

# Getting more out of black hole superradiance: a statistically rigorous approach to ultralight boson constraints from black hole spin measurements

Sebastian Hoof<sup>1</sup>,<sup>1,2</sup>★ David J. E. Marsh<sup>3</sup>,<sup>3</sup>★ Júlia Sisk-Reynés<sup>4,5</sup>,<sup>4,5</sup> James H. Matthews<sup>6</sup><sup>6</sup> and Christopher Reynolds<sup>7</sup><sup>7</sup>

<sup>1</sup>*Dipartimento di Fisica e Astronomia ‘Galileo Galilei’, Università degli Studi di Padova, Via F. Marzolo 8, I-35131 Padova, Italy*

<sup>2</sup>*Istituto Nazionale di Fisica Nucleare–Sezione di Padova, Via F. Marzolo 8, I-35131 Padova, Italy*

<sup>3</sup>*Theoretical Particle Physics and Cosmology, King’s College London, Strand, London WC2R 2LS, UK*

<sup>4</sup>*Center for Astrophysics|Harvard & Smithsonian, Cambridge, MA 02138, USA*

<sup>5</sup>*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

<sup>6</sup>*Department of Physics, Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*

<sup>7</sup>*Department of Astronomy, University of Maryland, College Park, MD 20742, USA*

Accepted 2025 September 11. Received 2025 August 12; in original form 2024 June 25

## ABSTRACT

Black hole (BH) superradiance can provide strong constraints on the properties of ultralight bosons (ULBs). While most of the previous work has focused on the theoretical predictions. Here, we investigate the most suitable statistical framework to constrain ULB masses and self-interactions using BH spin measurements. We argue that a Bayesian approach based on a simple time-scales analysis provides a clear statistical interpretation, deals with limitations regarding the reproducibility of existing BH analyses, incorporates the full information from BH data, and allows us to include additional nuisance parameters or to perform hierarchical modelling with BH populations in the future. We demonstrate the feasibility of our approach using mass and spin posterior samples for the X-ray binary BH M33 X-7 and, for the first time in this context, the supermassive BH IRAS 09149-6206. We explain the differences to existing ULB constraints in the literature and illustrate the effects of various assumptions about the superradiance process (equilibrium regime versus cloud collapse, higher occupation levels). As a result, our procedure yields the most statistically rigorous ULB constraints available in the literature, with important implications for the quantum chromodynamics (QCD) axion and axion-like particles. We encourage all groups analysing BH data to publish likelihood functions or posterior samples as supplementary material to facilitate this type of analysis, and for theory developments to compress their findings to effective time-scale modifications (<https://github.com/sebhoof/bhsr>).

**Key words:** astroparticle physics – black hole physics – elementary particles.

## 1 INTRODUCTION

This paper will introduce a simple yet rigorous statistical framework for constraining ultralight bosons (ULBs) using black hole (BH) superradiance (SR) from observational BH spin estimates. The framework introduces no new physics, using only previously developed models of BHSR. Our approach simplifies the constraint setting procedure to use only a time-scales analysis, rather than simulating BH populations and evolution. We argue that this simplified approach is appropriate given the inherent astrophysical uncertainties/systematics, and is user-friendly since the analysis is easy to repeat with modified time-scale assumptions. What we develop is a rigorous Bayesian framework to set limits on ULBs that allows to simply and directly use the posterior probabilities on BH model parameters published by observers, and advocate for future

publication of BH results in this format. The robust approach to data products will allow for transparent and easy compilation of future large BH data sets. Future improved understandings to the physics of BHSR should be incorporated via modifications to the characteristic time-scales involved, and we demonstrate this by comparing the bosonova model to the equilibrium model for ULB self-interactions.

Arvanitaki et al. (2010), Arvanitaki & Dubovsky (2011), and Arvanitaki, Baryakhtar & Huang (2015) showed that BHs with large spins can probe the existence of new bosonic degrees of freedom via BHSR. Due to a link between the BH mass and the boson mass, the resulting constraints are relevant for ULBs. Indeed, Arvanitaki et al. (2015), Cardoso et al. (2018), Stott & Marsh (2018), Stott (2020), Baryakhtar et al. (2021), and Mehta et al. (2021) placed constraints on ULBs with masses  $\mu \lesssim 10^{-10}$  eV using BHSR from BH spin measurements.

ULBs are of particular interest in particle astrophysics at present. The quantum chromodynamics (QCD) axion (Peccei & Quinn 1977a, b; Wilczek 1978; Weinberg 1978), an excellent dark matter (DM)

\* E-mail: [hoof@pd.infn.it](mailto:hoof@pd.infn.it) (SH); [david.j.marsh@kcl.ac.uk](mailto:david.j.marsh@kcl.ac.uk) (DJEM)

candidate (Preskill, Wise & Wilczek 1983; Abbott & Sikivie 1983; Dine & Fischler 1983; Turner 1983, 1986), is among this class of particles. There is a far-ranging experimental programme underway (see e.g. Marsh 2016; Irastorza & Redondo 2018; Chadha-Day, Ellis & Marsh 2022; Semertzidis & Youn 2022; Adams et al. 2022; O’Hare 2024, ) to search for QCD axions or axion-like particles (ALPs), which can also be DM candidates (see e.g. Arias et al. 2012). As we will review later, BHSR in the stellar mass range places constraints on the lower end of the allowed  $\mu$  range of the QCD axion. The ‘fuzzy’ DM model is also in the ULB class, and it is of interest as an alternative to cold DM with distinctive phenomenology on subgalactic scales (Khlopov, Malomed & Zeldovich 1985; Hu, Barkana & Gruzinov 2000; Marsh & Silk 2014; Schive, Chiueh & Broadhurst 2014; Hui et al. 2017), and as the DM candidate with the absolute lowest allowed mass (Alves Batista et al. 2021; Zimmermann et al. 2025). A combination of cosmological probes and galactic astrophysics probe the fuzzy DM mass from below, while BHSR in the supermassive mass range constrains it from above (Marsh 2016; Marsh & Hoof 2023). String/M-theory also predicts the existence of ULBs covering a wide mass range (Svrcek & Witten 2006; Conlon 2006; Cicoli, Goodsell & Ringwald 2012; Arvanitaki et al. 2010; Acharya, Bobkov & Kumar 2010), and it has recently become possible to use BHSR to rule out certain extra dimensional geometries, thus placing observational constraints on the string landscape (Mehta et al. 2021).

While BH spin estimates are a well-established probe of ULBs, there are a number of issues and subtleties related to the derivation of the constraints just discussed. Extracting the BH mass  $M$  and (dimensionless) spin  $a_*$  from BH data requires an involved analysis (McLure & Jarvis 2002; Casares & Jonker 2014; Bambi et al. 2021; Reynolds 2021), making this a highly specialized task where data sets need to be dealt with on a case-by-case basis. The resulting  $(M, a_*)$  estimates can be correlated and, since the strongest constraints come from BHs with close-to-maximal  $|a_*|$ , a standard Gaussian approximation may be inadequate. Moreover, the theoretical computation of BHSR-related effects is challenging and depends on the BH evolution. These complications require some compromises between a statistically rigorous analysis and the practical challenges: we wish to include uncertainties while – at the same time – leverage as much information as possible from the data, BH evolution, and BH population parameters.

We argue that a Bayesian framework admits all of these desiderata. We demonstrate the feasibility of our approach based on the X-ray binary BH M33 X-7 and the supermassive BH (SMBH) IRAS 09149-6206. These two sources were selected due to their well-understood nature and the availability of reliable  $(M, a_*)$  estimates. Ultimately, however, it is important to recognize that any BH mass and spin measurements – regardless of the BH system or measurement technique utilized – may be used for our method.

We use the  $(M, a_*)$  distributions from the two BHs above to obtain constraints on ULB parameters. In particular, these are the first constraints from IRAS 09149-6206. We compare the ULB limits arising from the full Bayesian treatment to other approximate statistical methods to explain the differences which arose in past works, and summarize strengths and weaknesses of each approach. While not the focus of this work, we also discuss the role of different computations of the BHSR rate – both with and without the inclusion of self-interactions.

At this point, let us emphasize that BH spin measurements are not the only way to search for ULBs using BHSR. Moreover, BH spins do not, individually, offer a route to discovery, except for in the case of population statistics (Arvanitaki & Dubovsky 2011;

Arvanitaki et al. 2017). Routes to discovery via BHSR can be pursued via dynamical means from either gravitational waves (e.g. Palomba et al. 2019; Zhu et al. 2020; Abbott et al. 2022a; Abbott et al. 2022b), binary orbits (e.g. Baumann, Chia & Porto 2019a; Hannuksela et al. 2019; Baumann et al. 2022; Xie & Huang 2024), or direct signatures of the ULB cloud (e.g. Foschi et al. 2023). These methods require significantly more involved modelling, and thus more inbuilt assumptions compared to constraints from BH spin time-scales, but the potential for discovery warrants the additional detailed treatment. Any alternative searches for ULBs through BHSR are complementary to the constraints we set here. In particular, our constraints would add additional information to the fitted ULB parameters in case of a discovery, given that the ULB’s existence has to be consistent with all available data.

The goal of this work is thus to provide rigorous and transparent BHSR constraints from observational BH spin estimates, which can be used with confidence by the community. Explaining the differences between past works contributes to understanding their underlying assumptions and statistical methodology, and hopefully encourages more groups analysing BH data to make their posterior samples or full analysis pipelines available. To facilitate this process, our code and the associated, redistributed BH data are publicly available (Hoof 2024).

## 2 BLACK HOLE EVOLUTION INCLUDING BHSR

SR arises due to an instability of the wave equation for a massive scalar field  $\phi$  with mass  $\mu$ ,

$$\square\phi - \mu^2\phi = 0, \quad (1)$$

where we initially neglect self-interactions (introduced in Section 2.3). The D’Alembertian is computed from the Kerr metric (Kerr 1963), for which we use the ‘mostly positive’ signature. The SR instability arises due to the presence of growing modes for  $\phi$ , which can be identified as imaginary eigenvalues of the associated Schrödinger equation. The energy required to create the ‘boson cloud’ is extracted from the mass and spin of the BH. The evolution of the mass and spin are treated quasi-statically, i.e. there is no backreaction on the Kerr space–time and one simply evolves the parameters in the Klein–Gordon equation (1).<sup>1</sup>

The SR rate  $\Gamma$  depends on  $\mu$  and is approximately maximized when the dimensionless ‘gravitational coupling’ of the atom-like, bosonic states around a spinning BH,

$$\alpha = GM\mu = 0.75 \left( \frac{M}{10M_\odot} \right) \left( \frac{\mu}{10^{-11} \text{ eV}} \right), \quad (2)$$

is  $\alpha \sim 1$ , i.e. it is approximately maximized when the boson Compton wavelength is of order the gravitational radius. Here,  $G$  is Newton’s constant. While we sometimes include the Planck mass  $M_{\text{P}} = 1/\sqrt{G} = 1.22 \times 10^{19} \text{ GeV}$  as a reference scale, we will generally set  $c = G = 1$  in what follows.

SR causes the spin  $a_*$  (and mass  $M$ ) of a BH to evolve in time,

$$\dot{a}_* = -\Gamma M_{\text{cloud}}/\omega^R \approx -\Gamma N_{\text{cloud}}, \quad (3)$$

<sup>1</sup>Backreaction is difficult to treat even in numerical relativity for a real scalar field due to the hierarchy of time-scales between the SR rate and the scalar field energy density free evolution time. This problem is not present for vector fields, which also undergo SR but have a slower evolution time for the free field energy density (see e.g. East & Pretorius 2017).

where the dimensionless BH spin is  $a_* \equiv a/M = J/M^2$ ,  $\omega^R$  is the ULB energy (the real part of the ULB frequency), and  $M_{\text{cloud}}$  and  $N_{\text{cloud}}$ , respectively, are the total mass and occupation number of the boson cloud forming around the BH. The latter evolves as

$$\dot{N}_{\text{cloud}} = \Gamma N_{\text{cloud}} + P_{\text{GW}}/\omega^R, \quad (4)$$

where  $P_{\text{GW}}$  is the small power radiated in gravitational waves (GW) from the cloud. We neglect  $P_{\text{GW}}$  in this work and refer the reader to Brito, Cardoso & Pani (2020, section 6.2.2) and Arvanitaki et al. (2017), Siemonsen, May & East (2023) for further details on the phenomenology of GW emission from boson clouds.

In the non-relativistic regime, the boson cloud can be thought of as levels in hydrogen atom-like states. The complete system is described by a set of coupled rate equations for the occupation numbers of each level, which grow due to BHSR, along with interactions between the levels due to scattering, and losses due to emissions to infinity and flux across the horizon (see e.g. Baryakhtar et al. 2021). An astrophysical BH will also be governed by other processes, having some typical time-scale  $\tau_{\text{BH}}$ , which we discuss further in Section 3.1.

## 2.1 Computing the free BHSR rates

Starting from Teukolsky (1972), the computation of the BHSR rates has seen a number of advances and improvements. In particular the analytical approximations of Detweiler (1980) provide a straightforward way to compute BHSR rates.<sup>2</sup> Semi-analytical results, e.g. by Dolan (2007), allowed a further refinement of the (semi-)analytical computations (see Baumann et al. 2019a; Baumann et al. 2019b; Bao, Xu & Zhang 2022).

It is illuminating to discuss the approximate analytical result. We use the metric of an uncharged, rotating BH with angular momentum  $a$  and mass  $M$  (Kerr BH; Kerr 1963), and the coordinate system of Boyer & Lindquist (1967). The outer horizon of a Kerr BH is  $r_+/M \equiv 1 + (1 - a_*^2)^{1/2}$ .

