

A Joint Radio–Infrared–X-ray Bound on Bondi-fed Accretion onto the Candidate Intermediate-Mass Black Hole in Omega Centauri

Tim Swanson

The Omega Centauri Society / Post Oak Labs

tim@postoaklabs.com

August 2026

v2.1 (last revised 2026-09-03) — sixth paper of the set (Paper F); companion to the Macro Transcension Hypothesis (Paper A), the inward-migration review (Paper B), the Omega Centauri campaign (Paper C), the migration economics (Paper D), the engineered-IMBH systems paper (Paper E), the residual X-ray census (Paper G), and the mass-tension paper (Paper H);
prepared for omegacentauri.me

Abstract

Three deep non-detections now constrain any accretion flow onto the candidate intermediate-mass black hole (IMBH) at the center of Omega Centauri: a 170-hour ATCA radio campaign reaching $1.1 \mu\text{Jy}$ rms at 7.25 GHz (Mahida et al., 2026), JWST NIRCам/MIRI imaging showing no accretion-like point source at any proposed center (Chen et al., 2026), and a 291-ks Chandra exposure bounding $L_X(0.5\text{--}7 \text{ keV}) \leq 1.6 \times 10^{30} \text{ erg s}^{-1}$ (Haggard et al., 2013). Chen et al. (2026) compare the infrared and radio limits qualitatively; no formal joint bound with propagated uncertainties has been published. We supply one, as a posterior. Defining the Bondi radiative efficiency $\varepsilon_B \equiv L_{\text{bol}}/(\dot{M}_B c^2)$, we build a forward model from (ε_B, M) and eight shared nuisance parameters (gas density, sound speed, distance, two SED band fractions, fundamental-plane scatter as a latent variable, inflow-suppression index, electron-heating fraction) to the three measured quantities, evaluate Gaussian likelihoods on the measurements themselves, and report the 95 per cent credible upper limit $\varepsilon_{95}(M)$ for three named accretion-flow families. For a radiatively inefficient (RIAF-like) flow, $\varepsilon_{95} = 2.9 \times 10^{-10}$ at the fast-star mass anchor of $8,200 M_\odot$ and 1.1×10^{-11} at $4 \times 10^4 M_\odot$; removing the fundamental-plane radio leg entirely relaxes these to 1.4×10^{-8} and 7.9×10^{-10} . Overplotting the expectation band for a natural outflow-suppressed hot flow at the same nuisance draws converts the curve into a calibrated verdict. Scoring each natural-flow draw against the measured data directly, the data reject 85 per cent of that parameter space at the fast-star anchor, 78 per cent at the pulsar-timing point-mass cap and 99 per cent at $4 \times 10^4 M_\odot$ (28, 20 and 76 per cent using no radio information). A Bondi-fed hot flow at a 47 Tucanae-like density is therefore disfavoured across the contested mass range rather than only at its top, and something must give at every anchor: the mass, the gas density, or the suppression physics. Two of those figures replace weaker ones. Drafts of this paper to v1.5 reported 31 and 90 per cent at the two upper anchors and a natural flow surviving comfortably at the fast-star anchor; that reading rested on an efficiency law linear in accretion rate where the fits it cited are close to square-root, and on an exclusion statistic that compared the natural draws against a limit marginalized over those same draws. Both are corrected here and both move the verdict the same way. The strengthened claim is also the more fragile one, since the corrected efficiency law is extrapolated three to five decades below the rates it was computed for; holding it flat at the edge of its computed range instead drives

every figure above 99 per cent, so the extrapolation is the conservative end of the bracket rather than the optimistic one. The limit's normalization is conditional on the unmeasured central gas density throughout, and we quantify that dependence three ways: prior-widening and prior-shift tests (factor ~ 6 each), a stellar-wind budget that bounds the density from above rather than below once its giant census is recounted, and a forecast showing that a pulsar dispersion-measure determination of the core density, feasible with the current 19-pulsar timing set, would harden the entire curve. The single observation that converts these limits from analogy-conditional to measured is a DM-gradient fit to the 19-pulsar set at the ± 10 per cent level (§8.1). Resolving the radio campaign into the 25 observing blocks its source paper tabulates also converts the duty-cycle loophole from an acknowledgement into a bounded region: a train of $100 \mu\text{Jy}$ flares lasting 2 hr is excluded for recurrence intervals shorter than 31 hr, while a 1 mJy flare lasting half an hour escapes the whole campaign if it recurs less often than every 3.4 days, a duty cycle of 6×10^{-3} . All results derive from a fixed-seed Monte Carlo shipped with the paper.

Keywords: intermediate-mass black holes; Omega Centauri; NGC 5139; Bondi accretion; radiatively inefficient accretion flows; fundamental plane of black hole activity

Contents

1	Introduction	4
2	Observational inputs	4
3	Method	6
3.1	Definitions and conventions	6
3.2	Shared nuisance parameters	6
3.3	Forward model and likelihoods	8
3.4	Accretion-flow families	8
3.5	The natural-flow expectation and its efficiency law	9
3.6	Defining the exclusion fraction	9
4	The gas supply, bounded from above	10
5	Results	11
5.1	The exclusion fraction as a continuous function of mass	13
5.2	Which band binds which family	14
5.3	Sensitivity budget	15
5.3.1	Which nuisance carries the spread	17
5.4	Where this sits against the survey population	18
6	The duty-cycle loophole	19
6.1	What the block list actually excludes	19
7	Forecasts	20
8	Scope of the adjudication	22
9	Conclusion	23
A	Conventions, unit conversions, and cross-paper reconciliation	24
A.1	Shared fiducial parameters across the paper set	26
B	Threshold-inversion cross-check	26

