adopt $k_{\text{bol}} = 15.8$ and derive the black hole mass values using the bulge stellar velocity dispersion relation.

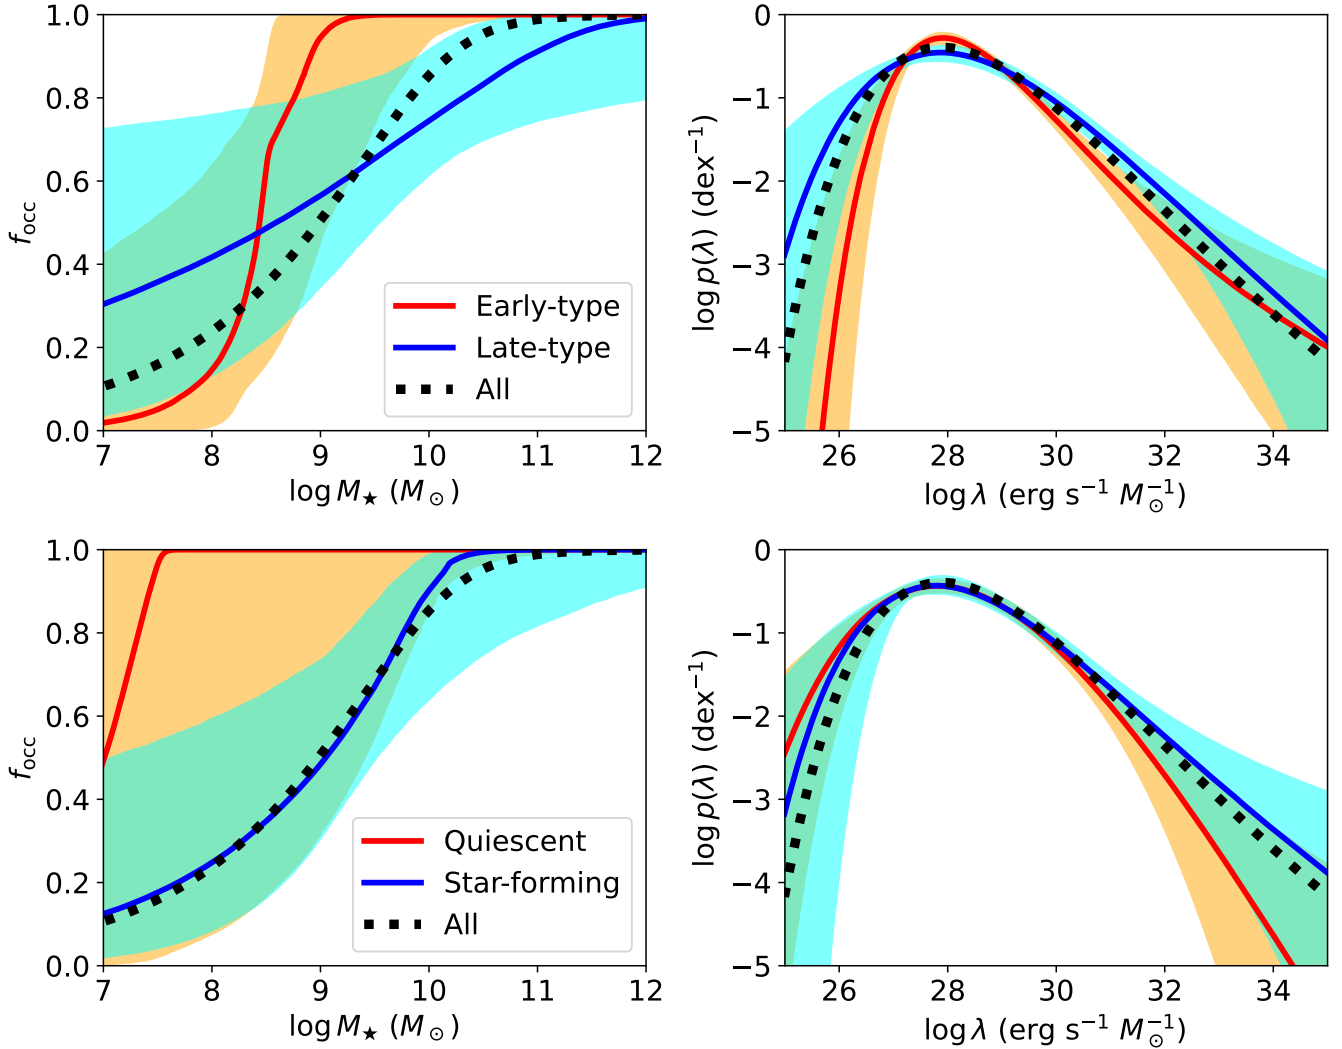


Figure 8. Sampled f_{occ} (left) and $p(\lambda)$ (right) for different types of galaxies separately. In the top panels, the red and blue solid curves represent the median of the posterior distributions for early-type and late-type galaxies, respectively, while in the bottom panels, they represent quiescent and star-forming galaxies, respectively. The black dotted curves correspond to all galaxies, referenced from Figures 6 and 7. The shaded regions indicate the 2σ uncertainty levels.

Similarly, we examine whether star-forming and quiescent galaxies have similar f_{occ} and $p(\lambda)$. A galaxy is classified as quiescent if its SFR is more than 1 dex below the star-forming main sequence described in J. Leja et al. (2022); otherwise, it is classified as star forming. Our sample comprises 625 quiescent galaxies (with 165 compact X-ray detections) and 750 star-forming galaxies (with 150 compact X-ray detections). The corresponding results are shown in the bottom panels of Figure 8. Star-forming galaxies exhibit a decreasing f_{occ} as M_\star decreases, but the constraints for quiescent galaxies are broad and consistent with $f_{\text{occ}} = 1$. Their $p(\lambda)$ remains largely consistent, although the quiescent $p(\lambda)$ appears suppressed at $\lambda \gtrsim 10^{31} \text{ erg s}^{-1} M_\odot^{-1}$. We have verified that the early-type, late-type, quiescent, and star-forming subsamples with compact nuclear X-ray detections generally have small P_{XRB} values (Section 2.1), with median values below 6%, indicating that XRB contamination is minimal.