The bosonic state around the BH can be labelled by ‘quantum numbers’  $|nlm\rangle$ , where  $n \geq 2$  is the principal quantum number,<sup>3</sup> and has the complex-valued frequency  $\omega_{nlm} = \omega_{nlm}^R + i\omega_{nlm}^I$ . The real part  $\omega_{nlm}^R$  corresponds to the ULB energy and the imaginary part  $\omega_{nlm}^I$ , if positive, sets the growth rate of the field amplitude. The BHSR rate is defined as the growth of the particle number and is thus  $\Gamma_{nlm} = 2\omega_{nlm}^I$ . The latter is the case, and the  $|nlm\rangle$  level is superradiant, if the SR condition is fulfilled,

$$\alpha/l \leq 1/2. \quad (5)$$

Hod (2016) and Richartz, Luís Rosa & Berti (2024) derived general bounds on  $\mu$  for superradiant states, which are also valid for levels with  $n \gg 1$ .

### 2.1.1 Non-relativistic approximation

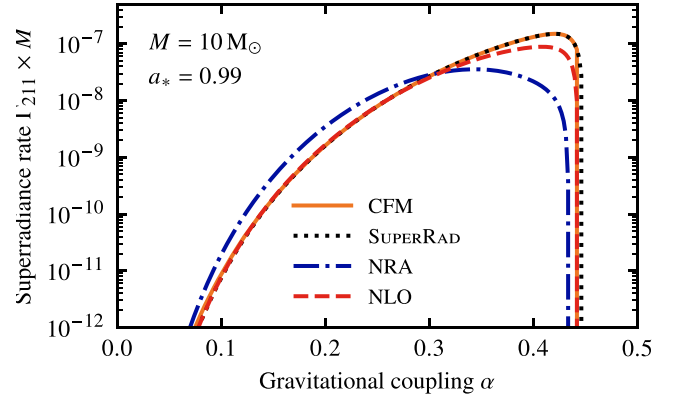
To first order in the non-relativistic approximation (NRA), we have (Arvanitaki & Dubovsky 2011)

$$\omega_{nlm}^R = \mu \left[ 1 - \frac{\alpha^2}{2n^2} + \mathcal{O}(\alpha^4) \right], \quad (6)$$

$$\omega_{nlm}^I \approx \mu (ma_* - 2r_+\mu) C_{nl} \Pi_{lm} \alpha^{4l+4} \quad (7)$$

<sup>2</sup>After correcting a missing factor 2 (e.g. Pani et al. 2012; Rosa 2013).

<sup>3</sup>Note that some authors use  $n' = n - l - 1 \geq 0$  instead.



**Figure 1.** The [211] SR rate (times the gravitational radius,  $M$ ), computed using different methods. For benchmark values of  $M = 10 M_{\odot}$  and  $a_* = 0.99$ , we compare the continued fraction method (CFM; orange line), adopted in this work, to the SUPERRAD code (dotted, black line), and our implementation of the next-to-leading-order corrections (NLO; dashed, blue line). For completeness, we also show the non-relativistic approximation (NRA; dash-dotted, red line).

$$\text{with } C_{nl} = \frac{2^{4l+2} (n+l)!}{n^{2l+4} (n-l-1)!} \left[ \frac{l!}{(2l)!(2l+1)!} \right]^2 \quad (8)$$

$$\text{and } \Pi_{lm} = \prod_{k=1}^l \left[ k^2 (1 - a_*^2) + (ma_* - 2r_+\mu)^2 \right]. \quad (9)$$

Terms at higher orders, which also introduce a dependence on the  $l$  and  $m$  quantum numbers, are e.g. discussed by Baumann et al. (2019b).

### 2.1.2 Higher-order corrections and continued fraction method

We can further refine the NRA result from Section 2.1.1. For instance, the PYTHON package SUPERRAD (Siemonsen et al. 2023) uses the QNM package (Stein 2019) to solve the radial Teukolsky equation using Leaver’s continued fraction method (CFM; Leaver 1985) in order to compute the relativistic BHSR rates. The numerical solution in the relativistic regime is valid at large  $\alpha$  and is matched consistently onto the Newtonian approximation at small  $\alpha$ , found from solving for hydrogen atom-like bound states and including terms up to  $\mathcal{O}(\alpha^5)$  (see Baumann et al. 2019b). SUPERRAD is fast and also has good control over the numerical accuracy of the methods, with a relative error on the SR time-scale estimated to be better than 1 per cent. It also includes gravitational wave emission.

However, the limitation of SUPERRAD is that it includes only modes with  $l = m = 1, 2$ , does not include self-interactions, and will not return rates at the highest  $a_*$  values. We thus use our own implementation of the CFM, following Dolan (2007).

Fig. 1 shows the results of our implementation for the [211] rate  $\Gamma_{211}$  (orange line), which compares well with the result of SUPERRAD (dotted, black line). To facilitate the numerical root-finding for the complex-valued root  $\omega_{211}$ , we also compute the next-to-leading-order (NLO) corrections (Bao et al. 2022) as a starting point, which we show as a dashed, red line for comparison. The computed rates differ at most by an order of magnitude, which, however, should be kept in mind when comparing different results in the literature. This applies in particular to the NRA (dash-dotted, blue line).

Since this work was first drafted, a study by Witte & Mummery (2025) appeared, which addresses the relativistic corrections in more

detail. For the dominant  $|211\rangle$  rate, the corrections are shown to be of order 20 per cent. It would be possible to update the rates used in our public code to include these corrections, but at the present accuracy we believe it would have a small effect. This conclusion is validated by the agreement between the limits we set with our Bayesian simplified time-scale comparison and the frequentist full trajectory analysis of Witte & Mummery (2025).<sup>4</sup>

## 2.2 Evolution of the free field

Due to equation (4), the SR instability causes the occupation number to grow exponentially,  $N_{\text{cloud}} \propto e^{\Gamma_{nlm} t}$ , where  $\Gamma_{nlm}$  denotes the ‘free’ SR rate  $\omega_{nlm}^I$  from equation (7). Eventually, however, the SR condition in equation (5) ceases to be valid and the cloud growth stops—even in the absence of self-interactions. By this time, the cloud will have extracted some spin  $\Delta a_*$  from the BH, which can be related to the cloud occupation number  $N_\Delta$  at the time when an amount of spin  $\Delta a_*$  has been extracted, assuming negligible initial occupation number of the cloud (named  $N_{\text{max}}$  in equation 8 Arvanitaki et al. 2015),

$$N_\Delta \approx 10^{76} \left(\frac{1}{m}\right) \left(\frac{\Delta a_*}{0.1}\right) \left(\frac{M}{10 M_\odot}\right)^2. \quad (10)$$

Equation (10) allows us to estimate whether or not a given cloud occupation level can be significantly reduce the observable spin of a BH. In particular, if we consider the exponential growth of the cloud for the  $|nlm\rangle$  level, some BH time-scale  $\tau_{\text{BH}}$  (further discussed in Section 3.1), and the SR time-scale  $\tau_{\text{SR}} = \Gamma_{nlm}^{-1}$ , we can only obtain constraints on ULB properties if

- (i)  $|nlm\rangle$  is superradiant according to equation (5),
- (ii) SR is fast enough to spin down the BH,  $\tau_{\text{SR}} < \tau_{\text{BH}} / \ln(N_\Delta)$ .

Note that the typical value of  $\Delta a_*$  adopted in equation (10) is of order the size of the observational errors on  $a_*$ . Since the condition on  $\tau_{\text{SR}}$  only logarithmically depends on  $N_\Delta$ , and thus  $\Delta a_*$ , even an order of magnitude difference in  $N_\Delta$  would not significantly change the derived limits. Still, since the choice of  $\Delta a_*$  is somewhat arbitrary, the BH evolution should be treated in more detail. This treatment goes beyond the present analysis, where we only compare time-scales.

## 2.3 Self-interactions

The effect of boson self-interactions has been considered in detail for the case of the dimensionless coupling  $\lambda$  from a  $\phi^4$  terms in the Lagrangian,

$$\lambda \equiv \frac{\mu^2}{4!} f^{-2}, \quad (11)$$

where  $f$  is the ULB decay constant, which is typically connected to the energy scale of new physics leading to the ULB’s existence (such as symmetry breaking or extra dimensions), while  $f^{-1}$  typically sets the scale of the ULB’s couplings to Standard Model particles. For a generic pseudo-Goldstone boson, equation (11) arises from Taylor expanding a generic potential with respect to the field periodicity  $f$ .

For a given value of  $\mu$ , experiments and astrophysical searches for ULBs typically place *upper* limits on the couplings or  $f^{-1}$ , while

<sup>4</sup>Witte & Mummery (2025) state that there appears to be an error in our analysis in the limit set at the low mass end. However, we have performed many consistency checks and believe the factor of two disagreement at low mass is entirely due to the statistical question being asked, i.e. Bayesian versus frequentist statistics.

BHSR places *lower* limits on  $f^{-1}$ , making the methodologies highly complementary. In the case of fuzzy DM, the relic abundance scales with  $f^2$  (Turner 1983, 1986; Marsh 2016), with future cosmological probes reaching  $f \sim 10^{16}$  GeV (see e.g. Bauer et al. 2021; Farren et al. 2022; Dvorkin et al. 2022), which overlaps with the BHSR sensitivity.

For QCD axions, the term in equation (11) arises from a Taylor expansion of the scalar potential in chiral perturbation theory, and di Cortona et al. (2016) found that  $\lambda = -0.346(22)$ . This prediction is possible since QCD axion models link  $\mu$  and  $f$ ,

$$\mu = \frac{\chi_{\text{top}}^2}{f} = (5.691 \pm 0.051) \text{neV} \left(\frac{10^{15} \text{GeV}}{f}\right), \quad (12)$$

where the QCD topological susceptibility  $\chi_{\text{top}}$  has been computed at NNLO (di Cortona et al. 2016; Gorghetto & Villadoro 2019). Note that these relations may not hold for a generic ALP. As we now discuss, self-interactions affect the SR rates.

### 2.3.1 The equilibrium regime

Initially, the ULBs will be near the vacuum and the rate  $\Gamma_{nlm}$  is the one for free bosons, as computed in Section 2.1. If the SR rate is faster than the inverse of the relevant BH time-scale  $\tau_{\text{BH}}^{-1}$ , the boson cloud can grow (see Section 2.2). Its subsequent evolution is determined by the strength of ULB self-interactions, for which there are two possibilities.

The first possibility is that the cloud enters an *equilibrium* between BHSR losses and level transitions caused by self-interactions and gravity (Baryakhtar et al. 2021). The BHSR rate in the equilibrium regime is effectively reduced compared to the vacuum rate by a factor of  $\eta_{\text{eq}} < 1$  for the equilibrium number density of bosons in the most superradiant level. In particular, the spin-down time-scale in the equilibrium regime is given by

$$\tau_{\text{eq}} = \frac{1}{\eta_{\text{eq}} \Gamma_{nlm}}, \quad (13)$$

where the factor  $\eta_{\text{eq}}$  can be computed following the procedure of Baryakhtar et al. (2021), who consider (transitions between) the  $|211\rangle$  and  $|322\rangle$  levels,

$$\eta_{\text{eq}} = \frac{2\sqrt{3}}{3} \frac{\sqrt{\Gamma_{211} \Gamma_{322 \times 322}^{211 \times \infty}}}{\Gamma_{211 \times 211}^{322 \times \text{BH}}}, \quad (14)$$

$$\Gamma_{322 \times 322}^{211 \times \infty} \approx 1.1 \times 10^{-8} \alpha^8 \mu f^{-4}, \quad (15)$$

$$\Gamma_{211 \times 211}^{322 \times \text{BH}} \approx 4.3 \times 10^{-7} \left(1 + \sqrt{1 - a_*^2}\right) \alpha^{11} \mu f^{-4}, \quad (16)$$

where  $\alpha$  is assumed to be sufficiently small and where  $f$  is the ULB decay constant in equation (11).<sup>5</sup>

Given values for  $(M, a_*)$ , we may exclude all ULB models for which, in addition to the conditions in Section 2.2,<sup>6</sup> the reduced SR rate in equilibrium is still faster than the typical BH time-scale. This can be encoded in the condition  $\tau_{\text{eq}} < \tau_{\text{BH}}$ .

<sup>5</sup>We will only consider the  $|211\rangle$  SR rate (and transitions to the  $|322\rangle$  level) for the equilibrium regime, making this a good assumption.

<sup>6</sup>The logarithm of  $N_\Delta$  in equation (10) is of similar size as the occupation number typically required to reach the equilibrium regime in the first place (Baryakhtar et al. 2021). We thus need to include such as condition to use the effective equilibrium rate in equation (13).

### 2.3.2 Bosenovae

The other possibility for the evolution is that a *bosenova* occurs (Arvanitaki & Dubovsky 2011; Kodama & Yoshino 2012; Yoshino & Kodama 2012; Arvanitaki et al. 2015; Yoshino & Kodama 2015). This happens when the occupancy of the cloud reaches a critical value  $N_{\text{crit}}$ , given by

$$N_{\text{crit}} \approx 10^{78} c_0 \left( \frac{M}{10 M_{\odot}} \right)^2 \left( \frac{f}{M_{\text{P}}} \right)^2 \frac{n^4}{\alpha^3}, \quad (17)$$

where  $c_0 \approx 5$  (equation 9 of Arvanitaki et al. 2015). Such an expression can be derived by comparing the order of magnitude size of different terms in the action (e.g. Mehta et al. 2021). Thus, in a single superradiant time-scale  $\tau_{\text{SR}}$  that would otherwise lead to a maximum occupancy  $N_{\Delta}$ , as given by equation (10), a number of bosenovae  $N_{\text{crit}}/N_{\Delta}$  occur. Since in each cycle the maximum occupancy is  $N_{\text{crit}}$ , the overall spin-down rate is reduced. This can be interpreted as an increase in the relevant time-scale (Arvanitaki et al. 2015):

$$\tau_{\text{BN}} = \frac{1}{\Gamma_{nlm}} \frac{N_{\Delta}}{N_{\text{crit}}} \ln(N_{\text{crit}}). \quad (18)$$

Given values for  $(M, a_*)$ , we may exclude all ULB models for which, in addition to the conditions in Section 2.2, the reduced SR rate from successive bosenovae is fast enough to reduce the spin. This can be written as  $\tau_{\text{BN}} < \tau_{\text{BH}}$ .