1 Introduction

Omega Centauri hosts the nearest strong IMBH candidate: fast-moving stars inside the central arcsecond require an enclosed dark mass $\gtrsim 8,200 M_\odot$ (Häberle et al., 2024), while pulsar-timing and kinematic modeling favor an extended dark remnant component of $\sim 2\text{--}3 \times 10^5 M_\odot$ and cap any point mass at $\lesssim 6,000 M_\odot$ at 3σ (Bañares-Hernández et al., 2025), a tension sharpened rather than resolved by the current 19-pulsar timing set (Colomo & Bernadich et al., 2026). The center is electromagnetically silent at the depths quantified below. The radio silence has a two-decade history: upper limits of $\sim 100 \mu\text{Jy}$ at GHz frequencies (Maccarone et al., 2005), a 2.5σ near-center excursion flagged as worth deeper follow-up (Lu and Kong, 2011), JVLA limits of $1.5\text{--}2.1 \mu\text{Jy beam}^{-1}$ in three other clusters that set the methodological template for radio IMBH searches (Strader et al., 2012), the 50-cluster MAVERIC survey null (Tremou et al., 2018), and now the deepest radio image of any globular cluster, 170 hours with ATCA reaching $1.1 \mu\text{Jy rms}$ at 7.25 GHz with no source at any proposed center and the nearest 5σ source 13 arcseconds away (Mahida et al., 2026). The Lu and Kong (2011) excursion does not recur at a factor ~ 6 deeper in rms; the 2.5σ peak they flagged would have registered at $\gtrsim 14\sigma$ in the new image, which retires that thread. In the infrared, JWST NIRCам and MIRI imaging of the core shows no point source with an accretion-like spectral energy distribution (Chen et al., 2026); in X-rays, the combined 291-ks Chandra exposure gives $L_X \leq 1.6 \times 10^{30} \text{ erg s}^{-1}$, which Haggard et al. (2013) translated to accretion efficiencies spanning $10^{-5}\text{--}10^{-9}$ across the $3 \times 10^3\text{--}3 \times 10^4 M_\odot$ mass window they consider. Their abstract also rounds that span to $10^{-6}\text{--}10^{-8}$; drafts of this paper to v1.5 quoted the rounded form as though it were the source’s range, and the wider figure is the one their text states.

Chen et al. (2026) place their infrared limits alongside the radio limits and note that the radio is the more restrictive at the high-mass end (their Figure 7). That comparison is qualitative: each band’s limit is quoted under its own model assumptions, at its own fiducial mass, with no shared error budget. The three limits constrain three different quantities, a radio flux routed through a jet or fundamental-plane relation, a band-limited infrared luminosity, and a band-limited X-ray luminosity, and all three depend on the same unmeasured gas density, the same contested black-hole mass, and the same distance. Multiplying or stacking them as if independent would overstate the joint constraint; quoting any of them at a single mass begs the question the mass tension leaves open.

This paper supplies the missing joint analysis, and in doing so surfaces a result the single-band treatments could not state: the combined data are now deep enough to exclude most of the parameter space of a natural Bondi-fed hot flow across the contested mass range. Section 2 fixes the observational inputs; Section 3 the forward model, the likelihoods, and the three accretion-flow families; Section 4 the gas-supply physics that bounds the dominant nuisance from above; Section 5 the posterior curves, the natural-flow comparison, and the prior-sensitivity budget; Section 6 the duty-cycle loophole; Section 7 forecasts; and Section 8 what the bound adjudicates. Appendix A collects every unit conversion and convention; Appendix B retains the simpler threshold-inversion construction of an earlier draft as a transparent cross-check. The Monte Carlo script, its fixed seed, and its full output accompany the paper.

2 Observational inputs

Radio. Mahida et al. (2026): ATCA project CX556 plus archival data, ~ 170 h at 5.5 and 9.0 GHz combining to an effective 7.25 GHz image with rms $\sigma_S = 1.1 \mu\text{Jy beam}^{-1}$; no detection at any of their three tested centers, A, B, and C, coincident with though not labelled as the van der Marel and Anderson (2010) and Noyola et al. (2008, 2010) centers. We treat the central flux density as measured to be $0 \pm \sigma_S$.

Infrared. Chen et al. (2026): JWST NIRCам and MIRI imaging of the core with no accretion-like

Table 1: JWST per-filter limits adopted for the infrared leg, from [Chen et al. \(2026\)](#), their Table 1, at the 95 per cent completeness row, at their adopted distance of 5.49 kpc. L_{lim} is band-integrated, as that table’s own definition $L = 4\pi d^2 F_\lambda \Delta\lambda$ states. $\Delta\lambda$ is not printed in the source and is recovered from the ratio L/F_λ across that filter’s completeness rows, which agree to 1–6 per cent. The likelihood predicts a band-peak νL_ν , so the conversion between the two enters each filter’s term: $\sigma_j = (\lambda_c/\Delta\lambda) L_{\text{lim},j}/1.645$ (§3.3). Drafts to v1.5 omitted that factor and compared a band-peak prediction against a band-integrated limit.