4.4. Eddington Ratio and Mass Accretion-rate Conversions

Our analysis utilizes $\lambda = L_X/M_\star$, rather than $\lambda_{\text{Edd}} = L_{\text{bol}}/M_{\text{BH}}$. This approach has the distinct advantage of avoiding the uncertainties associated with converting (L_X, M_\star) to

$(L_{\text{bol}}, M_{\text{BH}})$. In principle, the SARDF can be converted to the ERDF using the bolometric correction factor $k_{\text{bol}} = L_{\text{bol}}/L_X$, or even to a mass accretion-rate (\dot{M}) distribution via the radiative efficiency $\epsilon = L_{\text{bol}}/(\dot{M}c^2)$. Figure 7 adopts a nominal value of $k_{\text{bol}} = 25$ when plotting the ERDF, which is typical for standard thin-disk accretion in AGNs. The characteristic AGN radiative efficiency is usually assumed to be $\epsilon = 0.1$ (e.g., S. W. Davis & A. Laor 2011; W. N. Brandt & D. M. Alexander 2015). However, both k_{bol} and ϵ may vary for low- λ_{Edd} black holes, since the accretion mode is expected to transition to radiatively inefficient, hot accretion flows as λ_{Edd} decreases (e.g., R. Narayan & I. Yi 1995; R. Mahadevan 1997; F. Yuan & R. Narayan 2014 and references therein). Importantly, variations in k_{bol} and ϵ do not affect our determination of the SARDF and f_{occ} , as our results rely exclusively on L_X rather than L_{bol} .

With regard to k_{bol} , I. E. López et al. (2024) analyze the X-ray to IR SEDs of MBHs with $-8 \lesssim \log \lambda_{\text{Edd}} \lesssim -3$, similar to the λ_{Edd} range of our sources (see Figure 7). Using physical models of advection-dominated accretion flows (ADAFs) and truncated accretion disks, they find that k_{bol} converges to a value of nine in the low-luminosity regime. Although this is lower than the typical AGN k_{bol} , it remains consistent with

extrapolations of the AGN $k_{\text{bol}}-L_{\text{bol}}$ relation presented in F. Duras et al. (2020).³³ Changing k_{bol} from 25 to 9 shifts the ERDF in Figure 7 to the left by approximately 0.4 dex. Alternatively, one can adopt an L_{bol} -dependent k_{bol} (e.g., Equation (4) in I. E. López et al. 2024), which further modifies the shape of the AGN ERDF, since the dependence of k_{bol} on L_{bol} becomes significant primarily in the AGN regime. Since our SARDF is already determined, converting it to the ERDF using different k_{bol} values and the $M_{\text{BH}}-M_*$ relation is straightforward for interested readers. Overall, reasonable variations in k_{bol} cause a shift of $\lesssim 0.5$ dex, while adopting different $M_{\text{BH}}-M_*$ relations (e.g., F. Shankar et al. 2020b) can result in a shift of up to $\lesssim 1$ dex.

Regarding ϵ , much larger deviations from the canonical AGN value of 0.1 are expected. We denote λ_M as the ratio between \dot{M} and the Eddington mass accretion rate, $\dot{M}_{\text{Edd}} = 10L_{\text{Edd}}/c^2$, i.e., $\lambda_M = \dot{M}/\dot{M}_{\text{Edd}} = (0.1/\epsilon)\lambda_{\text{Edd}}$. F.-G. Xie & F. Yuan (2012) show that ϵ quickly decreases with decreasing λ_M when $\lambda_M \lesssim 0.01$, and the expected ϵ may be as low as $10^{-4}-10^{-3}$ when $\lambda_M = 10^{-6}$. Therefore, λ_M is generally larger than λ_{Edd} by a few orders of magnitude when $\lambda_{\text{Edd}} \ll 1$. The exact value of ϵ as a function of λ_M depends on the chosen ADAF model parameters, especially δ_{ADAF} , the fraction of viscously dissipated energy that heats electrons (see, e.g., F. Yuan & R. Narayan 2014). Varying δ_{ADAF} within the nominal range of 0.1–0.5 can introduce up to a 1 dex uncertainty in ϵ (F.-G. Xie & F. Yuan 2012). Adopting ϵ as a function of λ_M from Table 1 of F.-G. Xie & F. Yuan (2012) and $\delta_{\text{ADAF}} = 0.1$ (0.5), we find that $\log \lambda_{\text{Edd}} = -6$ corresponds to $\log \lambda_M = -4.6$ (-5.2) with $\epsilon = 0.03$ (0.01). Nevertheless, although ϵ may vary significantly with λ_M , its changes have no impact on our SARDF or ERDF because λ_{Edd} is instead based on k_{bol} , which has a limited variation for low- λ MBHs.

4.5. Implications for the Local Black Hole Mass Function

The distribution of black hole masses, particularly those in $z=0$ galactic nuclei, is key to understanding MBH growth over cosmic history. For black hole masses below $\approx 10^7 M_\odot$, this distribution can elucidate the assembly of MBH progenitors (at $z \gtrsim 15$; T. E. Woods et al. 2019), affect tidal disruption event rates (N. C. Stone & B. D. Metzger 2016), and distinguish between models for star formation suppression in dwarf galaxies (J. Silk 2017). LISA (P. Amaro-Seoane et al. 2017) will probe black holes from $M_{\text{BH}} = 10^4$ to $10^7 M_\odot$, though current knowledge is limited (E. Gallo & A. Sesana 2019; J. E. Greene et al. 2020).

The black hole mass function (BHMF), $\phi(M_{\text{BH}})$, is defined as the number of MBHs per unit comoving volume per unit $\log M_{\text{BH}}$. The integral of the BHMF over redshift yields the total number of MBHs in the observable Universe. In the local Universe, MBH masses are typically estimated through empirical scaling relations with a separate observational quantity, such as M_* , whose distribution is known. Accounting for the M_* -dependent MBH occupation fraction, the BHMF

takes the form

$$\phi(M_{\text{BH}}) = \int f_{\text{occ}}(M_*) \phi(M_*) \frac{1}{\sqrt{2\pi\sigma_M^2}} \times \exp\left(-\frac{[\log M_{\text{BH}} - \mathbb{E}(\log M_{\text{BH}}|M_*)]^2}{2\sigma_M^2}\right) d\log M_*, \quad (41)$$

where $\phi(M_*)$ is the galaxy stellar mass function, $\mathbb{E}(\log M_{\text{BH}}|M_*)$ is the chosen $M_{\text{BH}}-M_*$ scaling relation, and σ_M is its scatter.