### 2.3.3 Bosenova or equilibrium?

In both the bosenova and equilibrium scenarios, bosonic self-interactions slow down BHSR. In the case of a bosenova, the cloud is prevented from reaching a high occupancy and exponential growth over the entire SR time-scale and instead goes through many short periods where a smaller amount of spin is extracted. Stronger self-interactions in this case decrease the value of  $N_{\text{crit}}$  and make each cycle shorter, thus extracting less spin overall. In the equilibrium case, stronger self-interactions lead to more rapid boson scattering in superradiant levels, consequently reducing their equilibrium number density, and thus their ability to extract spin from the BH.

Baryakhtar et al. (2021) approximated the full superradiant evolution by including the most relevant levels and found that an equilibrium is typically reached before a bosenova occurs. This finding was also confirmed by Omiya et al. (2023), who included relativistic effects. More concretely, for  $\alpha \lesssim 0.2$ , Baryakhtar et al. (2021) found that a two-level description is sufficient to model the BH evolution, the inclusion of higher levels does not affect their results, and a bosenova does not occur. For  $\alpha \lesssim 0.3$ , they found an equilibrium for the three-level system  $|211\rangle \leftrightarrow |322\rangle \leftrightarrow |411\rangle$  and argued that a bosenova is also not likely in this case. For  $\alpha \gtrsim 0.3$ , higher SR levels become important and solving the related system of rate equations becomes increasingly challenging. Since the focus of the present work is to improve the statistical framework, we leave such computations for future work.

In summary, the bosenova scenario appears to be excluded for  $\alpha \lesssim 0.2$ . In this regime, we may use the modified rates based on the results from Baryakhtar et al. (2021). For larger  $\alpha$ , the modified equilibrium and bosenova rates become less robust, introducing a systematic uncertainty from the assumptions underlying the related computations. In order to compare the possible impact of the two scenarios on the ULB constraints, our results in Fig. 4 will use the SR rates as computed in the previous sections for all values of  $\alpha$ . We stress again that the rates computed in the previous sections are

only robust for values of  $\alpha \lesssim 0.2$  and that the bosenova scenario is strongly disfavoured in that regime.

In contrast to the above, Yoshino & Kodama (2012) performed simulations of the full Klein-Gordon equation in  $3 + 1$  dimensions on a fixed Kerr space-time with a cosine, i.e. axion-like potential, for the scalar field. This method differs from the rate equation methods, and treats axion self-interactions beyond the quartic approximation. Yoshino & Kodama (2012) provided evidence that a bosenova occurs, rather than saturation. Aurrekoetxea et al. (2024) performed full numerical relativity simulations including back-reaction for binary BHs with a *complex* self-interacting scalar field, and also found evidence of a bosenova.

Finally, as already noted, BHSR for scalars has not been studied in fully non-linear  $(3 + 1)$ -dimensional numerical relativity due to the challenging time-scales involved. Given this fact, and the presence of historic disagreement on the bosenova issue, we leave open the possibility of bosenova and equilibrium and present results for both in our analysis. Comparing both bosenova and equilibrium models demonstrates how our method is easily adaptable to changes in physics assumptions and allows easy comparison of the effect of different theoretical models on ULB constraints.

## 2.4 Regge trajectories

As visualized by Arvanitaki et al. (2015, fig. 2), the BH evolution will be dominated by the fastest SR rate for which the SR condition in equation (5) is met. This is typically the  $|211\rangle$  level but, as the BH evolves and its mass and spin decrease, the SR condition will eventually cease to be fulfilled. Once this happens, ULBs from higher levels can still drain the BH’s mass and spin – as long as they meet the SR condition and the SR rate is fast enough to make the change observable. Eventually, there will be no such superradiant levels left, and the existence of ULBs makes a striking prediction: any given BH should not have a spin higher than a certain critical value  $a_*^{\text{crit}}$ , which depends on  $M$ ,  $\tau_{\text{BH}}$ ,  $\mu$ , and  $f$ . The line  $a_*^{\text{crit}}(M, \tau_{\text{BH}}, \mu, f)$  is referred to as ‘Regge trajectory’.

If we knew the true values of  $(M, a_*, \tau_{\text{BH}})$ , we could exclude the existence of any ULBs whose parameters predict  $a_*^{\text{crit}} < a_*$ . BHs on or below Regge trajectories might simply have been ‘born’ with a lower spin, and this lack of knowledge about the BH’s history means that they are compatible with the existence of the corresponding ULBs. In reality, we can only estimate the BH parameters, and thus only produce a statistical estimate for the exclusion probability of a given ULB (see Section 4).

## 3 BLACK HOLE DATA

We select two BH data sets as examples to demonstrate the methodology proposed in this work. One BH is the high-mass eclipsing X-ray binary (XRB) M33 X-7, which has frequently been studied in the SR literature. The other one is the SMBH IRAS 09149-6206, which possesses well-defined, independent mass and spin posteriors. It is also the first time that this SMBH has been used to infer ULB constraints. The spin of the XRB was inferred via the thermal fit continuum method, while the spin of the SMBH was inferred from X-ray reflection spectroscopy. We discuss these two techniques in Sections 3.2 and 3.3, respectively, and refer the reader to Bambi et al. (2021) and Reynolds (2021) for detailed, recent reviews.

### 3.1 BH time-scales

BHs can acquire angular momentum (be ‘spun up’) due to accretion. The characteristic time-scale for a BH to acquire a given mass through accretion at the Eddington limit is given by the Salpeter time,  $\tau_{\text{Salp}} \approx 4.5 \times 10^7$  yr, assuming a canonical radiative efficiency of 10 per cent. The Salpeter time is also the appropriate characteristic time-scale for BH spin-up from accretion, and thus often one simply chooses  $\tau_{\text{BH}} = \tau_{\text{Salp}}$ . Although super-Eddington accretion is possible, this approach is conservative. X-ray binaries accrete at sub-Eddington luminosities even in outburst (Bahramian & Degenaar 2023), with only the rarer sub-population of ultraluminous X-ray sources (Bachetti et al. 2014; King, Lasota & Middleton 2023) producing super-Eddington luminosities.

Indeed, the high spin measurements in XRBs present a puzzle: it is impossible for the XRB to acquire near-maximal spin via accretion over its lifetime, thus requiring high ‘natal’ spin (for M33 X-7, see discussion by Liu et al. 2008). For active galactic nuclei (AGNs), the vast majority of radiatively efficient systems accrete with Eddington fractions of  $\lambda_{\text{Edd}} \equiv L_{\text{bol}}/L_{\text{Edd}} \sim 0.01\text{--}1$  (Suh et al. 2015; Temple et al. 2023), where  $L_{\text{bol}}$  and  $L_{\text{Edd}}$  are the bolometric and Eddington luminosities, respectively. Indeed, super-Eddington accretion is generally found to be rare in both observational (Schulze & Wisotzki 2010; Zhang et al. 2024) and theoretical/simulation (Bustamante & Springel 2019; Shirakata et al. 2019; Sala et al. ) studies, especially at low redshift (however, see also Valiante et al. 2017; Farrah et al. 2022; Liu et al. 2021; Bennett et al. 2024).

An approximate current Eddington fraction can be calculated from the estimated bolometric luminosities and BH masses for each object, with reports of  $\lambda_{\text{Edd}} \approx 0.1$  for M33 X-7 (Liu et al. 2008) and  $\lambda_{\text{Edd}} \approx 0.2$  for IRAS 09149-6206 (GRAVITY Collaboration 2020). However, accreting BHs are variable sources, so to accurately estimate  $\tau_{\text{BH}}$  for individual sources would require detailed knowledge of the accretion history over tens of Myr. For the SMBH case (IRAS 09149-6206), we assume here that  $\tau_{\text{BH}} = \tau_{\text{Salp}}/\bar{\lambda}_{\text{Edd}}$ , where  $\bar{\lambda}_{\text{Edd}}$  can be thought of as an appropriately weighted time-averaged Eddington ratio. We set  $\bar{\lambda}_{\text{Edd}} = 0.1$ , which is likely to be fairly conservative, but note the inherent astrophysical uncertainty associated with this value. For the stellar mass BH M33 X-7, we instead use the system age (see Section 3.2.3). If  $\tau_{\text{SR}} = \Gamma^{-1} < \tau_{\text{BH}}/\ln N_{\Delta}$ , then SR dominates the evolution of the BH and  $a_*$  will decrease. Observing a BH in a stable configuration for which we can determine  $a_*$  and  $\tau_{\text{BH}}$  then allows us to exclude ULBs models that would predict  $\tau_{\text{SR}} < \tau_{\text{BH}}$ .

For completeness, note that the longest relevant time-scale is the age of the observable Universe, which is of order the Hubble time  $\tau_{\text{H}} = 1/H_0 = 14.5$  Gyr (Planck Collaboration VI 2020). Using  $\tau_{\text{H}}$  would result in the strongest possible limits, and it is thus a useful time-scale for investigating the addressable ULB parameter space. While the ‘age’ of a number of SMBHs – or, perhaps more accurately, the time since they accumulated most of their mass – could be of order  $\tau_{\text{H}}$ , as also suggested by studies with the *James Webb Space Telescope* (e.g. Curtis-Lake et al. 2023; Maiolino et al. 2024), the value of  $\tau_{\text{BH}}$  is rather set by accretion and other processes in the BH evolution, as discussed above.

### 3.2 X-ray binary BHs and M33 X-7

X-ray binaries consist of a stellar mass BH accreting matter from a companion star. XRBs are classified as low-mass (LMXB) or high-mass (HMXB), depending on whether the companion is lower

or higher mass than the BH, respectively. The accretion process emits electromagnetic radiation, which can be detected using X-ray telescopes. Analysing the data from such measurements yields estimates for  $(M, a_*)$  with, in some cases, relatively high statistical precision but potentially large systematic uncertainties due to the modelling of the system.

#### 3.2.1 Mass estimates in stellar-mass BHs

As reviewed by Casares & Jonker (2014), the most robust method to estimate BH masses in XRBs obtains a *dynamical mass measurement*, using Kepler’s laws of motion and the presence of a stellar companion. This method is more direct and robust than the methods used for the majority of BHs because the dynamical impact of the BH on its companion can be measured. From spectroscopy, radial velocities of the emission or absorption lines from the companion star can be obtained, which allows a measurement of the orbital period  $P_{\text{orb}}$  and the radial velocity semi-amplitude  $K_c$ . The so-called mass function  $F(M)$  can then be defined, which explicitly relates the two radial velocity parameters and  $M$ ,

$$F(M) = \frac{K_c^3 P_{\text{orb}}}{2\pi G} = \frac{M \sin^3(\theta_i)}{(1+q)^2}, \quad (19)$$

where  $\theta_i$  is the binary inclination and  $q$  is the mass ratio. For known inclination and  $q$ , one can estimate  $M$ , whereas for unknown inclination  $F(M)$  provides a firm lower limit on the mass (since  $1+q > 1$  and  $\sin \theta_i \leq 1$ ).

Dynamical mass measurements are easier to obtain for LMXBs, because the systems typically go through outburst cycles and have long periods of very low accretion activity (quiescence). During quiescent periods, optical spectroscopy can be used to accurately characterize the radial velocity curve of the companion star, as has been carried out for about 20 LMXBs (Corral-Santana et al. 2016; Chaty 2022). In HMXBs, the accretion state phenomenology is more complex; of the over 100 known HMXBs, many of which host neutron stars, only a handful of dynamical mass measurements have been possible for BH systems. These include famous sources such as Cygnus X-1, and the source studied here, M33 X-7. In HMXBs hosting BHs, the X-ray source is persistent and true quiescent periods do not occur. However, as HMXB companion stars are, by definition, rather massive, they outshine the accretion disc in the optical band, making spectroscopy of the donor star possible. There are still various systematic effects that can complicate the dynamical BH mass measurements, the main three factors being line formation in (or contamination by) the strong stellar winds of the companion, uncertainties in the donor star mass, and unknown Roche lobe filling factors (see discussion by Casares & Jonker 2014). Fortunately, M33 X-7 is an eclipsing system, which mitigates some of the effects mentioned above. In addition, it would be possible in future to incorporate uncertainties on, e.g. the donor mass, Roche lobe filling factor, and inclination as additional nuisance parameters into the analysis.

#### 3.2.2 XRB spin estimates from the continuum fitting method

In addition to the X-ray reflection method (see Section 3.3.2), the continuum-fitting method has widely been used to estimate BH spins in XRBs. This relies on the influence of BH spin on the inner edge of the accretion disc, as higher (prograde) spins lead to an innermost stable circular orbit (ISCO) lying closer to the BH. The closer the ISCO radius ( $R_{\text{ISCO}}$ ) is to the event horizon (which depends on  $M$

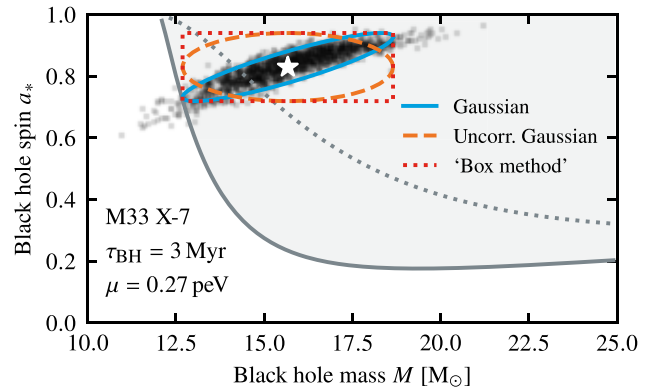
and  $a_*$ ; cf. e.g. equations 2–4 in Reynolds 2021), the more energy will be released per unit accreted matter. The maximum effective temperature of the accretion disc corresponds to that at the ISCO and is set by  $T_{\text{eff,max}}^4 \propto \lambda_{\text{Edd}} M R_{\text{ISCO}}^{-3}$  in the thin-disc model. For a canonical 10 per cent radiative efficiency,  $T_{\text{eff,max}}$  peaks in the soft X-ray band (i.e. below 2 keV) for an accreting stellar BH of  $M = 10 M_{\odot}$ . In contrast,  $T_{\text{eff,max}}$  peaks in the extreme UV band for an accreting SMBH with  $M = 10^6 M_{\odot}$  for the same canonical radiative efficiency. The interpretation of observed X-ray spectra of XRBs accreting at moderate rates relative to the Eddington limit ( $\lambda_{\text{Edd}} = 0.01 - 0.3$ ), paired with suitable accretion disc models, has provided spin constraints for around 10 XRBs. We refer to Steiner et al. (2009), McClintock, Narayan & Steiner (2014), and Reynolds (2021) for relevant reviews.