Filter	Vega mag limit	L_{lim} (erg s ^{−1})	$\Delta\lambda$ (Å)	$\lambda_c/\Delta\lambda$	σ_j (erg s ^{−1})
F200W	> 20.8	4.8×10^{30}	4,720	4.24	1.24×10^{31}
F444W	> 17.8	9.2×10^{30}	11,090	4.00	2.24×10^{31}
F770W	> 17.5	1.9×10^{30}	14,985	5.14	5.94×10^{30}
F1500W	> 14.5	4.9×10^{30}	26,850	5.59	1.66×10^{31}

point source. Their Table 1 states a limiting luminosity per filter at each of three completeness levels, and we use those numbers directly rather than a single band-summary constant (Table 1). Each filter enters the likelihood as its own Gaussian term on the predicted band-peak νL_ν , at the 95 per cent completeness row, the row that matches the X-ray leg’s stated confidence. Their tabulated limits are band-integrated, as their own definition $L = 4\pi d^2 F_\lambda \Delta\lambda$ states, so each σ_j carries the conversion between the two quantities as well as the quantile: $\sigma_j = (\lambda_c/\Delta\lambda) L_{\text{lim},j}/1.645$, a factor of 4.0 to 5.6 per filter (Table 1). Drafts to v1.5 compared the band-peak prediction against the band-integrated limit directly, which made the infrared terms tighter than the data support by that factor; correcting it loosens the RIAF anchors by 4 to 7 per cent, the jet anchors by 32 to 37 per cent and the radio-free anchors by 60 per cent. Completeness is a detection-recovery fraction rather than a confidence level, so reading it through the normal quantile is a convention; the 99.7 and 68 per cent rows are run alongside and are reported in Appendix A. F770W dominates the combination. Treating four filters as independent terms is optimistic, since they observe one source: the conservative reading, which keeps only the tightest single filter, sits within 2 per cent of the combination on the jet family and 6 per cent on the RIAF family, so nothing here rests on the independence assumption. A second correction to the same table is disclosed rather than made silently. The 99.7 per cent completeness row for F200W is transcribed in the provenance file with a luminosity entry equal to its flux entry, which breaks the ratio L/F_λ that the filter’s other three rows share to 1 per cent; the value that ratio implies, 7.5×10^{31} erg s^{−1}, is used in the 99.7 per cent ladder row of Appendix A. That row is not the primary one and no headline number depends on it. Limits are referenced at [Chen et al.](#)’s adopted 5.49 kpc and rescaled to the drawn distance. [Chen et al. \(2026\)](#) state their accretion constraint for $M \lesssim 10^4 M_\odot$, so the $4 \times 10^4 M_\odot$ anchor extrapolates their stated range by a factor of four; the SED-family treatment of Section 3.4 is what carries that anchor, and the radio-free curve there should be read with this caveat attached.

X-ray. [Haggard et al. \(2013\)](#): 290.9 ks of Chandra ACIS data, $f_X(0.5\text{--}7\text{ keV}) \leq 5.0 \times 10^{-16}$ erg cm^{−2} s^{−1} absorption-corrected. [Haggard et al. \(2013\)](#) quote this at 95 per cent confidence and never write a σ ; their aprates bound is on a non-negative count rate, so we read it as one-sided and take the measured central flux as $0 \pm f_{X,\text{lim}}/1.645$. An earlier draft of this paper labelled the limit 3σ and used $f_{X,\text{lim}}/3$, which is not the source’s convention. The two-sided 95 per cent and the $f_{X,\text{lim}}/3$ readings are carried as a ladder in Appendix A; they span 5 per cent on the RIAF anchors, so the choice inside the ladder is not load-bearing.

Positional coverage. The proposed centers ([van der Marel and Anderson, 2010](#); [Noyola et al., 2008, 2010](#)) differ by over an arcsecond, and the hole wanders about the potential minimum with r.m.s. amplitude $\lesssim 10^{-2}$ pc, about 0.4'' at 5.43 kpc, with a leptokurtic distribution whose excursions reach 1–2'' more often than a Gaussian would allow ([Di Cintio et al., 2023](#); [Chatterjee et al., 2002](#)). None of this strains the

limits; the reason is independent of beaming: each dataset images or covers the entire core region at its quoted depth, and in the radio the nearest 5σ source of any kind lies $13''$ from Center B, in [Mahida et al.](#)’s own phrasing ([Mahida et al., 2026](#)). Every position in the center list, and the full wander tail, inherits the per-band limits; no positional correction enters the error budget. The distance used here, 5.43 kpc, is this paper’s own likelihood input (Table 2); the companion papers in the set adopt 5.49 kpc as the set fiducial, and the resulting offset, 1.2σ of this paper’s distance prior, is disclosed in §3.2 rather than harmonized away.

3 Method

3.1 Definitions and conventions

Let $\dot{M}_B$ be the Bondi capture rate from gas at rest in the cluster core,

$$\dot{M}_B = 4\pi\lambda \frac{(GM)^2 \rho_{\text{gas}}}{c_s^3}, \quad \lambda(\gamma=5/3) = 0.25 (\text{Bondi, 1952}), \quad (1)$$

with $\rho_{\text{gas}} = \mu_e m_p n_e$ and $\mu_e = 1.17$ for an ionized plasma of solar-like composition. Gas in hydrostatic equilibrium shares the cluster’s rest frame, and the hole’s Brownian velocity, $\sigma(m_*/M)^{1/2} \sim 0.1 \text{ km s}^{-1}$, is negligible against c_s , so no bulk-velocity term enters the denominator; this is the convention of the comparison literature ([Strader et al., 2012](#); [Tremou et al., 2018](#); [Haggard et al., 2013](#)). A variant that treats the stellar velocity dispersion as a turbulent-velocity proxy in the denominator, the convention of Paper E’s fuel budget, is evaluated in Appendix A; it lowers $\dot{M}_B$ by a factor $\simeq 1.9$ net of the eigenvalue and weakens every limit by the same factor.