We adopt the stellar mass function derived by S. P. Driver et al. (2022) and the same scaling relation as presented in E. Gallo & A. Sesana (2019): $\mathbb{E}(\log M_{\text{BH}}|M_*) = 8.13 + 1.72(\log M_* - 11)$ with $\sigma_M = 0.61$ dex. This scaling relation is based on the combined AGN and inactive-galaxy sample from A. E. Reines & M. Volonteri (2015). The AGN-based scaling relation and the inactive-galaxy one are found to be different (e.g., A. E. Reines & M. Volonteri 2015), and we adopt the combined one to mitigate the underlying uncertainties. The resulting BHMF is shown in Figure 9. The gray shaded region represents the statistical uncertainty induced by the f_{occ} inference. The effect of a declining occupation fraction flattens the BHMF at $M_{\text{BH}} \lesssim 10^6 M_\odot$ (e.g., E. Gallo & A. Sesana 2019; H. Cho & J.-H. Woo 2024). The median BHMF within $10^4 M_\odot < M_{\text{BH}} < 10^{10} M_\odot$ can be well approximated by the following polynomial fit:

$$\log \phi = -0.0420 - 0.941 \log M_{\text{BH}} + 0.0592(\log M_{\text{BH}})^2 + 0.0142(\log M_{\text{BH}})^3 - 0.00156(\log M_{\text{BH}})^4. \quad (42)$$

The right panel of Figure 9 illustrates that the total mass budget of MBHs is dominated by supermassive black holes with $M_{\text{BH}} \gtrsim 10^8 M_\odot$. The contribution from MBHs with $M_{\text{BH}} \lesssim 10^6 M_\odot$ is expected to be several orders of magnitude lower, regardless of f_{occ} . We emphasize that the primary uncertainty in constraining the low-mass end of the BHMF arises from the mass scaling relation. The shaded cyan area at high masses in Figure 9 illustrates the extent of this uncertainty, which exceeds 1 dex, as noted by F. Shankar et al. (2016, 2020a, 2020b).

An important caveat is that the BHMF presented in this work refers only to central MBHs. It has been argued that a large fraction (potentially over 50%) of MBHs in dwarf galaxies could be off-nuclear (e.g., K. T. Voggel et al. 2019; J. E. Greene et al. 2020; M. Mezcua & H. Domínguez Sánchez 2020; J. M. Bellovary et al. 2021; A. Sacchi et al. 2024). In fact, even in massive galaxies, a significant portion (potentially up to 40%) of the total M_{BH} budget may reside in off-nuclear MBHs (e.g., A. Kulier et al. 2015; A. Ricarte et al. 2021b; R. Pechetti et al. 2022; M. Häberle et al. 2024; F. Zou et al. 2024a). Large volume cosmological simulations also predict the existence of a large population of wandering, off-nuclear MBHs (e.g., A. Ricarte et al. 2021b; E. J. Weller et al. 2022; T. Di Matteo et al. 2023); but see also F. van Dokkelaar et al. (2025). The off-nuclear MBH population remains poorly understood and is therefore not considered in this work. Incorporating an additional model component to account for the off-nuclear fraction will be necessary when connecting our observational results to theoretical models. Directly detecting these off-nuclear MBHs is challenging, as most are dormant and only rarely observed as hyperluminous X-ray sources, dual AGNs, or off-nucleus tidal disruption events. Nevertheless,

³³ It is worth noting that, for low- λ_{Edd} MBHs, the ratio of UV–optical luminosity to L_X deviates significantly from that of AGNs, as radiatively inefficient accretion flows lack the prominent UV–optical emission bump produced by standard thin disks. Nevertheless, k_{bol} is less suppressed because the IR emission, where ADAF SEDs peak, contributes substantially to L_{bol} .




Figure 9. The BHMf, plotted as $\phi(M_{\text{BH}})$ and $M_{\text{BH}}\phi(M_{\text{BH}})$ in the left and right panels, respectively. The right panel highlights the contributions from MBHs at different masses to the total M_{BH} budget. The black curve represents the BHMf based on the occupation fraction derived in this study, with the gray shaded regions indicating the 1σ , 2σ , and 3σ uncertainty ranges, accounting solely for uncertainties in f_{occ} . The cyan shaded area illustrates the possible BHMf range derived from F. Shankar et al. (2020b), which is extensive due to the uncertainties in the $M_{\text{BH}}-M_*$ scaling relation. The effect of a declining f_{occ} is to flatten $\phi(M_{\text{BH}})$ for $M_* \lesssim 10^6 M_\odot$.

some observational constraints are possible, as discussed by, for example, M. Guo et al. (2020) and A. Ricarte et al. (2021a).

5. Summary

In this work, we present joint observational constraints on the black hole occupation fraction f_{occ} of local galaxies and the probability $p(\lambda)$ of the SARDF of their central MBHs, based on over two decades of Chandra observations of 1606 nearby galaxies within 50 Mpc. Our main results are summarized as follows.