Generally, the accretion disc model assumed in continuum-fitting spin studies is the geometrically thin, optically thick, steady-state accretion disc model due to Novikov & Thorne (1973). This model is an extension to the also well-known Shakura–Sunyaev alpha thin-disc model (Shakura & Sunyaev 1973) to the relativistic regime. In the Novikov–Thorne model, mass loss due to winds in the inner disc is assumed to be negligible, and, as in other thin-disc models, the disc acquires non-molecular viscosity due to internal stresses and heat is dissipated locally. These two conditions have fairly sound theoretical and empirical motivation, given that thin-disc models can describe the putative disc component of the observed X-ray spectrum in XRBs during soft, thermally dominated accretion states. Furthermore, fully relativistic continuum-fitting models generally adopt a zero-torque boundary condition at the ISCO.

Continuum-fitting models predict local thermal spectra by combining the predicted thermal disc spectrum – subject to an accretion disc model – with the temperature-dependent colour correction factor. Although the temperature dependence of the latter can be inferred from sophisticated radiative transfer calculations (Davis et al. 2005), in practise, some continuum-fitting models implement an approximate temperature-dependent colour correction factor, such as that described by Done et al. (2012).

Spin estimates in XRBs from continuum fitting have been achieved with high statistical precision. However, the systematic errors emerging from the assumptions and approximations made by these models must be accounted for. First, as mentioned above, the thin-disc solution is only expected to be appropriate for thermally-dominated discs for bolometric luminosities  $L_{\text{bol}} \lesssim 0.3 L_{\text{Edd}}$  (Kulkarni et al. 2011). Secondly, even in this moderate luminosity regime, radiation could be emitted from within the ISCO, which is expected to induce an uncertainty in the recovered spin from continuum fitting of  $\sigma_{a_*} = 0.05$  (Zhu et al. 2012). Third, if the aforementioned zero-torque boundary condition at the ISCO is not satisfied, e.g. due to magnetic fields threading the inner disc, further systematic errors of  $\sigma_{a_*} \leq 0.1$  would be induced (Shafee et al. 2008b; Shafee, Narayan & McClintock 2008a).

Moreover, astrophysical parameters with inherent uncertainties – such as the inclination to the angular momentum vector of the binary, the BH mass, and the distance and accretion geometry of the XRB system – are also expected to contribute to the overall systematic uncertainty on the inferred BH spin. These nuances have led to variations in the analysis pipelines employed by different research groups amongst different data sets, potentially leading to further systematics from modelling and different statistical methods. Nevertheless, a consistent analysis of multiple XRBs would in principle be possible, thanks to publicly available tools such as the XSPEC spectral-fitting package (Arnaud 1996).



**Figure 2.** The sampled BH mass and spin distribution (black points) of M33 X-7 and Regge trajectories of the non-interacting (grey lines and shaded region) and self-interacting  $|211\rangle$  level (dotted, grey line; equilibrium regime,  $f^{-1} = 2 \times 10^{-15} \text{ GeV}^{-1}$ ,  $\mu = 0.27 \times 10^{-12} \text{ eV}$ , and  $\tau_{\text{BH}} = 3 \times 10^6 \text{ yr}$ ). The full distribution, whose mean is marked with a white star, is compared to its  $2\sigma$  error bars for an uncorrelated (dashed, orange line) and full (blue line) Gaussian, and a ‘ $2\sigma$  box’ (dotted, red line).

### 3.2.3 M33 X-7 posteriors

In this work, we use the spin posterior of the XRB system M33 X-7 of Liu et al. (2008), inferred using a fully relativistic continuum-fitting model. The analysis in Liu et al. (2008) was applied to a set of *Chandra* and *XMM-Newton* observations of the binary in the thermally dominated state and applied to epochs where  $L_{\text{bol}} < 0.3 L_{\text{Edd}}$ . Based on previously determined system inclination and BH mass values, Liu et al. (2008) constrained the BH spin value to  $a_* = 0.84 \pm 0.05$ . The authors noted that the assumptions made by their fully relativistic continuum-fitting model, added to the assumption that the angular momentum vector of the binary is approximately aligned with the BH spin vector, are expected to induce additional systematic uncertainties.

In Fig. 2, we show the parameter sampling distributions from bootstrapping (black dots) for the BH system M33 X-7, which we obtained by digitizing fig. 3 of Liu et al. (2008) (see also Orosz et al. 2007). Clearly, the  $2\sigma$  contour of an uncorrelated Gaussian (dashed, orange line) does not provide a good description of the distribution, and including correlations (blue line) is more appropriate. We also show the ‘box’ defined by the uncorrelated  $2\sigma$  errors, which captures all possible underlying correlations at the price of discarding the information contained therein. We will discuss the relative importance of these effects on the ULB limits in Section 5.1.

Note that, in other cases, the resulting  $(M, a_*)$  distribution may be highly non-Gaussian. For instance, the somewhat uniformly distributed posterior distribution of GRS 1124-683 (Chen et al. 2016) provides an extreme example.

Fig. 2 also shows the Regge trajectory (solid, grey line and shading; cf. Section 2.4) for the non-interacting  $|211\rangle$  level and an ULB mass of  $\mu = 0.15 \times 10^{-12} \text{ eV}$ . Since the estimated age  $\tau_{\text{age}} = 3 \times 10^6 \text{ yr}$  of M33 X-7 (Gou et al. 2011) is below  $0.1 \tau_{\text{Salp}}$ , we choose  $\tau_{\text{BH}} = \tau_{\text{age}}$ . To illustrate the effect of self-interactions, in this case assuming the equilibrium regime and  $f^{-1} = 4 \times 10^{-16} \text{ GeV}^{-1}$ , we also include the modified Regge trajectory as a dotted, grey line. We can exclude ULBs if most of the  $(M, a_*)$  samples can be found above the Regge trajectory.

### 3.2.4 Effects of the companion star

The companion of M33 X-7 has a mass of  $M_2 \approx 70 M_\odot$  (Valsecchi et al. 2010) and an orbital period of  $T = 3.45$  d (Pietsch et al. 2006). Using the equations from appendix I of Baryakhtar et al. (2021), we verified that the companion does not have a significant effect on the exclusion regions for the ULB parameters.<sup>7</sup> For other systems, however, the companion might play a more important role. In such cases, our proposed statistical framework can include and propagate the related uncertainties on the companion parameters, which would be particularly relevant for the large uncertainty of about 10 per cent on  $M_2$  (Valsecchi et al. 2010, see Section 6 for other such nuisance parameters).

## 3.3 Supermassive BHs and IRAS 09149-6206

Supermassive black holes are believed to reside at the centres of all galaxies, and a compelling way to study them observationally is by probing the activity of their associated AGNs. AGNs are nuclear regions located at the centre of about 1–10 per cent of galaxies. Their integrated luminosity exceeds the stellar luminosity of their host galaxies as a consequence of the accretion onto the SMBH from the surrounding accretion disc. Given the high-quality observations needed, obtaining spin and mass estimates for SMBHs is rather challenging. Here, we summarize how relativistic X-ray reflection and infrared interferometry can, respectively, be used to constrain the spin and mass of SMBHs in AGNs.

### 3.3.1 SMBH mass estimates

The mass of Sgr A\*, the closest SMBH to Earth, can be measured dynamically from the motion of stars within the BH gravitational potential (Ghez et al. 2008; Genzel, Eisenhauer & Gillessen 2010). However, in more distant AGNs, the BH's sphere of influence cannot be resolved and its mass is often estimated from broad emission lines: under the assumption that the broad-line region (BLR) gas is virialized, the width of the line,  $\Delta v$  is related to the BH mass through the relation  $M = f R_{\text{BLR}} (\Delta v)^2 / G$ , where  $f$  is the so-called virial factor,  $f \sim \mathcal{O}(1)$ , that accounts for BLR geometry and projection effects. Using this virial estimator requires an estimate for the line formation radius  $R_{\text{BLR}}$ , which ideally would be estimated from reverberation lags (Kaspi et al. 2000, 2005). If  $R_{\text{BLR}}$  is not known, then it is usually estimated from the luminosity-size relation,  $R_{\text{BLR}} \propto L^{1/2}$  – where  $L$  is an appropriately chosen proxy for the bolometric or ionising luminosity – calibrated with the low-redshift empirical relation from reverberation mapped AGNs (e.g. Bentz et al. 2013). There are numerous other ways of estimating BH mass, which we cannot describe exhaustively, each with their own characteristic uncertainties and systematics; notable methods include the use of scaling relations between  $M$  and the stellar velocity dispersion (Ferrarese & Merritt 2000; Gebhardt et al. 2000) or bulge mass (Magorrian et al. 1998), and the observation of H<sub>2</sub>O megamasers (Greene et al. 2010).

Recently, the GRAVITY interferometric beam-combiner (GRAVITY Collaboration 2017) on the *Very Large Telescope* has made it possible to estimate  $R_{\text{BLR}}$  and model the BLR geometry from the differential phases in infra-red interferometry. These measurements are extremely useful for BH masses and have allowed more accurate

<sup>7</sup>This conclusion is consistent with those reached by Witte & Mummery (2025).

estimates for a number of luminous AGNs (Gravity Collaboration 2018, 2021). Additionally, while subject to various assumptions, the modelling provides a well-defined posterior probability distribution for  $M$  for our analysis. One such object observed is IRAS 09149-6206, which is our chosen SMBH case study (see Section 3.3.3). Our use of the mass posterior from GRAVITY observations mitigates some, but not all, of the systematic uncertainties associated with BH mass estimates.

### 3.3.2 SMBH spin estimates from X-ray reflection spectroscopy

In this work, we adopt the BH spin posterior distribution for IRAS 09149-6206 inferred by Walton et al. (2020), who used the X-ray reflection method. In the absence of Compton-thick absorption and of spectral degeneracies induced by warm absorbers and ultra-fast outflows, the detailed modelling of reflected features in X-ray spectrum of AGNs provides a powerful pathway to probing the strong gravity regime in their central SMBHs. We note that several alternative observational techniques have been used to generally probe spins of BHs accreting at lower rates or hosting strong radio-jets (Tamburini, Thidé & Della Valle 2020; Ünal & Loeb 2020; Wen et al. 2021; Daly 2022; Ricarte et al. 2023). We refer to Davoudiasl & Denton (2019) for a discussion on BHSR from the spin estimate of M87\* from the interpretation of *Event Horizon Telescope* data; and to Ng et al. (2021b) for joint constraints on the ULB mass and spin distribution at the time of binary BH formation from the Gravitational-Wave Transient Catalog (GWTC-2).

The disc reflection spectrum arises from the reprocessing of the primary X-ray continuum from the inner and outer regions of the ionized accretion disc, as well as from distant neutral and non-relativistic material surrounding the accretion disc. The primary continuum is thought to originate from the Compton up-scattering of seed photons in the innermost regions of the disc by the electrons in the hot X-ray corona. A fraction of the direct coronal emission towards the observer will back-scatter on the disc and be re-radiated away, with the non-thermal emission of coronal electrons governing the high-energy cutoff tail of the broadband spectral energy distribution (usually in the 100–300 keV range) in type-1 AGNs.

Amongst other features in the X-ray reflection spectrum, the most prominent one is the Fe K  $\alpha$  line, as iron is the most abundant element in the disc and has a high fluorescent yield (George & Fabian 1991). Generally, the Fe K  $\alpha$  line receives contributions from reprocessing from the inner regions of the disc, as well as cold, non-relativistic matter surrounding the immediate vicinity of the AGN. A narrow Fe K  $\alpha$  reflection feature (centred at 6.4 keV in the rest frame) often arises due to the reflection from cold matter surrounding the outer accretion disc. However, the most important feature for BH spin inference is the broad Fe K  $\alpha$  feature, centred around 6.4–6.97 keV rest-frame energy, which likely indicates the presence of reprocessed emission from the innermost regions of the accretion disc.

Given that this reprocessing is thought to take place in close proximity to the BH event horizon, the reflected emission is subject to standard Doppler and general relativity effects in the disc, with the latter including light bending and gravitational redshift. The spectral imprints emanating from these effects, especially a distinct red wing to the line, will be more dramatic and noticeable the closer the emission takes place to the event horizon (Fabian et al. 1989). In addition, relativistic reflection from the innermost regions of the disc will also induce a forest of reflection features down to the soft X-ray band (below 2 keV). Finally, another characteristic feature of the disc reflection spectrum in AGNs is the Compton (reflection)

hump, which arises in the range of 20–40 keV due to photo-electric absorption of low-energy photons and the multiple reprocessing of high-energy (coronal) electrons (Zdziarski et al. 1990).

Current relativistic X-ray reflection models can provide a powerful probe of the BH spin if the inner accretion disc extends down to  $R_{\text{ISCO}}$ . Under this condition and assuming an accretion disc model, these models can then be used to make predictions of reflected spectra parametrized by  $a_*$ , given assumptions about

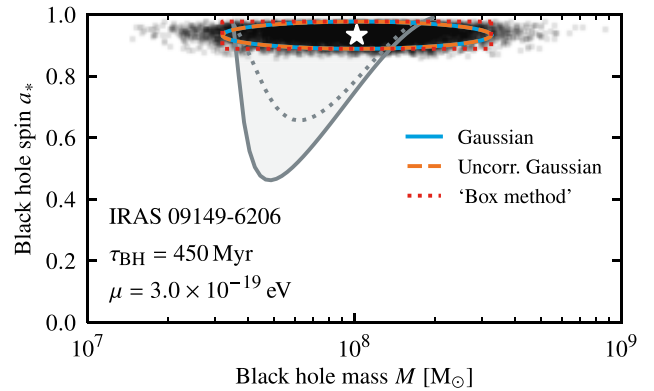
- (i) the geometry of the corona,
- (ii) the disc density at the mid-plane and/or disc ionization state,
- (iii) and the metric of space–time around the central BH.