We define the *Bondi radiative efficiency*

$$\varepsilon_B \equiv \frac{L_{\text{bol}}}{\dot{M}_B c^2}, \quad (2)$$

the bolometric radiative output per unit of captured rest-mass energy, the quantity the data constrain directly. It factors as $\varepsilon_B = \eta f_B$, with f_B the fraction of captured gas reaching the horizon and η the radiative efficiency of the arriving flow. In natural hot flows the inflow declines inward, $\dot{M}(R) \propto R^s$, with $0 \leq s \leq 1$ admissible on general grounds ([Blandford and Begelman, 1999](#)) and modeled values spanning $s \simeq 0.3\text{--}0.8$ ([Yuan and Narayan, 2014](#)); we adopt $s \in [0.3, 0.5]$ as this paper’s working sub-range within that modeled span. With $r_B = GM/c_s^2$ and the gravitational radius $r_g = GM/c^2$, so that $r_B/r_g = (c/c_s)^2$ and the bracket depends on c_s and s alone,

$$f_B = \left(\frac{r_B}{r_g}\right)^{-s} \in [10^{-4.4}, 10^{-2.6}] \quad \text{for } s \in [0.3, 0.5], \quad (3)$$

consistent with Paper E’s fiducial $s \simeq 0.3$ suppression of ~ 500 at $2 \times 10^4 M_\odot$. Published per-band efficiencies use other conventions: [Mahida et al. \(2026\)](#) bound a model-specific accretion fraction ($\lesssim 4 \times 10^{-3}$), [Haggard et al. \(2013\)](#) an efficiency under fixed gas density and mass ($10^{-5}\text{--}10^{-9}$ over their stated mass window). Appendix A prints the conversion table; none of those numbers is ε_B , and no direct numerical comparison should be made without it.

3.2 Shared nuisance parameters

Table 2 lists the priors. One draw of the full vector serves all three bands in each realization: the gas density that scales the radio prediction is the same density that scales the infrared and X-ray ones, the SED

Table 2: Monte Carlo priors (4×10^4 draws, fixed seed; script `figs/ff_posterior_v4.py`). SED band fractions are family-specific (Table 3).

Parameter	Prior	Basis
n_e	log-normal, median 0.23 cm^{-3} , 0.5 dex	47 Tuc analogy (Freire et al., 2001; Abate et al., 2018); §4
c_s	uniform, 11.7–16.6 km s^{-1}	10^4 K photoionized gas, isothermal ($\mu=0.6$) to adiabatic (pure H)
d	normal, $5.43 \pm 0.05 \text{ kpc}$	cluster distance; rescales all bands
f_X, f_{IR}	family-specific (Table 3)	hot-flow SED literature
FP latent offset	normal, $0 \pm 0.88 \text{ dex}$	fundamental-plane intrinsic scatter (Merloni et al., 2003)
s	uniform, 0.3–0.5	inflow-suppression index (Yuan and Narayan, 2014); Paper E fiducial 0.3; enters ε_{nat} only, not the likelihood (§5.3.1)
δ	uniform over $\{10^{-3}, 10^{-2}, 0.1, 0.5\}$	electron-heating fraction; the four values Xie and Yuan (2012) tabulate, drawn discretely because the efficiency fit exists only at those values; enters ε_{nat} only (§3.5)
σ	fixed, 21 km s^{-1}	paper-set fiducial (enters r_{infl} only)

fraction f_X that converts the X-ray measurement also drives the fundamental-plane prediction (the two legs constrain the same source’s SED, so drawing f_X once induces the correlation reality does), and the distance rescales all three measured quantities together.

The gas density is the weakest input and it multiplies everything, so its treatment is spread across three places: the prior here (a 47 Tucanae analogy, since $\omega \text{ Cen}$ has no direct ionized-gas measurement; the one $\omega \text{ Cen}$ -specific figure, $n_e = 1.94 \text{ cm}^{-3}$ from sodium absorption, is flagged by its authors as a foreground-dominated upper limit and is not used; Wang et al. 2025), the wind budget of Section 4, which bounds the local supply from above, and the sensitivity budget of Section 5.3. Mahida et al. (2026) adopt $n_e = 0.2 \pm 0.1 \text{ cm}^{-3}$ for the same cluster on the same kind of reasoning, from pulsar-derived electron densities in other globular clusters, which is an independent arrival at the median used here and an indication of how narrow the defensible range is once the analogy is accepted at all. Every predicted flux in the forward model scales linearly with $n_e \varepsilon_B$, so the data constrain the product, and the normalization of every curve in this paper is conditional on the n_e prior. Every part of the ε_B curve depends on the 47 Tuc analogy; what can be done, and is done below, is to bound the analogy’s plausible range astrophysically and to show the answer across that range.

One further input is disclosed rather than changed. The distance prior is the paper set’s fiducial $5.43 \pm 0.05 \text{ kpc}$, while Mahida et al. (2026) adopt the kinematic distance $5.494 \pm 0.061 \text{ kpc}$. The offset is 1.2σ of the prior and would loosen every leg by about 2 per cent, since a more distant source is fainter at fixed luminosity and every predicted flux falls. We hold the fiducial value across this paper set rather than harmonize it silently per paper, and record the direction and size of the effect here so that a reader who prefers the larger distance can apply it.