1. $f_{\text{occ}} \approx 1$ for massive galaxies with $M_* \gtrsim 10^{10} M_\odot$, but it decreases rapidly at lower M_* . We obtain $f_{\text{occ}}(10^8 M_\odot < M_* < 10^9 M_\odot) = 33\%_{-9\%}^{+13\%}$ and $f_{\text{occ}}(10^9 M_\odot < M_* < 10^{10} M_\odot) = 66\%_{-7\%}^{+8\%}$. The $f_{\text{occ}}(M_*)$ function is tabulated in Table 2. See Section 3.4.
2. The MBH SARDF, $p(\lambda)$, can be well described by a Box-Cox distribution for inactive galaxies. Its main features include a peak at $\log \lambda \approx 28$ and a power-law tail at $\log \lambda \gtrsim 29$, which can further connect with the AGN $p(\lambda)$ at $\log \lambda \gtrsim 31.5$. $p(\lambda)$ also shows no significant shifts with respect to M_* . We demonstrate that this newly parameterized $p(\lambda)$ is more consistent with the data compared to normal or power-law distributions. See Sections 3.4, 4.1, and 4.2.
3. We examine f_{occ} and $p(\lambda)$ separately for different types of galaxies, including morphological classifications (early versus late type) and star-forming properties (quiescent versus star forming). We find that the conclusion of decreasing f_{occ} at $M_* \lesssim 10^{10} M_\odot$ also holds for early-type, late-type, and star-forming subsamples, but we do not detect statistically significant evidence for quiescent galaxies. Nevertheless, these different types of galaxies share statistically consistent f_{occ} and $p(\lambda)$, although late-type and star-forming galaxies tentatively have higher $p(\lambda)$ at $\lambda \gtrsim 10^{29} \text{ erg s}^{-1} M_\odot^{-1}$. See Section 4.3.
4. We derive the BHMf from the galaxy stellar mass function and the $M_{\text{BH}}-M_*$ scaling relation, with f_{occ}

included. The BHMf at $M_{\text{BH}} \lesssim 10^6 M_\odot$ becomes flat. We find that the f_{occ} uncertainty is subdominant compared to the scaling-relation uncertainty when determining the BHMf normalization. See Section 4.5.

To the extent that MBH occupation in nearby galaxies serves as a viable diagnostic for MBH seeds at high redshifts (see Section 1 and references therein), the results presented in this work appear to favor a heavy seeding scenario. However, further theoretical work is necessary to quantitatively link f_{occ} with MBH seeding models. First, as discussed in Section 4.5, there is an additional layer of uncertainty due to the off-nuclear MBH population, which remains poorly understood. Second, the extrapolation of both light-seed and heavy-seed models to low redshifts is subject to significant uncertainties. For example, A. Ricarte & P. Natarajan (2018) argue that light-seed models may result in a low f_{occ} if most seeds undergo little growth and thus remain undetectable at $z=0$. Conversely, A. K. Bhowmick et al. (2025) contends that heavy-seed models could yield a high f_{occ} if MBH seed formation is more prevalent. Despite these challenges, this work provides a crucial observational benchmark for constraining this fundamental quantity, and future research can further refine MBH seeding models based on f_{occ} .

Acknowledgments

We thank the anonymous referee for constructive comments and suggestions. We are grateful to Riccardo Arcodia and Feng Yuan for helpful insights. This research has made use of data obtained from the Chandra Source Catalog, provided by the Chandra X-ray Center (CXC). Support for this work was provided by the National Aeronautics and Space Administration through Chandra Award Number 21700098 issued by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for and on behalf of the National Aeronautics Space Administration under contract NAS8-03060. A.C.S acknowledges support from National Science Foundation astronomy and astrophysics grant AST-2108180.

W.N.B. acknowledges support from NSF grant AST-2407089. A.E.R. gratefully acknowledges support provided by NSF through CAREER award 2235277.

Appendix A X-Ray-detected Sources with $M_\star < 10^8 M_\odot$

Three sources with $M_\star < 10^8 M_\odot$ are detected by Chandra. Such cases, where possible MBHs reside in very low-mass galaxies, are often interesting individually, and there may not be a consensus regarding their nature. We do not aim to incorporate their complexities in this work and only briefly comment on their properties as noted in the literature.

1. *PGC 027864*. R.A. = 09:44:01.87, decl. = -00:38:32.1, distance = 21 Mpc, $\log M_\star = 7.4$, and $\log L_X = 39.0$. This galaxy is metal poor and discussed in detail by M. Reefe et al. (2023). It contains another even brighter X-ray source 3" away (P. Senchyna et al. 2020) that is not included in our analyses. The galaxy nucleus has a broad H α line, implying an estimated $M_{\text{BH}} \approx 3150 M_\odot$. M. Reefe et al. (2023) showed that the galaxy nucleus has an [Fe x] $\lambda 6374$ coronal line, which is strong evidence indicating AGN activity; however, this line may also be a misidentified Si II $\lambda 6371$ line (E. C. Herenz et al. 2023).
2. *PGC 031259 (also known as Mrk 1434)*. R.A. = 10:34:10.15, decl. = +58:03:49.1, distance = 36 Mpc, $\log M_\star = 7.1$, and $\log L_X = 39.9$. This galaxy is metal poor and discussed specifically by E. Thygesen et al. (2023). It contains two luminous X-ray sources that can only be explained by either ULXs or AGNs. Its high L_X also leads to its independent selection as a dwarf AGN by XMM-Newton (K. L. Birchall et al. 2020; J. M. Cann et al. 2024) and eROSITA (S. D. Bykov et al. 2024).
3. *PGC 1206166*. R.A. = 15:08:22.69, decl. = +01:47:55.0, distance = 24 Mpc, $\log M_\star = 8.0$, and $\log L_X = 38.8$. This galaxy is an early-type member of the NGC 5846 group (A. Marino et al. 2016). Its X-ray emission was reported by B. Miller et al. (2012) and is unlikely to be explained by XRBs.

Appendix B Selection Effects

Two main selection effects are likely to affect our sample. First, the parent galaxy catalog, the 50MGC, is incomplete for dwarf galaxies. Second, Chandra ACIS may be biased in favor of accretion-powered MBHs. For the first selection effect, we find that the galaxies considered in this work become systematically more compact than the expected mass-size relation in S. G. Carlsten et al. (2021) at $M_\star \lesssim 10^8 M_\odot$, indicating that we may be missing low surface brightness systems. We expect that this effect has only a limited systematic impact on our inference of f_{occ} . To examine this, we consider the subsample of galaxies within 30 Mpc and with $M_\star > 10^8 M_\odot$; as shown in Section 4.2 of D. Ohlson et al. (2024), galaxies in this regime should be complete. The resulting ACIS-covered sample size is reduced to 1007 galaxies (63% of the original sample), with 277 compact X-ray detections. We rerun our analyses for this mass-complete subsample and verify that the median f_{occ} is similar to that of the original sample, though with somewhat larger uncertainties (not explicitly plotted here).