State-of-the-art X-ray reflection models (such as the widely used `relxill` family of models, described by García et al. 2014; Dauser et al. 2014, 2016) have provided spin constraints for about 20 stellar-mass BHs and 50 SMBHs (Reynolds 2021). The simplifications made by these reflection models, and their possible translation into systematic biases on the inferred BH spin were thoroughly reviewed by Reynolds (2021, section 3.2) and Bambi et al. (2021, sections 3 and 4). We also note that, although the majority of current spin estimates drawn from X-ray reflection spectroscopy rely on the detection and interpretation of the broadened Fe K  $\alpha$  line, several SMBH spins in AGNs have been inferred via the reflection interpretation of the soft excess – notably, in Narrow-line Seyfert 1 galaxies (Jiang et al. 2019; Mallick et al. 2022).

### 3.3.3 IRAS 09149-6206 posterior samples

In this work, we make use of spin posteriors from a specific analysis of IRAS 09149-6206 by Walton et al. (2020), who analysed the broadband X-ray data from the *Swift*, *XMM-Newton*, and *NuSTAR* telescopes. As is commonly the case, this AGN is found to have a complex X-ray spectrum with signatures of X-ray reflection from the inner accretion disc as well as both neutral and ionized absorption by more distant circumnuclear matter along the line of sight. Walton et al. (2020) construct a 23-parameter spectral model to describe the primary continuum reflection from a BH accretion disc, which then passes through multiple layers of absorption (see table 4 of Walton et al. 2020). Using that model, and the multisatellite data set, the authors computed the posterior distribution of the parameters with the Goodman–Weare algorithm.

Marginalizing over all other parameters of this global model, the inferred BH spin is  $a_* = 0.94^{+0.02}_{-0.07}$ . The principal ‘nuisance’ parameters that have bearing on the spin are the inclination of the accretion disc, the iron abundance of the disc, the ionization state of the surface layers of the disc, and the geometry of the X-ray-emitting corona (characterized as a ‘lamppost’ height, also known as the coronal geometry). For typical AGN parameters, the inclination of the disc is readily determined by the high-energy ‘edge’ of the iron line profile, resulting in little degeneracy with the inferred spin (that results from the shape of the low-energy ‘tail’ of the iron line). The iron abundance principally influences the relative strength of the iron line and the Compton reflection hump with secondary impacts on the inferred spin, but is well constrained in data sets that cover both features (such as is the case here). The ionization state of the accretion disc photosphere strongly influences the shape of the low-energy ( $E < 4$  keV) reflection spectrum and shifts the energy of the iron line but, again, is well constrained by data sets that cover the 0.5–50 keV band. It is the lamppost height that has the greatest degeneracy with the inferred spin; this is included in the quoted errors on the spin measurement (see figs A1–A2 of Walton et al. 2020, for the associated corner plots).



**Figure 3.** The sampled BH mass and spin distribution (black points) of IRAS 09149-6206 and Regge trajectories of the non-interacting (grey lines and shaded region) and self-interacting |211) level (dotted, grey line; equilibrium regime,  $f^{-1} = 3 \times 10^{-16} \text{ GeV}^{-1}$ ,  $\mu = 3 \times 10^{-19} \text{ eV}$ , and  $\tau_{\text{BH}} = 4.5 \times 10^8 \text{ yr}$ ). The full distribution, whose median is marked with a white star, is compared to its  $2\sigma$  error bars for an uncorrelated (dashed, orange line) and full (blue line) Gaussian, and a ‘ $2\sigma$  box’ (dotted, black lines).

Our mass posteriors come from the BLR modelling in the analysis of interferometric data by the GRAVITY Collaboration (2017, 2020). This mass estimate is consistent with the estimate from timing by Walton et al. (2020).

We show the IRAS 09149-6206 posterior samples for  $M$  (GRAVITY Collaboration 2017, 2020) and  $a_*$  (Walton et al. 2020) in Fig. 3, assuming a BH time-scale of  $\tau_{\text{BH}} = \tau_{\text{Salp}}/0.1 = 4.5 \times 10^8 \text{ yr}$ . Note that, due to the two independent measurements for  $M$  and  $a_*$ , there are no correlations between the two parameters in this case. As for M33 X-7, we again include the Regge trajectories for non-interacting (solid, grey line) and self-interacting (dotted, grey line) ULBs.

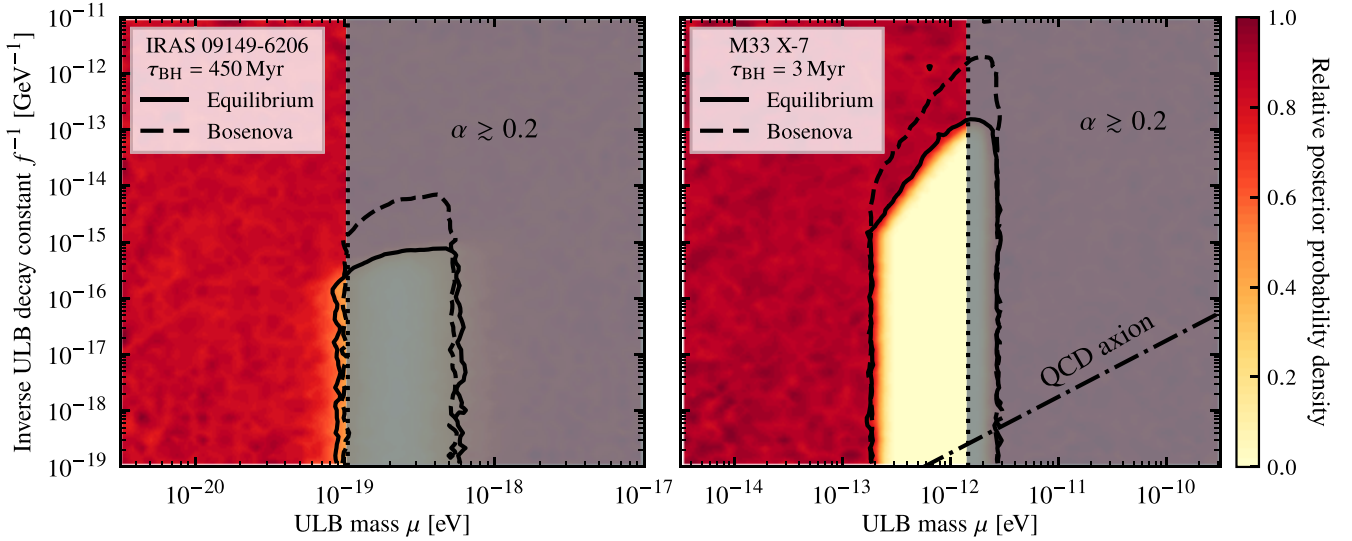
Since IRAS 09149-6206 has a spin close to the theoretical maximum value, the ‘ $2\sigma$  region’ of the  $a_*$  distribution surpasses that value. Instead of a regular Gaussian, one could use a truncated Gaussian to implement this physical upper limit on  $a_*$ .

We also checked that, as is presumably the case for many SMBHs, the  $M$  distribution is better described by a Gaussian in  $\ln(M/M_\odot)$  than in  $M$ . The mean value of  $\ln(M/M_\odot)$  corresponds to the median for the lognormal distribution of  $M$ , and Gaussian error propagation,  $\sigma_M/M \approx |\sigma_{\ln(M/M_\odot)}|$ , is not necessarily a good approximation. Converting quoted errors on  $\ln(M/M_\odot)$  into an error on  $M$  is thus potentially problematic – ven more so when considering the ‘ $2\sigma$  region’. We will see the effect of this in Section 5.1.

### 3.4 Gravitational-wave binary BH mergers

Gravitational-wave (GW) data can be used to estimate the mass and spin of binary black hole (BBH) mergers. In principle, these data can thus be used to constrain ULB properties.

However, current detectors only measure a mass-weighted combination of the spins along the orbital angular momentum (Ng et al. 2018). This complicates the determination of the individual spins at the time of the merger, while the unobservable pre-merger history also introduces additional challenges (Ng et al. 2021a; Baryakhtar et al. 2021). These caveats apply to ULB constraints from BBH mergers based on the inferred pre-merger spins alone, as e.g. presented by Stott & Marsh (2018), Fernandez, Ghalsasi & Profumo (2019), and Stott (2020).



**Figure 4.** Normalized posterior probability density distributions for  $(\mu, f^{-1})$  for IRAS 09149-6206 (left) and M33 X-7 (right). We show the 95 per cent credible regions at highest posterior density for both the equilibrium (solid line) and bosenova (dashed lines) scenarios and, in the right panel, the QCD axion model line (dash-dotted line) predicted by equation (12). Note that we highlight (grey shading) the  $\mu$  region where more than 5 per cent of BH mass samples imply  $\alpha > 0.2$  (dotted line) and, meaning that the computation of the SR rate involves large theoretical uncertainties and may not be valid.

To ameliorate these issues, Ng et al. (2021a, 2021b) and Cheng, Zhang & Bao (2023) propose using a Bayesian hierarchical analysis to account for the pre-merger history using hyperprior distributions. Cheng et al. (2023) in particular investigate the influence of prior choices and consider an average of possible BBH lifetimes in the range of  $10^6$ – $10^{10}$  yr. Since the constraints are mainly driven by a few of the observed BBH systems, the authors conclude that more data is needed for more robust conclusions once a broad range of pre-merger scenarios is included in the analysis. We thus leave the including of GW data from BBH mergers for future work.

#### 4 STATISTICAL FRAMEWORK

With the BHSR rates from Sections 2.3.1 and 2.3.2, and the posterior distributions for the BH parameters from Sections 3.2 and 3.3, we can now introduce our procedure for computing exclusion regions for ULB models.

As discussed in Section 2.4, we can exclude ULB models<sup>8</sup>  $\alpha = (\mu, f^{-1})$  for which the *true* values of  $\beta = (M, a_*, \tau_{\text{BH}}, \dots)$  lie above the Regge trajectory, i.e. when  $a_* > a_*^{\text{crit}}(M, \tau_{\text{BH}}, \dots, \alpha)$ . The ellipsis in  $\beta$  is to indicate that there may be more parameters of the BH environment that have an influence on the Regge slope. For this work, however, we will not consider any such parameters and assume  $\tau_{\text{BH}}$  to be a fixed value. Also note that any other BH nuisance parameters, such as the inclination angle  $i$  or virial factor  $f$ , can already be marginalized out in the original BH analysis.

In terms of probabilities, we can write the condition on  $a_*$  as a Heaviside function,  $p(a_*|M, \alpha) = \Theta(a_*^{\text{crit}} - a_*)$ , i.e.  $p(a_*|M, \alpha) = 0$  if the BH lies above the Regge slope, and otherwise  $p(a_*|M, \alpha) = 1$ . In particular, the BH’s parameters only depend on the ULB parameters when SR is active and the condition on  $a_*$  applies, such that the conditional probability  $p(\beta|\alpha)$  factorizes as  $p(\beta|\alpha) = p(\beta)p(a_*|M, \alpha)$ .

<sup>8</sup>Thanks to equation (12), QCD axion models effectively only have one free parameter (either  $\mu$  or  $f^{-1}$ ), in addition to the (well-constrained) topological susceptibility  $\chi_{\text{top}}$  as a nuisance parameter.

While we do not know the true values of  $\beta$ , we can obtain a sampling distribution from the BH data  $\mathcal{D}$  to compute the posterior  $p(\alpha|\mathcal{D})$ . Using the law of conditional probabilities and Bayes’ theorem, we find that

$$p(\alpha|\mathcal{D}) = \int p(\alpha, \beta|\mathcal{D}) d\beta = \int p(\beta|\mathcal{D}) p(\alpha|\beta, \mathcal{D}) d\beta \quad (20)$$

$$= \int p(\beta|\mathcal{D}) \frac{p(\beta|\alpha) p(\alpha)}{p(\beta)} d\beta \quad (21)$$

$$= p(\alpha) \int p(\beta|\mathcal{D}) p(a_*|M, \alpha) d\beta \quad (22)$$

$$\approx \frac{p(\alpha)}{N} \sum_{i=1}^N p(a_*^i|M^i, \alpha), \quad (23)$$

where we use a slash to indicate the removal of the redundant dependence on the data  $\mathcal{D}$  and where, in the last step, we assume that the BH analysis provides  $N$  equally weighted samples  $\beta^i = (M^i, a_*^i)$  from the posterior  $p(\beta|\mathcal{D})$ , which we can use to compute the integral in equation (22) via Monte Carlo integration.

Since  $p(\alpha|\mathcal{D}) \propto p(\alpha) p(\mathcal{D}|\alpha)$ , we may re-interpret the sum in equation (23) as a marginal likelihood. Furthermore, the computation of  $p(a_*|M, \alpha)$  requires a computation of the BHSR rates and assumptions about the BH evolution. As we will discuss in Section 5.2, both our computations and assumptions could be refined further. Even then, the framework presented here is general enough to allow for such modifications without having to change to the framework itself.

#### 5 RESULTS AND DISCUSSION

Our main results can be found in Fig. 4, where we show the posterior distribution for  $(\mu, f^{-1})$  for IRAS 09149-6206 (left panel) and M33 X-7 (right panel).

We assume log-uniform prior distributions for  $\mu$  and  $f^{-1}$ , limited by the ranges shown in Fig. 4. The ranges approximately correspond to the relevant (observable) parameter space, except that we restrict  $f < M_p$ . The latter is expected in string theory, and more generally if

the effective field theory is free of corrections from quantum gravity (Arkani-Hamed et al. 2007).

We use EMCEE (Foreman-Mackey et al. 2013) with 48 walkers to generate, after a burn-in run, 60 000 samples per chain i.e. a total of around  $10^7$  samples (total runtime around 600 CPUh). Finally, we compute the 95 per cent credible regions (CRs) at highest posterior density (HPD) for the ULB parameters.