3.3 Forward model and likelihoods

For each trial (ε_B, M) and nuisance draw θ : $L_{\text{bol}} = \varepsilon_B \dot{M}_B(M, \theta) c^2$. The predicted observables are

$$f_X^{\text{pred}} = f_X L_{\text{bol}} / (4\pi d^2), \quad (4)$$

$$(\nu L_\nu)_{\text{IR}}^{\text{pred}} = f_{\text{IR}} L_{\text{bol}}, \quad (5)$$

$$\log L_R^{\text{pred}} = 0.60 \log(k f_X L_{\text{bol}}) + 0.78 \log(M/M_\odot) + 7.33 + \delta_{\text{FP}}, \quad (6)$$

where $k = 0.610$ converts the 0.5–7 keV band fraction to the 2–10 keV band the fundamental plane is calibrated in (power law, $\Gamma = 2$; [Haggard et al.](#)’s own $\Gamma = 2.3$ gives $k = 0.462$ and is carried as a ladder row in Appendix A), and $\delta_{\text{FP}} \sim \mathcal{N}(0, 0.88 \text{ dex})$ is the plane’s intrinsic scatter as a latent variable. The plane is used only in the forward direction, the direction it was fit in ([Merloni et al., 2003](#)); no inverse regression is performed, which removes the calibration-problem objection that attaches to algebraic inversion of a scattered forward relation ([Plotkin et al., 2012](#)). The likelihood is the product of six Gaussians on the measured values: central radio flux density $0 \pm 1.1 \mu\text{Jy}$, central X-ray flux $0 \pm f_{X,\text{lim}}/1.645$, and four infrared band-peak luminosities $0 \pm \sigma_j$ from Table 1, the infrared terms rescaled to the drawn distance from [Chen et al.](#)’s 5.49 kpc. Where a band’s published product is a limit rather than a measured value, a Gaussian whose σ is that limit divided by the quantile of the confidence the source itself states is the reconstruction we adopt for the two legs published as limits, and Appendix A measures what it costs; the data-availability section requests exactly those numbers from the source teams, and the machinery ingests them unchanged.

The posterior at each M is computed as the plain Monte Carlo mean of the likelihood over the 4×10^4 nuisance draws, which is an unweighted average because the draws come from the prior itself and no proposal reweighting is applied, on a log-spaced ε_B grid (10^{-14} – 10^{-2} , log-uniform prior), and ε_{95} denotes the 95 per cent credible upper limit of the normalized posterior. Re-running the anchors under eight nuisance seeds, each with its own δ stream, puts the Monte Carlo spread of ε_{95} at 1.4 to 9.6 per cent of the mean across the three anchors and the three families, widest for the RIAF family and narrowest for the radio-free curve, and the spread of P_{excl} at 0.13 to 0.88 percentage points. Both sit an order of magnitude under the analysis-choice ranges of Section 5.3. The seed of record sits at the low end of the ε_{95} range and the high end of the P_{excl} range; the eight-seed minima and maxima ship in `figs/ff_v4_seedsread.json`. Upper limits from non-detections are sensitive to the prior’s lower cutoff; Section 5.3 quantifies that dependence (a factor of 9 to 40 when the floor rises from 10^{-14} to 10^{-11}) and every quoted ε_{95} states its grid. This likelihood product, not a distribution of inverted thresholds, is what Paper E’s Bayesian adjudication machinery consumes; the marginal likelihood over ε_B ships as a machine-readable table with the source.

3.4 Accretion-flow families

A single SED bracket cannot serve all flows, and under a common bracket for f_X and f_{IR} the X-ray leg dominates the infrared leg in every draw, leaving JWST unconstraining in this band; an earlier draft of this paper had exactly that structure, and the referee who caught it was right that it made the three-band claim cosmetic. We therefore evaluate three named families with distinct band-fraction priors (Table 3): a radiatively inefficient ADAF-like flow (X-ray-bright among the bands considered), a jet-dominated, synchrotron-peaked flow (infrared-bright, X-ray-faint, the family in which the infrared leg comes closest to binding), and a thin disk (radiatively efficient; included so the reader sees by how many decades it is excluded, and evaluated without the fundamental-plane leg since the plane is a hard-state/quiescent relation).

Table 3: Accretion-flow families. Band fractions are log-uniform within the stated ranges, set conservatively by construction to bracket each family’s characteristic SED shape rather than fit to a specific source.

Family	f_X (0.5–7 keV)	f_{IR} (band peak)	radio leg
RIAF (ADAF-like)	0.03–0.3	0.03–0.3	fundamental plane
Jet-dominated	0.01–0.1	0.1–0.5	fundamental plane
Thin disk	0.05–0.2	0.05–0.3	none

3.5 The natural-flow expectation and its efficiency law

The limit curve becomes a verdict only against a prediction, and the prediction is $\varepsilon_{\text{nat}} = \eta_{\text{ADAF}}(\dot{m}_{\text{hor}}) f_B$, with $\dot{m}_{\text{hor}} = f_B \dot{M}_B / \dot{M}_{\text{Edd}}$ the rate reaching the horizon in Eddington units. Xie and Yuan (2012) give η for hot flows as a piecewise power law,

$$\eta(\dot{m}) = \eta_0 \left(\frac{\dot{m}}{10^{-2}} \right)^a, \quad (7)$$

with (η_0, a) tabulated per branch and per δ , the fraction of turbulent dissipation delivered directly to electrons. In the lowest branch, the one relevant here, a runs from 0.47 to 0.71 and η_0 from 0.065 to 1.58 across $\delta = 10^{-3}$ to 0.5. We adopt their Table 1 as printed and carry δ as an eighth nuisance, drawn uniformly over the four tabulated values, since no measurement of it exists for this object. Like the suppression index s , δ enters ε_{nat} and not the likelihood, so its whole influence runs through the comparison rather than through ε_{95} .