The second selection effect likely has a more direct impact, as the galaxies in our sample were initially proposed for Chandra observations for a variety of scientific reasons, a fraction of which could be related to the presence of MBHs. The Chandra proposal subject categories include "solar system and exoplanets," "stars and white dwarfs," "white-dwarf binaries and cataclysmic variables," "black hole and neutron-star binaries," "supernova, supernova remnant, and isolated neutron star," "gravitational wave event," "normal galaxies, active galaxies, and quasars," "clusters of galaxies," "extragalactic diffuse emission and surveys," and "galactic diffuse emission and surveys." Among these, only "active galaxies and quasars," and in some cases "normal galaxies," are directly relevant to central MBHs; the other categories are generally unrelated. Therefore, for simplicity, we assume that the Chandra

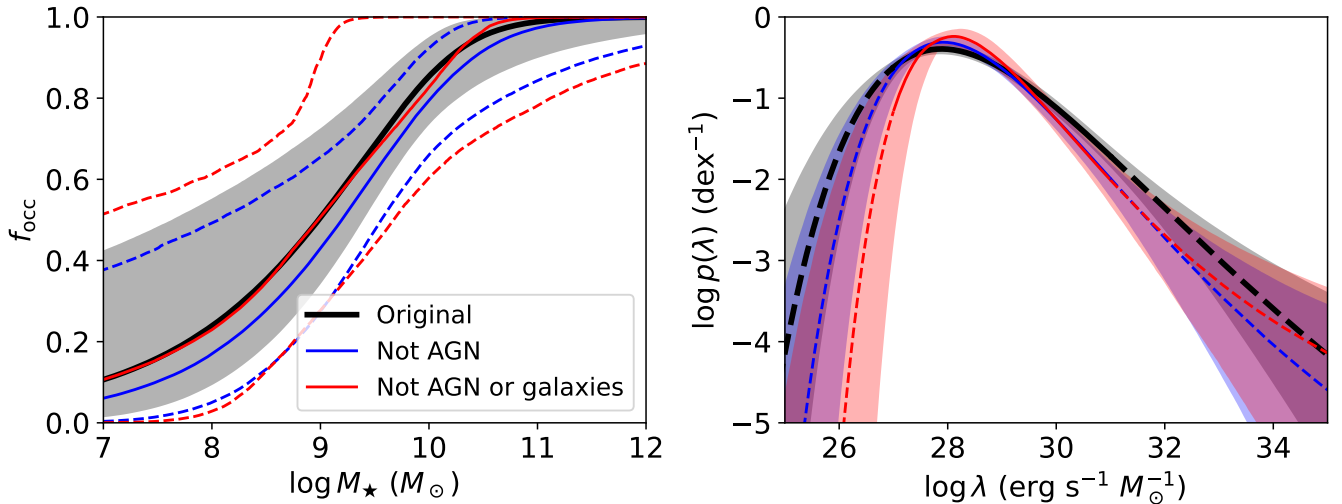


Figure 10. Comparison of f_{occ} (left) and $p(\lambda)$ (right). Black curves represent the original results. Blue curves correspond to the sample excluding all sources within 1' of the observational centers with the Chandra proposal subject category "active galaxies and quasars;" red curves indicate the sample further excluding sources based on both the "active galaxies and quasars" and "normal galaxies" categories. In the left panel, solid curves show the medians of the posterior distributions, while the shaded regions and dashed curves indicate the 2σ ranges. In the right panel, curves represent the medians, with dashed segments showing extrapolations beyond the data range (as in Figure 7), and shaded regions correspond to the 2σ ranges.

selection effects are primarily driven by targeted observations of individual sources under the “active galaxies and quasars” and (potentially) “normal galaxies” categories. To account for this, we attempt to remove all sources within $1'$ of the centers of these specific observations. If we exclude sources only from the “active galaxies and quasars” category, 1312 sources (82% of the original sample size) remain in our sample, with 263 compact detections. If we instead exclude sources from both the “active galaxies and quasars” and “normal galaxies” categories, 847 sources (53% of the original sample size) remain, with 120 compact detections. We note that this removal procedure may introduce an opposite selection bias, namely, that high- λ sources could become underrepresented in the final sample.

We rerun our analysis and present the corresponding results in Figure 10. Compared to the original results, the median f_{occ} values are similar, and $p(\lambda)$ is only slightly suppressed at the high- λ end. Overall, the measurements remain largely consistent, indicating that the selection effects associated with intentional Chandra observations of MBH-related phenomena are smaller than the statistical uncertainties. However, it is unlikely that selection effects can be fully eliminated. While they may be mitigated by using samples from dedicated surveys of specific sky regions (e.g., the Virgo Cluster; E. Gallo et al. 2010), the resulting smaller sample size, as demonstrated by M15, would lack sufficient constraining power.

Appendix C Results Based on B. P. Miller et al. (2015)

We apply the methodology presented in the main text to the data in M15 to examine if we can obtain consistent results. Their sample included approximately 200 optically selected early-type galaxies within 30 Mpc, which is a subsample of ours. We adopt their “clean” sample from Table 1 and present the f_{occ} results in Figure 11. Both our measured f_{occ} and the original M15 f_{occ} are consistent with $f_{\text{occ}} = 1$; the M15 f_{occ} constraint appears narrower outside the range $10^8 M_{\odot} \lesssim M_{\star} \lesssim 10^{10} M_{\odot}$ because M15 assumed $f_{\text{occ}}(M_{\star} < 10^7 M_{\odot}) \approx 0$ and $f_{\text{occ}}(M_{\star} > 10^{10} M_{\odot}) \approx 1$. The overall consistency serves as further evidence that the low f_{occ} inferred in the main text is indeed genuine based on the improved data—if we downgrade the data quality, we would not detect the decreasing trend of f_{occ} at $M_{\star} < 10^{10} M_{\odot}$.




Figure 11. f_{occ} results when applying our methodology to the M15 data. The solid and dashed curves, along with their corresponding shaded regions, represent the 1σ and 2σ ranges, respectively. The original 1σ measurements from M15 are shown as the orange region for comparison. In agreement with M15, our results show no evidence of decreasing f_{occ} at low M_{\star} .