We can see from Fig. 4 that the equilibrium regime and the bosonova condition lead to similar disfavoured regions of parameter space in  $\mu$ . In particular, the lower  $\mu$  cutoff is set by the condition to achieve a significant cloud occupation while, at larger  $\mu$ , the limits are cut off by the SR condition. It has been noted before, and we will see this explicitly in Section 5.1, that the range of excluded  $\mu$  may be extended by including higher occupation levels with  $n > 2$ . However, here we only show results considering the  $|211\rangle$  since our computed SR rates are most reliable for this case. To further emphasize this, we also show the  $\mu$  value for which we mostly, i.e. for at least 95 per cent of the  $M$  samples for a given BH, obtain  $\alpha < 0.2$ . The ULB constraints are most robust in this regime where, moreover, the occurrence of bosonovae is strongly disfavoured (see Section 2.3.3).

From the relative posterior densities, it appears that the constraints on IRAS 09149-6206 are weaker than for M33 X-7, which we will explore in more detail in Section 5.1. This highlights the relevance of a precise determination of SMBH masses to obtain meaningful constraints. While only BHs with a high spin can give strong limits on ULB parameters, the connection between  $M$  and  $\mu$  via  $\alpha = M\mu$  implies that we can only locate these constraints if  $M$  is measured precisely enough. Large uncertainties on  $M$  prevent us from doing so and, although it is clear that such a BH would give strong constraints on ULBs, an improvement in the BH mass measurement is required to tell us about the corresponding  $\mu$  range. The primary and secondary criteria for suitable BH candidates for obtaining ULB constraints should thus be a high  $a_*$  and precisely determined  $M$ , respectively.

Note that the adopted priors implicitly assume the existence of ULBs within their possible parameter range, and that we can always define CRs of the posterior distribution to exclude part of the ULB parameter space. Such constraints are only meaningful if the prior is sensible – i.e. in particular *if such ULBs exist*. In order to assess the probability of ULBs’ existence, one has to perform a model comparison, e.g. using Bayes’ factors. Here, we are interested in setting limits, and thus do not perform this additional step.

More generally, the choice of prior on  $(\mu, f^{-1})$  will necessarily affect the exclusion regions. Hoof et al. (2019) analysed this in more detail for axion models and found that the prior dependence can be very strong. Still, at least with regards to changing the range of the log-uniform priors, we do not expect extreme deviations as for e.g. the choice between log-uniform and uniform priors. In particular, when combining data from multiple BHs, the overall prior range will increase when considering the ‘sensitive  $\mu$  region’ for all BHs, making a joint fit increasingly conservative.

We also want to stress that the marginal distributions on  $\mu$  and  $f^{-1}$  will not be very constraining:<sup>9</sup> there is practically no limit

on  $f^{-1}$  alone due to the available  $\mu$  parameter space with high posterior density outside of the region related to the inferred  $M$  range. Similarly, the region of  $\mu$  that can be constrained when marginalizing  $f^{-1}$  is much narrower than suggested by the two-dimensional exclusion region. When considering self-interactions, it is important to show the two-dimensional contour instead of assuming that the excluded one-dimensional  $\mu$  range will correspond to the range suggested by the two-dimensional contour at low  $f^{-1}$ . This would even be more so the case in a frequentist framework.

## 5.1 Comparison with other works

Surveying the literature, we identify two main strategies to obtain ULB constraints, which we refer to as the ‘box method’ and ‘Gaussian distance method’, respectively.

### 5.1.1 The ‘box method’

Arvanitaki et al. (2015) applied the following checks for BH parameter estimates  $(\hat{M}, \hat{a}_*)$  with associated uncertainties  $(\sigma_{\hat{M}}, \sigma_{\hat{a}})$  and  $\tau_{\text{BH}} = \min(\tau_{\text{age}}, 40 \text{ Myr})$ . A given ULB model can be excluded if at least one  $|nlm\rangle$  with  $l \leq 5$  (i.e.  $n \leq 6$ ) and all  $M \in [\hat{M} - 2\sigma_{\hat{M}}, \hat{M} + 2\sigma_{\hat{M}}]$  and  $a = \hat{a} - 2\sigma_{\hat{a}}$  fulfil the SR and bosonova conditions (see Section 2.3.2). Comparing to Figs 2 and 3, these conditions demand that the rectangle defined by the mass and spin measurement  $\pm 2\sigma$  in each direction is contained within the envelope of the Regge trajectory, and that bosonovae do not spoil the spin-down of the BHs. Baryakhtar et al. (2021) also used the ‘box method’, but based on the equilibrium scenario and conditions (see Section 2.3.1).

### 5.1.2 The ‘Gaussian distance method’

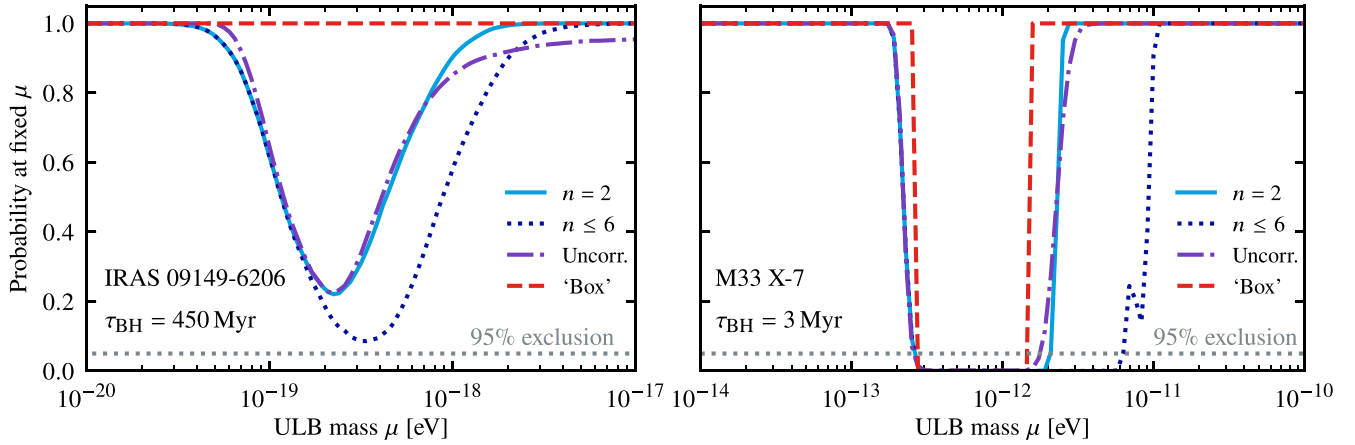
Stott & Marsh (2018) and Stott (2020) proposed to interpret  $(\hat{M}, \hat{a}_*)$  estimates and their errors as uncorrelated Gaussian measurements. They then defined ‘effective errors’ with respect to the Regge trajectory, which allowed them to compute a distance measure between the uncorrelated Gaussian distribution of the measurement and the area above the Regge trajectory. Their goal was to compute the weighted overlap between Gaussian distribution and the Regge trajectory ‘in a numerically efficient manner’, as visualized in Stott & Marsh (2018, fig. 17). In this sense, their underlying logic is similar to ours with the difference being their assumption about the uncorrelated Gaussian distribution and their computational scheme which will, as they pointed out, only approximately give the desired result when considering levels with  $n > 2$  (Appendix B Stott & Marsh 2018).

### 5.1.3 Comparison with our method

To facilitate a direct comparison with the literature in Fig. 5, we first ignore self-interactions – i.e. take  $f \rightarrow \infty$  or, equivalently, set  $\lambda = 0$  – and compute the posterior probability at fixed values of  $\mu$ .<sup>10</sup> Instead of implementing the ‘Gaussian distance method’, we will use our method, but with an uncorrelated Gaussian instead of the full posterior distributions. This is equivalent to what Stott & Marsh (2018) wanted to compute. For illustrative purposes, for the uncorrelated Gaussian, we choose estimators in  $M$  for both BHs, even though the underlying distribution of the SMBH is better

<sup>9</sup>This is different for QCD axions, due to the connection between  $\mu$  and  $f$  from equation (12). Choosing a Gaussian prior on  $\chi_{\text{top}}$ , we can compute the posterior distribution of  $(f^{-1}, \chi_{\text{top}})$  for M33 X-7, which constrains the relevant portion of parameter space (cf. Fig. 4). In both the equilibrium and bosonova scenarios, we find that  $0.6 \text{ peV} < \mu < 3 \text{ peV}$  for the 95 per cent credible interval, where the lower end of the interval is effectively set by the prior ( $f < M_{\text{p}}$ ).

<sup>10</sup>This technically corresponds to repeating the analysis multiple times with different ‘ $\delta$  priors’.



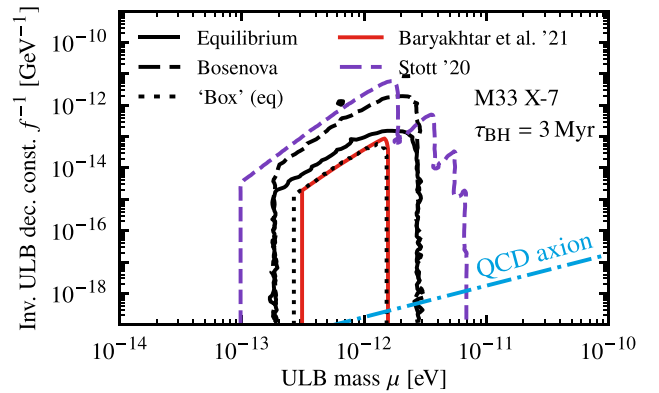
**Figure 5.** Comparison of exclusion methods at fixed mass  $\mu$  for non-interacting ULBs for IRAS 09149-6206 (*left*) and M33 X-7 (*right*). We compare our approach for the full posterior distribution with  $n = 2$  (solid blue lines) and  $n \leq 6$  (dotted blue lines) compared to uncorrelated Gaussians ('Uncorr.'; dash-dotted, purple line) and the 'box method' ('Box'; dashed, red line).

described by a Gaussian in  $\log(M/M_\odot)$ . Finally, note that the 'box method' gives a binary outcome; the probability is either  $P = 0$  or  $P = 1$ .

Consider the left panel of Fig. 5 for IRAS 09149-6206. The 'box method' finds  $P = 1$  for all ULB masses and, indeed, the lines for the other methods and assumptions also remain above the exclusion threshold. However, thanks to our probabilistic interpretation, our method could yield constraints by combining the data from multiple BHs—even when they individually would not constraint ULB parameters. Further note that the deviations between the full and the approximated results mainly come from the approximate Gaussian error propagation, as the data themselves are uncorrelated. Fig. 5 also demonstrates explicitly that the inclusion of more levels can extend the bounds to higher ULB masses. This is because the higher  $n$ , and thus higher  $l$ , allow us to fulfil the SR condition in equation (5) for larger  $\mu$ . For these constraints, we only consider the SR rates  $\Gamma_{nlm}$  of the individual levels with no regards to the BH evolution and time required to fill these levels.

The case of M33 X-7 is shown in the right panel of Fig. 5. Given the clearly visible correlations in the  $(M, a_*)$  distribution in Fig. 2, it is somewhat surprising that neglecting them appears to be an acceptable approximation. This might only be a somewhat lucky accident for M33 X-7, but at least in this case the approximation seems to be inconsequential. The inclusion of levels with  $n \leq 6$  again extends the ULB mass range of the constraints.

Fig. 6 presents a direct comparison with ULB constraints from M33 X-7 in the  $(\mu, f^{-1})$  plane. Stott (2020) used the bosenova prescription for self-interactions, a Gaussian approximation for the BH errors, and included levels up to  $n = 6$  (dashed, purple line). The result broadly agrees with our implementation (dashed, black line) of the bosenova prescription in terms of  $f^{-1}$ , while we observe that the inclusion of higher SR levels extends the limits from Stott (2020) to higher values of  $\mu$ . Our results disagree with Stott (2020) at the low mass end by a factor of  $\mathcal{O}(2)$  on  $\mu$ . This is due to Stott (2020) (and the related works Stott & Marsh 2018; Mehta et al. 2021) neglecting the large  $\ln(N_\Delta)$  factor required to extract significant spin, and the effect of the companion. Baryakhtar et al. (2021) used the 'box' statistical method, and the equilibrium model for self-interactions (solid, red line). As mentioned before, our result (solid, black line) gives rise to a slightly wider exclusion region on  $\mu$ , due to the overly conservative nature of the box method. Note that our implementation



**Figure 6.** Comparison of exclusion contours for  $(\mu, f^{-1})$  with different methods and assumptions for M33 X-7. We compare our results from Fig. 4 to two recent works using the 'box' and 'Gaussian distance' methods. The QCD axion model line is shown as a dash-dotted, blue line.

of the box method (dotted, black line) closely reproduces the results by Baryakhtar et al. (2021) for both  $\mu$  and  $f^{-1}$ .

Advantages of our approach include that we do not assume Gaussianity, which is not always justified – in particular for the treatment of upper limits on  $a_*$ . It is important to stress again that, for SMBHs, Gaussianity only appears to be justified for  $\ln(M/M_\odot)$  rather than  $M$ . While past work also ignored potential correlations between  $M$  and  $a_*$ , we do not find this to be a major issue in our examples. Our approach also turns out to be computationally feasible with Monte Carlo integration, alleviating concerns of Stott & Marsh (2018).

As a final remark, note that, in some sense, the 'box method' is conservative as it accounts for all possible correlations and at least some non-Gaussian features of the  $(M, a_*)$  distribution. This naturally leads to weaker exclusion limits, i.e. smaller exclusion regions for  $\mu$ , as can be seen in Fig. 5. However, the  $(M, a_*)$  distribution contains additional information, and discarding it seems unnecessary. Moreover, the ' $2\sigma$  interval' relies on a (possibly ill-defined) extrapolation from the ' $1\sigma$  error' of  $\hat{M}$  and  $\hat{a}_*$ .<sup>11</sup> Due to

<sup>11</sup>Note that Arvanitaki et al. (2015) directly use the 95 per cent intervals for  $a_*$  instead of an extrapolation.

this, it is not straightforward to apply the ‘box method’ to multiple BHs or to compare to limits obtained at a different exclusion level.