Drafts of this paper to v1.5 used $\eta = 0.1 \min(1, \dot{m}/10^{-2})$, linear in $\dot{m}$, and described it as a single-break simplification of the same fits. It is not a simplification of them: at the rates this problem reaches it is lower than any tabulated δ by one to three decades, because $a = 1$ falls away far faster than $a \approx 0.6$. Correcting it raises the expectation band by about two decades and is the larger of the two reasons the verdict of Section 5 differs from the earlier drafts’.

A caveat travels with the correction and is not removed by it. The horizon-scale rates here are $\dot{m}_{\text{hor}} \approx 10^{-9}$ to 10^{-6} , with a median of 1.8×10^{-8} at the fast-star anchor, while the lowest rate Xie and Yuan compute is 1.6×10^{-5} to 9.4×10^{-5} depending on δ . Every number in this paper’s expectation band therefore extrapolates their lowest branch three to five decades below its computed range, as the linear law it replaces also did. The direction of the extrapolation matters for how the verdict should be read: continuing the power law downward gives the smallest η of any reading, so it is the choice least favourable to exclusion. Holding η flat at the value on the fit’s own lower edge, the opposite bracket, drives every exclusion fraction in Table 6 to 0.99 or above. The extrapolated law is what the headline numbers use, and the two rows below the rule in that table are the bracket.

3.6 Defining the exclusion fraction

Each natural-flow draw is a fully specified model: $\varepsilon = \varepsilon_{\text{nat}}(\theta)$ with that draw’s own nuisances. It predicts a flux in every observed band, and the measured data, zero central flux at the published σ_j , then have a joint chi-square under it, $\chi^2 = \sum_b (\text{pred}_b / \sigma_b)^2$ over the six terms of Section 3.3, distributed as χ_6^2 if that draw is the truth. We call a draw excluded when the data reject it at the 5 per cent level, and report

$$P_{\text{excl}} = P[p(\chi^2) < 0.05] \quad (8)$$

over the nuisance draws, with its Monte Carlo standard error.

Drafts to v1.5 reported a different quantity under the same name: the fraction of natural draws exceeding ε_{95} , with ε_{95} itself marginalized over those same draws. That statistic compares a distribution against a

scalar built from it, uses common random numbers on both sides, and has no coverage reading; it also produced the artefact that widening the density prior cost the verdict almost nothing, since a log-normal's median does not move when its width does, so the exceedance count barely changed. Equation 8 avoids all three. It is per-draw, it uses the likelihood rather than a summary of it, and it never reads ε_{95} , which is why the efficiency prior's lower cutoff, a factor of 9 to 40 on ε_{95} , cannot touch it. The old statistic is retained in Table 4 under its correct name, recomputed on independent draws so that the common random numbers are gone; it runs 0.03 to 0.04 above the posterior-predictive figure at the two lower anchors and reaches 1.00 at the high anchor where the posterior-predictive value is 0.99.

4 The gas supply, bounded from above

The n_e prior needs an astrophysical anchor, and the cluster's own stellar winds supply one. Drafts of this paper to v1.5 used that budget as a floor. A recount of its inputs (`figs/ff_calc7_windfloor.py`) shows that it is a ceiling, and the correction runs to two decades.

The chain is the standard one. Inside the influence radius $r_{\text{infl}} = GM/\sigma^2 = 0.39 \text{ pc}$ of a $4 \times 10^4 M_\odot$ hole, a Plummer model of the cluster ($M = 4 \times 10^6 M_\odot$, $a = 3.74 \text{ pc}$) encloses $4.5 \times 10^3 M_\odot$ of stars, about 7.4×10^3 stars at a luminous-mass mean of $0.6 M_\odot$. Of those, the Kroupa weight of the turnoff window is 0.035 and the fraction of a turnoff star's main-sequence life spent above $10^3 L_\odot$ is $(0.35+0.07)/13 = 0.032$, so the standing giant population inside r_{infl} is 8.4 stars, and 5.4 to 11.4 across the phase-lifetime and mean-mass variations tested. At $10^{-8} M_\odot \text{ yr}^{-1}$ each the aggregate injection is $8.4 \times 10^{-8} M_\odot \text{ yr}^{-1}$, and if nothing removes it the steady state at outflow speed $v_{\text{out}} \sim 50 \text{ km s}^{-1}$ is

$$n_e \sim \frac{\dot{M}_{\text{wind}}}{4\pi r_{\text{infl}}^2 \mu_e m_p v_{\text{out}}} \simeq 0.030 \text{ cm}^{-3}, \quad (9)$$

with 0.019 to 0.040 cm^{-3} over the same variations, refilled on the crossing time $r_{\text{infl}}/v_{\text{out}} \sim 10^4 \text{ yr}$. Earlier drafts carried 10^3 giants, $10^{-5} M_\odot \text{ yr}^{-1}$ and $n_e \sim 3 \text{ cm}^{-3}$ at this point. The arithmetic linking those three was correct, and the recount reproduces 3.5 cm^{-3} from the 10^3 -giant input; the input itself was too large by a factor of 119, having counted the stars rather than the ones currently on the giant branch.