Appendix D Inference for Mock Data

To validate our methodology, we generate mock data and examine our fitting results. Section 2.2 of B25 presented similar simulations and validated the overall consistency when recovering f_{occ} . In this appendix, we primarily focus on examining whether we can still reasonably recover $p(\lambda)$ and f_{occ} if the low- λ $p(\lambda)$ instead follows a power law. We extrapolate the SARDF at $M_{\star} = 10^{10} M_{\odot}$ and $z = 0.05$ from F. Zou et al. (2024b) to low λ using a power law. Since the integral of $p(\lambda)$ is 1 by definition, the distribution must have a sharp decrease below a minimum $\log \lambda$, which we infer to be 28.4. We set $f_{\text{occ}}(M_{\star} \geq 10^{9.5} M_{\odot}) = 1$ and consider two extreme cases for dwarf galaxies: $f_{\text{occ}}(M_{\star} < 10^{9.5} M_{\odot}) = 0$ or 1. For each M_{\star} value in our sample, we randomly assign a λ value as per Equation (1). We then adopt an L_X threshold of $10^{39.4} \text{ erg s}^{-1}$, ensuring that the fraction of massive galaxies with $M_{\star} > 10^{10} M_{\odot}$ above the threshold matches the fraction of uncensored data points at $M_{\star} > 10^{10} M_{\odot}$ in our sample (54%). All sources with L_X below this threshold are treated as upper limits (note that the minimum $\log \lambda$ and the L_X threshold are higher than those in the main text because a power-law




Figure 12. The recovered f_{occ} (left) and $p(\lambda)$ (right) from mock data generated based on the $p(\lambda)$ shown as the black curve in the right panel. Our inputs are $f_{\text{occ}}(M_{\star} \geq 10^{9.5} M_{\odot}) = 1$ and $f_{\text{occ}}(M_{\star} < 10^{9.5} M_{\odot}) = 0$ (blue) or 1 (red). The blue and red solid curves represent the medians of the fitted posteriors, and the shaded regions indicate the 2σ levels. The vertical dashed line in the left panel marks $M_{\star} = 10^{9.5} M_{\odot}$. The fitted results are largely consistent with the inputs.

extrapolation of $p(\lambda)$ itself is inconsistent with observations, as discussed in Section 4.2).

We use the same method as in Section 3 to measure $p(\lambda)$ and f_{occ} for the mock data, and the results are presented in Figure 12. The fitted f_{occ} values are in good agreement with the input f_{occ} in both the low- and high- f_{occ} cases. For $p(\lambda)$, our fitted curves capture the general trend well: the peak is slightly larger than the minimum λ , and the high- λ tail broadly resembles the power-law tail. Small differences are present because our Box-Cox distribution family does not exactly replicate a power-law distribution with a sharp drop.

ORCID iDs

Fan Zou <https://orcid.org/0000-0002-4436-6923>
 Elena Gallo <https://orcid.org/0000-0001-5802-6041>
 Anil C. Seth <https://orcid.org/0000-0003-0248-5470>
 Edmund Hodges-Kluck <https://orcid.org/0000-0002-2397-206X>
 David Ohlson <https://orcid.org/0009-0004-9457-2495>
 Tommaso Treu <https://orcid.org/0000-0002-8460-0390>
 Vivienne F. Baldassare <https://orcid.org/0000-0003-4703-7276>
 W. N. Brandt <https://orcid.org/0000-0002-0167-2453>
 Jenny E. Greene <https://orcid.org/0000-0002-5612-3427>
 Piero Madau <https://orcid.org/0000-0002-6336-3293>
 Dieu D. Nguyen <https://orcid.org/0000-0002-5678-1008>
 Richard M. Plotkin <https://orcid.org/0000-0002-7092-0326>
 Amy E. Reines <https://orcid.org/0000-0001-7158-614X>
 Alberto Sesana <https://orcid.org/0000-0003-4961-1606>
 Jong-Hak Woo <https://orcid.org/0000-0002-8055-5465>
 Jianfeng Wu <https://orcid.org/0000-0001-7349-4695>

References

Aalen, O. 1978, *AnSta*, 6, 701
 Agarwal, B., Smith, B., Glover, S., Natarajan, P., & Khochfar, S. 2016, *MNRAS*, 459, 4209
 Aird, J., Coil, A. L., & Georgakakis, A. 2018, *MNRAS*, 474, 1225
 Akritas, M. G., Murphy, S. A., & LaValley, M. P. 1995, *JASA*, 90, 170
 Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv:1702.00786
 Babak, S., Gair, J., Sesana, A., et al. 2017, *PhRvD*, 95, 103012