## 5.2 Limitations and possible extensions

As explored before, more powerful limits on ULB masses can be derived using higher SR levels. While we can compute the free SR rates for  $n > 2$ , the most reliable limits on the self-coupling are currently only possible for  $n = 2$ . One important extension of existing work is to compute the equilibrium rates when higher levels are relevant. Furthermore, one may include more interactions, such as stimulated decay in the case of the axion-photon coupling  $g_\gamma \propto f^{-1}$ .

From the data modelling side, our method can also easily include additional parameters to propagate their uncertainties into the ULB exclusion regions. One such parameter is the BH time-scale  $\tau_{\text{BH}}$ , as discussed in Section 3.1, which could be included with a physically-informed prior distribution or constrained by related data sets. We also did not consider possible interactions between the boson cloud and the SMBH environment, which could affect the cloud growth. For instance, Du et al. (2022, see also Arvanitaki et al. 2015) computed the boson cloud’s perturbative stability in a thin accretion disc or in the presence of a stellar halo and found that their stability estimator worsens as  $\alpha^q$  with  $q \gtrsim 5$ .

Additional nuisance parameters can also be included when inferring the mass and spin of a given BH to better reflect their uncertainties. In Section 3, we provided a summary of various aspects of the BH analysis to highlight the underlying assumptions and possible sources of uncertainties. In the case of X-ray reflection modelling, several parameters, e.g. the inclination and iron abundance of the accretion disc, are well-known to be degenerate with the spin in state-of-the-art reflection models. In addition, in the case of continuum-fitting models used in XRBs, spin inference is affected by intrinsic astrophysical uncertainties, such as the inclination angle, distance, or the BH’s mass. More information on these parameters could be provided by involving more constraints, such as data from current and upcoming X-ray polarimetry missions including the currently operating *IXPE* observatory.

Moreover, in addition to including more parameters, we could also employ a hierarchical Bayesian analysis, considering both individual BH data as well as BH population parameters and relations. This has already been explored in the context of BBH mergers (see Section 3.4), which we did not consider in this work. The past work assumed hyperpriors for the initial masses and spins of the merging BH, which we could also employ for SMBHs and XRBs. This would allow us to compute the final BH mass and spin from the full BH evolution including SR, which can be compared to the BH data (potentially using methods from simulation-based inference; see e.g. Cranmer, Brehmer & Louppe 2020). Finally, note that various nuisance parameters, such as the virial factor  $f$  (see e.g. Onken et al. 2004, fig. 2) or a possible  $M$ – $L$  relationship between the BH mass and luminosity (see e.g. Ding et al. 2020, figs 6–7) might be exploited to refine the  $(M, a_*)$  inference for BHs. While the complexity of our analysis would increase when employing hierarchical model, these are natural extensions of a Bayesian framework, which—apart from refining the accuracy of its inference—could help to improve convergence and gain new insights into BH physics.

Compared to both alternative methods, our Bayesian statistical approach more easily generalises to a joint analysis of multiple BHs and the inclusion of other constraints: the Bayesian logic is naturally extendable while allowing us to include any correlations or additional uncertainties through nuisance parameters where necessary. This is useful when considering BHSR constraints in global fits of axion

models, such as the analysis by Hoof et al. (2019) within the GAMBIT global-fitting framework. Apart from previously used BH data (see e.g. Stott 2020), such a global fit could also consider the well-studied XRB GRS 1915 + 105, whose spin has been estimated from both the X-ray reflection and continuum-fitting methods (McClintock et al. 2006; Shreeram & Ingram 2020). Another promising prospect is the inclusion of future data on tidal disruption events near SMBHs from the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), which depend on the BH spin and can constrain ULBs in the range of  $10^{-20} \text{ eV} < \mu < 10^{-18} \text{ eV}$  (Du et al. 2022).

## 6 SUMMARY AND CONCLUSIONS

We have presented a universal approach for obtaining constraints on ULBs from BH data. We exemplified our methodology with two well-studied BHs: M33 X-7, a stellar-mass BH, and IRAS 09149-6206, an SMBH. For each of these, we used posterior samples for the BH’s mass and spin  $(M, a_*)$ , and computed the posterior probability on the ULB model parameters  $(\mu, f^{-1})$ . In particular, these are the first ULB constraints derived from IRAS 09149-6206.

Our method introduces a rigorous statistical framework, extending previous approaches and preserving their advantages while, at the same time, giving more accurate or powerful constraints. In particular, our method appears to be the only proposed methodology to self-consistently derive ULB constraints from SR in multiple BHs. We only rely on samples from the  $(M, a_*)$  posterior distribution alone and need not reproduce the entire BH data analysis tool chain.

Our approach also allowed us to clarify the statistical nature and reasons for disagreement between previous works (see Section 5.1). In particular, we improve on previous methods by capturing correlations and non-Gaussianities of the  $(M, a_*)$  distribution, providing a more complete treatment of *statistical* uncertainties from the BH analyses. Still, we emphasise again that BH inference is affected by potentially large *systematic* uncertainties related to the assumptions and limitations of accretion disc models, BH evolution, bosonova versus equilibrium scenarios, or computation of the SR rates.

Let us summarize the most important results from our work and the comparison to the literature. In this initial study we observe that

- (i) the proposed Bayesian approach automatically includes all information from the data while being easily realisable,
- (ii) neglecting correlations between  $M$  and  $a_*$  may be acceptable,
- (iii) the Gaussian approximation holds for  $\ln(M/M_\odot)$  instead of  $M$  for SMBHs,
- (iv) the Gaussian approximation is poor when  $|a_*| \sim 1$ ,
- (v) the ‘box method’ leads to overly conservative limits and does not leverage the full information contained in the BH data,
- (vi) inclusion of higher levels up to  $n > 2$  can potentially constrain higher masses by up to an order of magnitude,
- (vii) compared to the bosonova, the equilibrium method for self-interactions reduces the maximum value of  $f^{-1}$  probed by any given BH by approximately one order of magnitude.

From the theoretical side, we are still limited by the availability of multi-level equilibrium BH evolution predictions, which would allow us to extend constraints for larger  $\mu$  values. From the data side, a wider applicability of our method is limited by the availability of a (sampled) distribution of  $(M, a_*)$ , and we thus encourage all groups deriving these constraints to make their likelihoods or posteriors publicly available. Each new data set will at worst linearly increase the computational time required for the sampling and more likely be sublinear due to faster convergence and due to abandoning parameter

points that fall below some likelihood threshold after considering part of the data. Overall, the computational effort should be manageable.

Our method has improved on the statistical interpretation of BHSR constraints compared to previous works (Stott & Marsh 2018; Baryakhtar et al. 2021), and we have identified causes of disagreement in the literature due to rates models and statistics. Our conclusions have consequences for BHSR constraints on the string theory landscape. Mehta et al. (2021) computed constraints on the landscape using a version of the implementation of rates and statistics by Stott (2020). We have verified the statistical approximation of Stott & Marsh (2018), Stott (2020) as self-consistent and valid. An error in the low mass end rates of Stott & Marsh (2018), Stott (2020) is relatively small and should not affect conclusions about constraints on the landscape, due to the logarithmic nature of the  $\mu$  distribution. Likewise, using the equilibrium rather than bosenova model for self-interactions will lead to only a slight shift in the maximum topological numbers probed by BHSR, due to the relatively small shift in  $f^{-1}$  and slow trends in landscape models on the variation of this parameter. On the other hand, inclusion of higher levels in Mehta et al. (2021) significantly widens the  $\mu$  range that BHSR can probe. The effect this will have on conclusions relating to the landscape is unclear. Mehta et al. (2021) used a large number of overlapping BHs in the stellar regime, so the effect from any single BH will be small, and the limits should only be affected by the inclusion of higher levels on the lowest mass BH in the sample. Nonetheless, this underscores the importance of computing full BH Regge trajectories in future improvements of BHSR constraints, in order to leverage the constraining power of higher levels accurately, and the importance of extending our methods to large numbers of BHs if the posterior samples can be made available.

Possible extensions of our Bayesian method are discussed in detail in Section 5.2, including additional parameters related to the BH time-scale, Bayesian hierarchical modelling of BH populations, or a companion in binary systems such as M33 X-7. Assuming hyperprior distributions for the initial BH mass and spin would further allow us to compute the final BH mass and spin from the full SR evolution. This treatment is more refined than simply comparing the associated SR and BH time-scales. Similarly, this approach can also incorporate the constraints from binary BH mergers, following Ng et al. (2021a), Ng et al. (2021b), Cheng et al. (2023). We can also include the physics of direct gravitational wave emission from the boson cloud in our model, which has been considered by Arvanitaki et al. (2017), Brito et al. (2017), Baumann et al. (2019a), Palomba et al. (2019), Abbott et al. (2022a), Omiya et al. (2024).

Moreover, constraints on  $\mu$  and  $f$  from BHSR are highly complementary to experimental and astrophysical probes of ULBs, with particular application to the QCD axion, fuzzy DM, and string theory. Our systematic approach will allow this complementarity to be leveraged in global statistical analyses.

To further facilitate these, and possibly other, future extensions, we made our software code available on Github (Hoof 2024).

## ACKNOWLEDGEMENTS

We thank Masha Baryakhtar, Marios Galanis, Olivier Simon, and Sam Witte for helpful, substantial input on our draft and discussions related to their ongoing and past works; Lijun Gou and Xueshan Zhao for ( $M, a_*$ ) samples from Chen et al. (2016); Jinyi Shangguan for  $a_*$  samples from the GRAVITY Collaboration (2017, 2020); Dominic J. Walton for  $a_*$  samples from Walton et al. (2020); and Jakob van den Eijnden, Rob Fender, and Matt Jarvis for helpful discussions on BH masses.

SH has received funding from the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101065579. DJEM is supported by an Ernest Rutherford Fellowship from the STFC, Grant No. ST/T004037/1 and by a Leverhulme Trust Research Project (RPG-2022-145). JSR acknowledges the support from the Science and Technology Facilities Council (STFC) under grant ST/V50659X/1 (project reference 2442592) and the Institute of Astronomy, as well as from a NASA Astrophysics Data Analysis Program (grant 80NSSC24K0617). JHM acknowledges funding from a Royal Society University Research Fellowship. We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support.

We made use of the BIBCOM (Hoof 2022) and WEBPLOTDIGITIZER (Rohatgi 2022) tools, as well as the PYTHON packages NUMBA, NUMPY (Harris et al. 2020), and SCIPY (Virtanen et al. 2020).

## DATA AVAILABILITY

The Python software code and data underlying this article are available on Github at <https://github.com/sebhoof/bhsr>.

## REFERENCES

- Abbott L. F., Sikivie P., 1983, *Phys. Lett. B*, 120, 133  
 Abbott R. et al., 2022a, *Phys. Rev. D*, 105, 102001  
 Abbott R. et al., 2022b, *Phys. Rev. D*, 106, 042003  
 Acharya B. S., Bobkov K., Kumar P., 2010, *JHEP*, 11, 105  
 Adams C. B. et al., 2022, preprint (arXiv:2203.14923)  
 Alves Batista R. et al., 2021, preprint (arXiv:2110.10074)  
 Arias P., Cadamuro D., Goodsell M., Jaeckel J., Redondo J., Ringwald A., 2012, *J. Cosmology Astropart. Phys.*, 2012, 013  
 Arkanī-Hamed N., Motl L., Nicolis A., Vafa C., 2007, *J. High Energy Phys.*, 2007, 060  
 Arnaud K. A., 1996, in Jacoby G. H., Barnes J., eds, *Astronomical Society of the Pacific Conference Series Vol. 101, Astronomical Data Analysis Software and Systems V*. p. 17  
 Arvanitaki A., Dubovsky S., 2011, *Phys. Rev. D*, 83, 044026  
 Arvanitaki A., Dimopoulos S., Dubovsky S., Kaloper N., March-Russell J., 2010, *Phys. Rev. D*, 81, 123530  
 Arvanitaki A., Baryakhtar M., Huang X., 2015, *Phys. Rev. D*, 91, 084011  
 Arvanitaki A., Baryakhtar M., Dimopoulos S., Dubovsky S., Lasenby R., 2017, *Phys. Rev. D*, 95, 043001  
 Aurrekoetxea J. C., Marsden J., Clough K., Ferreira P. G., 2024, *Phys. Rev. D*, 110, 083011  
 Bachetti M. et al., 2014, *Nature*, 514, 202  
 Bahramian A., Degenaar N., 2023, in *Handbook of X-ray and Gamma-ray Astrophysics*. Edited by Cosimo Bambi and Andrea Santangelo. Springer, Singapore, p. 120  
 Bambi C. et al., 2021, *Space Sci. Rev.*, 217, 65  
 Bao S.-S., Xu Q.-X., Zhang H., 2022, *Phys. Rev. D*, 106, 064016  
 Baryakhtar M., Galanis M., Lasenby R., Simon O., 2021, *Phys. Rev. D*, 103, 095019  
 Bauer J. B., Marsh D. J. E., Hložek R., Padmanabhan H., Laguë A., 2021, *MNRAS*, 500, 3162  
 Baumann D., Chia H. S., Porto R. A., 2019a, *Phys. Rev. D*, 99, 044001  
 Baumann D., Chia H. S., Stout J., ter Haar L., 2019b, *J. Cosmology Astropart. Phys.*, 2019, 006  
 Baumann D., Bertone G., Stout J., Tomaselli G. M., 2022, *Phys. Rev. Lett.*, 128, 221102  
 Bennett J. S., Sijacki D., Costa T., Laporte N., Witten C., 2024, *MNRAS*, 527, 1033  
 Bentz M. C. et al., 2013, *ApJ*, 767, 149  
 Boyer R. H., Lindquist R. W., 1967, *J. Math. Phys.*, 8, 265