The recount changes three things, and two of them cut against this paper's earlier reading.

The wind budget bounds the density from above. A no-removal steady state is the most gas the winds can sustain, so 0.030 cm^{-3} bounds the wind-fed density from above. The $n_e \geq 0.05 \text{ cm}^{-3}$ floor of earlier drafts sits a factor 1.7 above it. The claims that accompanied it, that the floor lay a factor 60 below the no-removal steady state and a factor 4.6 below the 47 Tuc value, are withdrawn, and with them the statement that the floor returns the widened prior almost to baseline. Section 5.3 now reports a truncation at $n_e \geq 0.02 \text{ cm}^{-3}$, inside the recounted range, and states what it costs rather than what it recovers.

The classical gas problem does not arise here. Globular clusters are gas-poorer than their stellar winds imply, and in 47 Tuc a removal agent, most plausibly millisecond-pulsar winds and heating, is required to explain it (Freire et al., 2001; Abbate et al., 2018). At 8 giants inside 0.39 pc there is little to remove: the wind supply and the 47 Tuc analogy value differ by a factor of 8 in the direction that needs no removal agent at all.

The recount bears on the prior's centre. The 0.23 cm^{-3} median is imported from 47 Tuc, and $\omega \text{ Cen}$'s own giants inside the influence radius cannot supply it without a source this budget omits. An independent route reaches a similar figure: Haggard et al. (2013) adopt $n_e = 0.038 \text{ cm}^{-3}$ for this cluster from a Pfahl–Rappaport wind budget, within 30 per cent of the recount here despite a different formalism and radius. Two readings survive. Either the central density is set by material the influence-radius wind budget does not track, in which case the 47 Tuc analogy stands and the prior is right, or it is set by the




Figure 1: 95 per cent credible upper limits on the Bondi radiative efficiency ε_B of any central accretion flow in ω Cen, per accretion-flow family, with the expectation band for a natural outflow-suppressed hot flow (grey: 5–95 percentile of $\varepsilon_{\text{nat}} = \eta_{\text{ADAF}}(\dot{m}) f_B$ over the same nuisance draws, on the Xie and Yuan (2012) efficiency of §3.5) rising while the limits fall. The band now sits above the RIAF limit across the whole range rather than crossing it inside the contested masses, which is why the exclusion fractions of Table 4 are high at every anchor; the vertical separation is not itself the exclusion fraction, which §3.6 defines from the likelihood rather than from this figure. Drafts to v1.5 drew this band two decades lower, on an efficiency linear in $\dot{m}$. Vertical lines mark the pulsar-timing point-mass cap (Bañares-Hernández et al., 2025), the fast-star lower bound (Häberle et al., 2024), and the upper kinematic range. All curves conditional on the n_e prior of Table 2.

local supply, in which case the prior median is high by a factor of 6 to 8 and every limit in this paper is correspondingly loose. We do not collapse that into one number. The prior is left at its published centre so that the curves remain comparable with the earlier drafts and with the comparison literature, the n_e median/10 row of Table 5 is the quantitative price of the second reading, and Section 7 states again that a pulsar dispersion-measure gradient is the measurement that ends the argument. On the second reading the density evidence sits at the weaker end of this paper’s exclusion range, which is the direction Table 6’s prior-shift row prices.

5 Results

Figure 2 shows the full joint posterior behind these curves over the two parameters of interest and the two dominant nuisances, RIAF family only.

Figure 1 and Table 4 carry the content. Three statements summarize them.

The radio leg dominates wherever the fundamental plane applies. The plane’s mass term rises as $M^{0.78}$ while the Bondi denominator rises as M^2 , which sets the direction; the realized slope of the marginalized limit is shallower than that single-draw asymptote, $\varepsilon_{95} \propto M^{-2.1}$ between the two upper anchors and $M^{-2.2}$ fitted across the whole grid, because the plane’s 0.88-dex scatter and the two flux legs flatten the response. Drafts to v1.5 quoted the asymptote itself, $M^{-3.3}$, as though it were the realized slope. The




Figure 2: Joint posterior corner plot for the RIAF family over the two parameters of interest ($\log_{10} \epsilon_{95}$, the Bondi efficiency whose 95 per cent credible upper limit Figure 1 reports, and $\log_{10} M$) and the two dominant nuisances ($\log_{10} n_e$ and the distance). Samples from the shipped Monte Carlo (same seed, priors, and likelihood; ϵ_B log-uniform over the efficiency grid, per-sample weights $e^{\ln L}$; effective sample size 1.05×10^5 of 4×10^5 draws; script `figs/fF_corner_plot_v2.py`). This figure alone adds a prior on M , log-uniform over the mass grid’s range, because a corner plot needs a joint sample and the rest of the paper deliberately reports ϵ_{95} at fixed M instead of marginalizing over a mass prior the kinematic literature does not agree on. The M panel therefore shows what that particular choice implies and is not a result of the paper. Contours enclose 68 and 90 per cent of the posterior weight. The ϵ_B – M degeneracy is the flux-limit constraint: the predicted X-ray, infrared, and radio luminosities all scale as $\epsilon_B M^2$ (Bondi) or steeper, so the posterior’s upper envelope in ϵ_B falls as M rises, the same structure Figure 1 displays curve by curve. The distance posterior reproduces its prior: the limits are flux measurements, and the drawn distance enters the prediction rather than the limit.