Baganoff, F. K., Maeda, Y., Morris, M., et al. 2003, *ApJ*, 591, 891
 Baldassare, V. F., Geha, M., & Greene, J. 2020, *ApJ*, 896, 10
 Beckmann, R. S., Dubois, Y., Volonteri, M., et al. 2023, *MNRAS*, 523, 5610
 Bellovary, J. M., Hayoune, S., Chaffa, K., et al. 2021, *MNRAS*, 505, 5129
 Benet, L., & Sanders, D. P. 2019, *JOSS*, 4, 1043
 Betancourt, M. 2017, arXiv:1701.02434
 Bhowmick, A. K., Blecha, L., Torrey, P., et al. 2025, *MNRAS*, 538, 518
 Bi, S., Feng, H., & Ho, L. C. 2020, *ApJ*, 900, 124
 Birchall, K. L., Watson, M. G., & Aird, J. 2020, *MNRAS*, 492, 2268
 Birchall, K. L., Watson, M. G., Aird, J., & Starling, R. L. C. 2022, *MNRAS*, 510, 4556
 Bogdán, Á., Goulding, A. D., Natarajan, P., et al. 2024, *NatAs*, 8, 126
 Bongiorno, A., Schulze, A., Merloni, A., et al. 2016, *A&A*, 588, A78
 Boquien, M., Burgarella, D., Roehly, Y., et al. 2019, *A&A*, 622, A103
 Box, G. E. P., & Cox, D. R. 1964, *Journal of the Royal Statistical Society: Series B (Methodological)*, 26, 211
 Brandt, W. N., & Alexander, D. M. 2015, *A&ARv*, 23, 1
 Brandt, W. N., & Yang, G. 2022, in *Handbook of X-Ray and Gamma-Ray Astrophysics*, ed. C. Bambi & A. Sanganello (Berlin: Springer), 78
 Bruzual, G., & Charlot, S. 2003, *MNRAS*, 344, 1000
 Burke, C. J., Liu, X., Shen, Y., et al. 2022, *MNRAS*, 516, 2736
 Burke, C. J., Natarajan, P., Baldassare, V. F., & Geha, M. 2025, *ApJ*, 978, 77
 Bykov, S. D., Gilfanov, M. R., & Sunyaev, R. A. 2024, *MNRAS*, 527, 1962
 Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, *ApJ*, 533, 682
 Cann, J. M., Weaver, K. A., Pfeifle, R. W., et al. 2024, *ApJ*, 961, 178
 Carlsten, S. G., Greene, J. E., Greco, J. P., Beaton, R. L., & Kado-Fong, E. 2021, *ApJ*, 922, 267
 Chabrier, G. 2003, *PASP*, 115, 763
 Cho, H., & Woo, J.-H. 2024, *ApJ*, 969, 93
 Dale, D. A., Helou, G., Magdis, G. E., et al. 2014, *ApJ*, 784, 83
 Dashyan, G., Silk, J., Mamon, G. A., Dubois, Y., & Hartwig, T. 2018, *MNRAS*, 473, 5698
 Davis, S. W., & Laor, A. 2011, *ApJ*, 728, 98
 Di Matteo, T., Ni, Y., Chen, N., et al. 2023, *MNRAS*, 525, 1479
 Dickey, C. M., Geha, M., Wetzel, A., & El-Badry, K. 2019, *ApJ*, 884, 180
 Driver, S. P., Bellstedt, S., Robotham, A. S. G., et al. 2022, *MNRAS*, 513, 439
 Duras, F., Bongiorno, A., Ricci, F., et al. 2020, *A&A*, 636, A73
 Evans, I. N., Evans, J. D., Martínez-Galarza, J. R., et al. 2024, *ApJS*, 274, 22
 Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2022, *ApJL*, 930, L15
 Feigelson, E. D., & Babu, G. J. 2012, arXiv:1205.2064
 Foord, A., Gallo, E., Hodges-Kluck, E., et al. 2017, *ApJ*, 841, 51
 Freeman, J., & Modarres, R. 2006, *Statistics & Probability Letters*, 76, 764
 Fryer, C. L., Woosley, S. E., & Heger, A. 2001, *ApJ*, 550, 372
 Gallo, E., Treu, T., Marshall, P. J., et al. 2010, *ApJ*, 714, 25
 Gallo, E., & Sesana, A. 2019, *ApJL*, 883, L18
 Gallo, E., Hodges-Kluck, E., Treu, T., et al. 2023, arXiv:2311.09161
 Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, *ApJ*, 689, 1044
 Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009, *ApJ*, 692, 1075

- Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, *ApJ*, 964, 39
- Greene, J. E., Strader, J., & Ho, L. C. 2020, *ARA&A*, 58, 257
- Guetzoyan, P., Aird, J., Georgakakis, A., et al. 2025, *MNRAS*, 536, 79
- Guo, M., Inayoshi, K., Michiyama, T., & Ho, L. C. 2020, *ApJ*, 901, 39
- Häberle, M., Neumayer, N., Seth, A., et al. 2024, *Natur*, 631, 285
- Habouzit, M., Volonteri, M., Latif, M., Dubois, Y., & Peirani, S. 2016, *MNRAS*, 463, 529
- Helsel, D. R. 2011, *Statistics for Censored Environmental Data Using Minitab and R*, Vol. 77 (New York: Wiley)
- Herenz, E. C., Micheva, G., Weibacher, P. M., et al. 2023, *RNAAS*, 7, 99
- Ho, L. C. 2009, *ApJ*, 699, 626
- Hodges-Kluck, E. J., Gallo, E., Seth, A., Greene, J., & Baldassare, V. 2020, *ApJ*, 898, 106
- Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv:1008.4686
- Inayoshi, K., Visbal, E., & Haiman, Z. 2020, *ARA&A*, 58, 27
- Isserlis, L. 1918, *Biometrika*, 12, 134
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
- Karachentsev, I. D., & Kaisina, E. I. 2013, *AJ*, 146, 46
- Karachentsev, I. D., Makarov, D. I., & Kaisina, E. I. 2013, *AJ*, 145, 101
- Kelly, B. C. 2007, *ApJ*, 665, 1489
- Kormendy, J., Fisher, D. B., Cornell, M. E., & Bender, R. 2009, *ApJS*, 182, 216
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, 51, 511
- Kovlakas, K., Zezas, A., Andrews, J. J., et al. 2020, *MNRAS*, 498, 4790
- Kulier, A., Ostriker, J. P., Natarajan, P., Lackner, C. N., & Cen, R. 2015, *ApJ*, 799, 178
- Lee, N., Gallo, E., Hodges-Kluck, E., et al. 2019, *ApJ*, 874, 77
- Lehmer, B. D., Eufrazio, R. T., Basu-Zych, A., et al. 2021, *ApJ*, 907, 17
- Lehmer, B. D., Eufrazio, R. T., Tzanavaris, P., et al. 2019, *ApJS*, 243, 3
- Leja, J., Speagle, J. S., Ting, Y.-S., et al. 2022, *ApJ*, 936, 165
- Licquia, T. C., & Newman, J. A. 2015, *ApJ*, 806, 96
- López, I. E., Yang, G., Mountrichas, G., et al. 2024, *A&A*, 692, A209
- Mahadevan, R. 1997, *ApJ*, 477, 585
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024, *A&A*, 691, A145
- Marino, A., Mazzei, P., Rampazzo, R., & Bianchi, L. 2016, *MNRAS*, 459, 2212
- Mezcua, M., Civano, F., Marchesi, S., et al. 2018, *MNRAS*, 478, 2576
- Mezcua, M., & Domínguez Sánchez, H. 2020, *ApJL*, 898, L30
- Mezcua, M., Siudek, M., Suh, H., et al. 2023, *ApJL*, 943, L5
- Millard, S. P. 2013, *EnvStats: An R Package for Environmental Statistics* (Berlin: Springer)
- Miller, B., Gallo, E., Treu, T., & Woo, J.-H. 2012, *ApJ*, 747, 57
- Miller, B. P., Gallo, E., Greene, J. E., et al. 2015, *ApJ*, 799, 98
- Moustakas, J., Lang, D., Dey, A., et al. 2023, *ApJS*, 269, 3
- Narayan, R., & Yi, I. 1995, *ApJ*, 452, 710
- Natarajan, P. 2014, *GRGr*, 46, 1702
- Natarajan, P., Pacucci, F., Ricarte, A., et al. 2024, *ApJL*, 960, L1
- Nguyen, D. D., Seth, A. C., Neumayer, N., et al. 2018, *ApJ*, 858, 118
- Nguyen, D. D., Seth, A. C., Neumayer, N., et al. 2019, *ApJ*, 872, 104
- Ohlson, D., Seth, A. C., Gallo, E., Baldassare, V. F., & Greene, J. E. 2024, *AJ*, 167, 31
- Pacucci, F., Mezcua, M., & Regan, J. A. 2021, *ApJ*, 920, 134
- Pardo, K., Goulding, A. D., Greene, J. E., et al. 2016, *ApJ*, 831, 203
- Partmann, C., Naab, T., Lahén, N., et al. 2025, *MNRAS*, 537, 956
- Pechetti, R., Seth, A., Kamann, S., et al. 2022, *ApJ*, 924, 48
- Pellegrini, S. 2010, *ApJ*, 717, 640
- Penny, S. J., Masters, K. L., Smethurst, R., et al. 2018, *MNRAS*, 476, 979
- Petersson, J., Hirschmann, M., Tress, R. G., et al. 2025, arXiv:2504.08035
- Pucha, R., Juneau, S., Dey, A., et al. 2025, *ApJ*, 982, 10
- Reefe, M., Satyapal, S., Sexton, R. O., et al. 2023, *ApJL*, 946, L38
- Reines, A. E., Greene, J. E., & Geha, M. 2013, *ApJ*, 775, 116
- Reines, A. E., & Volonteri, M. 2015, *ApJ*, 813, 82
- Reines, A. E. 2022, *NatAs*, 6, 26
- Reynolds, C. S., Kara, E. A., Mushotzky, R. F., et al. 2023, *Proc. SPIE*, 12678, 126781E
- Ricarte, A., & Natarajan, P. 2018, *MNRAS*, 481, 3278
- Ricarte, A., Tremmel, M., Natarajan, P., & Quinn, T. 2021a, *ApJL*, 916, L18
- Ricarte, A., Tremmel, M., Natarajan, P., Zimmer, C., & Quinn, T. 2021b, *MNRAS*, 503, 6098
- Richards, F. J. 1959, *JEBot*, 10, 290
- Sacchi, A., Bogdán, Á., Chadayammuri, U., & Ricarte, A. 2024, *ApJ*, 974, 14
- Sakurai, Y., Yoshida, N., Fujii, M. S., & Hirano, S. 2017, *MNRAS*, 472, 1677
- Sartori, L. F., Schawinski, K., Treister, E., et al. 2015, *MNRAS*, 454, 3722
- Senchyna, P., Stark, D. P., Mirocha, J., et al. 2020, *MNRAS*, 494, 941
- Shankar, F., Allevato, V., Bernardi, M., et al. 2020a, *NatAs*, 4, 282
- Shankar, F., Bernardi, M., Sheth, R. K., et al. 2016, *MNRAS*, 460, 3119
- Shankar, F., Weinberg, D. H., Marsden, C., et al. 2020b, *MNRAS*, 493, 1500
- Sharma, R. S., Brooks, A. M., Tremmel, M., et al. 2022, *ApJ*, 936, 82
- She, R., Ho, L. C., & Feng, H. 2017, *ApJ*, 842, 131
- Silk, J. 2017, *ApJL*, 839, L13
- Stone, N. C., & Metzger, B. D. 2016, *MNRAS*, 455, 859
- Tak, H., Ghosh, S. K., & Ellis, J. A. 2018, *MNRAS*, 481, 277
- Thygesen, E., Plotkin, R. M., Soria, R., et al. 2023, *MNRAS*, 519, 5848
- Ueda, Y., Akiyama, M., Hasinger, G., Miyaji, T., & Watson, M. G. 2014, *ApJ*, 786, 104
- Valiante, R., Schneider, R., Volonteri, M., & Omukai, K. 2016, *MNRAS*, 457, 3356
- van Dokelaar, F., Mayer, L., Capelo, P. R., & Tamfal, T. 2025, *MNRAS*, 538, 2255
- Voggel, K. T., Seth, A. C., Baumgardt, H., et al. 2019, *ApJ*, 871, 159
- Volonteri, M. 2012, *Sci*, 337, 544
- Wasleske, E. J., & Baldassare, V. F. 2024, *ApJ*, 971, 68
- Weigel, A. K., Schawinski, K., Caplar, N., et al. 2017, *ApJ*, 845, 134
- Weller, E. J., Pacucci, F., Hernquist, L., & Bose, S. 2022, *MNRAS*, 511, 2229
- Wilson, E. B. 1927, *JASA*, 22, 209
- Woods, T. E., Agarwal, B., Bromm, V., et al. 2019, *PASA*, 36, e027
- Xie, F.-G., & Yuan, F. 2012, *MNRAS*, 427, 1580
- Yang, G., Boquien, M., Buat, V., et al. 2020, *MNRAS*, 491, 740
- Yang, G., Brandt, W. N., Vito, F., et al. 2018, *MNRAS*, 475, 1887
- Yuan, F., & Narayan, R. 2014, *ARA&A*, 52, 529
- Yuan, F., Yoon, D., Li, Y.-P., et al. 2018, *ApJ*, 857, 121
- Zou, F., Brandt, W. N., Chen, C.-T., et al. 2022, *ApJS*, 262, 15
- Zou, F., Brandt, W. N., Gallo, E., et al. 2024a, *ApJ*, 976, 6
- Zou, F., Brandt, W. N., Ni, Q., et al. 2023, *ApJ*, 950, 136
- Zou, F., Yu, Z., Brandt, W. N., et al. 2024b, *ApJ*, 964, 183