- Brito R., Ghosh S., Barausse E., Berti E., Cardoso V., Dvorkin I., Klein A., Pani P., 2017, *Phys. Rev. D*, 96, 064050
- Brito R., Cardoso V., Pani P., 2020, Superradiance, 2nd edn. Lecture Notes in Physics Vol. 906. Springer, Cham preprint (arXiv:1501.06570)
- Bustamante S., Springel V., 2019, *MNRAS*, 490, 4133
- Cardoso V., Dias Ó. J. C., Hartnett G. S., Middleton M., Pani P., Santos J. E., 2018, *J. Cosmology Astropart. Phys.*, 2018, 043
- Casares J., Jonker P. G., 2014, *Space Sci. Rev.*, 183, 223
- Chadha-Day F., Ellis J., Marsh D. J. E., 2022, *Sci. Adv.*, 8, eabj3618
- Chaty S., 2022, Accreting Binaries; Nature, formation, and evolution. IOP Publishing, Bristol, England
- Chen Z., Gou L., McClintock J. E., Steiner J. F., Wu J., Xu W., Orosz J. A., Xiang Y., 2016, *ApJ*, 825, 45
- Cheng L.-d., Zhang H., Bao S.-s., 2023, *Phys. Rev. D*, 107, 063021
- Cicoli M., Goodsell M., Ringwald A., 2012, *JHEP*, 10, 146
- Conlon J. P., 2006, *JHEP*, 05, 078
- Corral-Santana J. M., Casares J., Muñoz-Darias T., Bauer F. E., Martínez-Pais I. G., Russell D. M., 2016, *A&A*, 587, A61
- Cranmer K., Brehmer J., Louppe G., 2020, *Proc. Natl. Acad. Sci.*, 117, 30055
- Curtis-Lake E. et al., 2023, *Nature Astronomy*, 7, 622
- Daly R. A., 2022, *MNRAS*, 517, 5144
- Dauser T., Garcia J., Parker M. L., Fabian A. C., Wilms J., 2014, *MNRAS*, 444, L100
- Dauser T., Garcia J., Walton D. J., Eikmann W., Kallman T., McClintock J., Wilms J., 2016, *A&A*, 590, A76
- Davis S. W., Blaes O. M., Hubeny I., Turner N. J., 2005, *ApJ*, 621, 372
- Davoudiasl H., Denton P. B., 2019, *Phys. Rev. Lett.*, 123, 021102
- Detweiler S., 1980, *Phys. Rev. D*, 22, 2323
- di Cortona G. G., Hardy E., Vega J. P., Villadoro G., 2016, *J. High Energy Phys.*, 2016, 34
- Dine M., Fischler W., 1983, *Phys. Lett. B*, 120, 137
- Ding X. et al., 2020, *ApJ*, 888, 37
- Dolan S. R., 2007, *Phys. Rev. D*, 76, 084001
- Done C., Davis S. W., Jin C., Blaes O., Ward M., 2012, *MNRAS*, 420, 1848
- Du P., Egaña-Ugrinovic D., Essig R., Fragione G., Perna R., 2022, *Nat. Commun.*, 13, 4626
- Dvorkin C. et al., 2022, preprint (arXiv:2203.07064)
- East W. E., Pretorius F., 2017, *Phys. Rev. Lett.*, 119, 041101
- Fabian A. C., Rees M. J., Stella L., White N. E., 1989, *MNRAS*, 238, 729
- Farrah D. et al., 2022, *MNRAS*, 513, 4770
- Farren G. S., Grin D., Jaffe A. H., Hložek R., Marsh D. J. E., 2022, *Phys. Rev. D*, 105, 063513
- Fernandez N., Ghalsasi A., Profumo S., 2019, preprint (arXiv:1911.07862)
- Ferrarese L., Merritt D., 2000, *ApJ*, 539, L9
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Foschi A. et al., 2023, *MNRAS*, 524, 1075
- GRAVITY Collaboration, 2017, *A&A*, 602, A94
- GRAVITY Collaboration, 2020, *A&A*, 643, A154
- GRAVITY Collaboration, 2021, *A&A*, 648, A117
- García J. et al., 2014, *ApJ*, 782, 76
- Gebhardt K. et al., 2000, *ApJ*, 539, L13
- Genzel R., Eisenhauer F., Gillessen S., 2010, *Rev. Modern Phys.*, 82, 3121
- George I. M., Fabian A. C., 1991, *MNRAS*, 249, 352
- Ghez A. M. et al., 2008, *ApJ*, 689, 1044
- Gorghetto M., Villadoro G., 2019, *J. High Energy Phys.*, 2019, 33
- Gou L. et al., 2011, *ApJ*, 742, 85
- Gravity Collaboration, 2018, *Nature*, 563, 657
- Greene J. E. et al., 2010, *ApJ*, 721, 26
- Hannuksela O. A., Wong K. W. K., Brito R., Berti E., Li T. G. F., 2019, *Nat. Astron.*, 3, 447
- Harris C. R. et al., 2020, *Nature*, 585, 357
- Hod S., 2016, *Phys. Lett. B*, 758, 181
- Hoof S., 2022, *BibCom – A BibTeX bibliography creator*. Accessed 3/10/25. Available on Github at <https://github.com/sebhoof/bibcom>
- Hoof S., 2024, BHSR – Statistical inference for ultralight bosons from black hole superradiance. Accessed 3/10/25. Available on Github at <https://github.com/sebhoof/bhsr>
- Hoof S., Kahlhoefer F., Scott P., Weniger C., White M., 2019, *J. High Energy Phys.*, 2019, 191
- Hu W., Barkana R., Gruzinov A., 2000, *Phys. Rev. Lett.*, 85, 1158
- Hui L., Ostriker J. P., Tremaine S., Witten E., 2017, *Phys. Rev. D*, 95, 043541
- Irastorza I. G., Redondo J., 2018, *Prog. Part. Nucl. Phys.*, 102, 89
- Jiang J. et al., 2019, *MNRAS*, 489, 3436
- Kaspi S., Smith P. S., Netzer H., Maoz D., Jannuzi B. T., Giveon U., 2000, *ApJ*, 533, 631
- Kaspi S., Maoz D., Netzer H., Peterson B. M., Vestergaard M., Jannuzi B. T., 2005, *ApJ*, 629, 61
- Kerr R. P., 1963, *Phys. Rev. Lett.*, 11, 237
- Khlopov M. I., Malomed B. A., Zeldovich I. B., 1985, *MNRAS*, 215, 575
- King A., Lasota J.-P., Middleton M., 2023, *New A Rev.*, 96, 101672
- Kodama H., Yoshino H., 2012, preprint (arXiv:1108.1365)
- Kulkarni A. K. et al., 2011, *MNRAS*, 414, 1183
- Leaver E. W., 1985, *Proc. Roy. Soc. Lond. Ser. A*, 402, 285
- Liu J., McClintock J. E., Narayan R., Davis S. W., Orosz J. A., 2008, *ApJ*, 679, L37
- Liu H., Luo B., Brandt W. N., Brotherton M. S., Gallagher S. C., Ni Q., Shemmer O., Timlin J. D. III, 2021, *ApJ*, 910, 103
- Magorrian J. et al., 1998, *AJ*, 115, 2285
- Maiolino R. et al., 2024, *Nature*, 627, 59
- Mallick L. et al., 2022, *MNRAS*, 513, 4361
- Marsh D. J. E., 2016, *Phys. Rep.*, 643, 1
- Marsh D. J. E., Hoof S., 2023, in Kimball D. F. J., van Bibber K., eds, The Search for Ultralight Bosonic Dark Matter. Springer, Cham, p. 73
- Marsh D. J. E., Silk J., 2014, *MNRAS*, 437, 2652
- McClintock J. E., Shafee R., Narayan R., Remillard R. A., Davis S. W., Li L.-X., 2006, *ApJ*, 652, 518
- McClintock J. E., Narayan R., Steiner J. F., 2014, *Space Sci. Rev.*, 183, 295
- McLure R. J., Jarvis M. J., 2002, *MNRAS*, 337, 109
- Mehta V. M., Demirtas M., Long C., Marsh D. J. E., McAllister L., Stott M. J., 2021, *J. Cosmology Astropart. Phys.*, 2021, 033
- Ng K. K. Y., Vitale S., Zimmerman A., Chatziioannou K., Gerosa D., Haster C.-J., 2018, *Phys. Rev. D*, 98, 083007
- Ng K. K. Y., Hannuksela O. A., Vitale S., Li T. G. F., 2021a, *Phys. Rev. D*, 103, 063010
- Ng K. K. Y., Vitale S., Hannuksela O. A., Li T. G. F., 2021b, *Phys. Rev. Lett.*, 126, 151102
- Novikov I. D., Thorne K. S., 1973, in Black Holes (Les Astres Occlus), Gordon and Breach, New York, p. 343
- O’Hare C. A. J., 2024, *PoS, COSMICWISPer*, 040
- Omiya H., Takahashi T., Tanaka T., Yoshino H., 2023, *J. Cosmology Astropart. Phys.*, 2023, 016
- Omiya H., Takahashi T., Tanaka T., Yoshino H., 2024, 110, 044002
- Onken C. A., Ferrarese L., Merritt D., Peterson B. M., Pogge R. W., Vestergaard M., Wandel A., 2004, *ApJ*, 615, 645
- Orosz J. A. et al., 2007, *Nature*, 449, 872
- Palomba C. et al., 2019, *Phys. Rev. Lett.*, 123, 171101
- Pani P., Cardoso V., Gualtieri L., Berti E., Ishibashi A., 2012, *Phys. Rev. D*, 86, 104017
- Peccei R. D., Quinn H. R., 1977a, *Phys. Rev. D*, 16, 1791
- Peccei R. D., Quinn H. R., 1977b, *Phys. Rev. Lett.*, 38, 1440
- Pietsch W., Haberl F., Sasaki M., Gaetz T. J., Plucinsky P. P., Ghavamian P., Long K. S., Pannuti T. G., 2006, *ApJ*, 646, 420
- Planck Collaboration VI, 2020, *A&A*, 641, A6
- Preskill J., Wise M. B., Wilczek F., 1983, *Phys. Lett. B*, 120, 127
- Reynolds C. S., 2021, *ARA&A*, 59, 117
- Ricarte A., Tiede P., Emami R., Tamar A., Natarajan P., 2023, *Galaxies*, 11, 6
- Richartz M., Luis Rosa J., Berti E., 2024, preprint (arXiv:2405.01003)
- Rohatgi A., 2022, Webplotdigitizer. Accessed 3/10/25, <https://automeris.io/WebPlotDigitizer>
- Rosa J. G., 2013, *J. High Energy Phys.*, 2013, 14

- Sala L., Valentini M., Biffi V., Dolag K., 2024, *A&A*, 685, A92
- Schive H.-Y., Chiueh T., Broadhurst T., 2014, *Nature Phys.*, 10, 496
- Schulze A., Wisotzki L., 2010, *A&A*, 516, A87
- Semertzidis Y. K., Youn S., 2022, *Sci. Adv.*, 8, eabm9928
- Shafee R., Narayan R., McClintock J. E., 2008a, *ApJ*, 676, 549
- Shafee R., McKinney J. C., Narayan R., Tchekhovskoy A., Gammie C. F., McClintock J. E., 2008b, *ApJ*, 687, L25
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
- Shirakata H., Kawaguchi T., Oogi T., Okamoto T., Nagashima M., 2019, *MNRAS*, 487, 409
- Shreeram S., Ingram A., 2020, *MNRAS*, 492, 405
- Siemonsen N., May T., East W. E., 2023, *Phys. Rev. D*, 107, 104003
- Stein L., 2019, *J. Open Source Softw.*, 4, 1683
- Steiner J. F., McClintock J. E., Remillard R. A., Narayan R., Gou L., 2009, *ApJ*, 701, L83
- Stott M. J., 2020, preprint (arXiv:2009.07206)
- Stott M. J., Marsh D. J. E., 2018, *Phys. Rev. D*, 98, 083006
- Suh H., Hasinger G., Steinhardt C., Silverman J. D., Schramm M., 2015, *ApJ*, 815, 129
- Svrcek P., Witten E., 2006, *JHEP*, 06, 051
- Tamburini F., Thidé B., Della Valle M., 2020, *MNRAS*, 492, L22
- Temple M. J. et al., 2023, *MNRAS*, 523, 646
- Teukolsky S. A., 1972, *Phys. Rev. Lett.*, 29, 1114
- Turner M. S., 1983, *Phys. Rev. D*, 28, 1243
- Turner M. S., 1986, *Phys. Rev. D*, 33, 889
- Únal C., Loeb A., 2020, *MNRAS*, 495, 278
- Valiante R., Agarwal B., Habouzit M., Pezzulli E., 2017, *PASA*, 34, e031
- Valsecchi F., Glebbeek E., Farr W. M., Fragos T., Willems B., Orosz J. A., Liu J., Kalogera V., 2010, in Kalogera V., van der Sluys M., eds, *AIP Conf. Ser. Vol. 1314, International Conference on Binaries: in celebration of Ron Webbink's 65th Birthday*. Am. Inst. Phys., New York, p. 285
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Walton D. J. et al., 2020, *MNRAS*, 499, 1480
- Weinberg S., 1978, *Phys. Rev. Lett.*, 40, 223
- Wen S., Jonker P. G., Stone N. C., Zabludoff A. I., 2021, *ApJ*, 918, 46
- Wilczek F., 1978, *Phys. Rev. Lett.*, 40, 279
- Witte S. J., Mummery A., 2025, *Phys.Rev.D*, 111, 083044
- Xie N., Huang F. P., 2024, *Sci. China Phys. Mech. Astron.*, 67, 210411
- Yoshino H., Kodama H., 2012, *Prog. Theor. Phys.*, 128, 153
- Yoshino H., Kodama H., 2015, *Class. Quant. Grav.*, 32, 214001
- Zdziarski A. A., Ghisellini G., George I. M., Svensson R., Fabian A. C., Done C., 1990, *ApJ*, 363, L1
- Zhang H., Behroozi P., Volonteri M., Silk J., Fan X., Aird J., Yang J., Hopkins P. F., 2024, *MNRAS*, 529, 2777
- Zhu Y., Davis S. W., Narayan R., Kulkarni A. K., Penna R. F., McClintock J. E., 2012, *MNRAS*, 424, 2504
- Zhu S. J., Baryakhtar M., Papa M. A., Tsuna D., Kawanaka N., Eggenstein H.-B., 2020, *Phys. Rev. D*, 102, 063020
- Zimmermann T., Alvey J., Marsh D. J. E., Fairbairn M., Read J. I., 2025, *Phys.Rev.Lett.*, 134, 151001

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.