Table 4: ε_{95} at the three mass anchors (evaluated at the anchor masses, not grid neighbours), and the posterior-predictive exclusion fraction P_{excl} of §3.6: the fraction of natural-flow draws whose own predicted observables the measured data reject at the 5 per cent level. Monte Carlo standard errors on P_{excl} are below 0.003 throughout. The last column is the prior-exceedance fraction that drafts to v1.5 reported under the name P_{excl} , recomputed here on independent draws.

$M (M_{\odot})$	RIAF	jet	thin disk	RIAF, no radio	P_{excl} (RIAF)	prior exc. (RIAF)
6,000	6.1×10^{-10}	1.4×10^{-9}	2.3×10^{-8}	2.5×10^{-8}	0.78	0.82
8,200	2.9×10^{-10}	6.5×10^{-10}	1.3×10^{-8}	1.4×10^{-8}	0.85	0.89
4×10^4	1.1×10^{-11}	2.1×10^{-11}	7.3×10^{-10}	7.9×10^{-10}	0.99	1.00

radio-driven curve separates from the radio-free curve by 1.9 decades at the top of the mass range. This is the single-object version of the population-level result of Maccarone et al. (2005), Strader et al. (2012), and Tremou et al. (2018): for IMBH masses, radio provides the tightest constraint on quiescent accretion. The plane’s extrapolation into deep quiescence is contested (Plotkin et al., 2012), its 0.88-dex scatter is carried as a latent variable rather than a point estimate, and readers who distrust it entirely should quote the red curve: $\varepsilon_{95} = 1.4 \times 10^{-8}$ at 8,200 $M_{\odot}$ and 7.9×10^{-10} at $4 \times 10^4 M_{\odot}$, from Chandra and JWST with no shared emission model.

Against the natural-flow expectation, the verdict is no longer mass-dependent in the way earlier drafts reported. The grey band is the prediction, $\varepsilon_{\text{nat}} = \eta_{\text{ADAF}}(\dot{m}) f_{\text{B}}$, evaluated draw by draw on the same nuisances, with η_{ADAF} now the piecewise fit of Xie and Yuan (2012) as summarized in Yuan and Narayan (2014) and δ carried as a nuisance over the four values they tabulate (§3.5). This expectation carries no δ_{FP} term, unlike ε_{95} for the RIAF and jet families, whose radio-leg prediction does (Section 3.3); the fundamental-plane scatter widens the limit but not the band it is compared against, so a reader who distrusts the plane should read the exclusion fractions off the radio-free column rather than treating the asymmetry as a wash. The band rises with mass while every limit falls, and it now sits above the RIAF limit across the whole contested range rather than crossing inside it. At the pulsar-timing cap and the fast-star anchor the data reject 78 and 85 per cent of natural-flow draws under the RIAF family, and 20 and 28 per cent with the radio leg removed; at $4 \times 10^4 M_{\odot}$ the figures are 99 and 76 per cent. Drafts to v1.5 reported 19, 31 and 90 per cent for the RIAF column, on a linear efficiency law and a statistic that compared the natural draws against a scalar built from those same draws; both are corrected here, and both corrections move the verdict the same way. Table 6 carries the figures across the sensitivity ladder. A Bondi-fed hot flow at the fiducial density is now disfavoured across the contested mass range and not only at its top, which is a stronger claim than the one the earlier drafts made and a more fragile one, since it rests on extrapolating a published efficiency fit well below the rates it was computed for. Section 3.5 states what that extrapolation is worth and Table 6’s last two rows bracket it.

The thin disk is dead everywhere. ε_{95} for the disk family sits four to five decades below the $\eta \sim 0.06$ –0.4 a disk radiates at; no Bondi-fed thin disk at any mass in the range survives, quantifying what every prior single-band paper assumed informally.

5.1 The exclusion fraction as a continuous function of mass

Table 4 evaluates P_{excl} at three masses because those are the masses the kinematic literature argues about. The Monte Carlo evaluates it on the whole grid, and the continuous version (Figure 3) is the more useful object for a reader who holds a different mass prior than any of the three anchors. It also locates the transitions. Under the RIAF family the exclusion fraction passes 50 per cent at $2.5 \times 10^3 M_{\odot}$ and 90 per cent at $1.1 \times 10^4 M_{\odot}$; under the jet family the same thresholds fall at 3.3×10^3 and $1.4 \times 10^4 M_{\odot}$.




Figure 3: Fraction of natural-flow parameter space excluded, P_{excl} of Eq. 8, as a continuous function of mass: every natural-flow draw scored against the measured data at every grid mass rather than at three. Drafts to v1.5 plotted the prior-exceedance fraction $P[\varepsilon_{\text{nat}} > \varepsilon_{95}]$ here; that curve is carried alongside in the shipped output. Vertical lines mark the pulsar-timing point-mass cap (Bañares-Hernández et al., 2025), the fast-star lower bound (Häberle et al., 2024), and the upper kinematic range. The three anchor values of Table 4 sit on these curves by construction. Coordinates from `figs/fF_expand_exclcurve_v2.json`. The quantity plotted is the posterior-predictive fraction of §3.6; that file also carries the prior-exceedance curve the v1.5 figure plotted, so the change is auditable.

Using no radio information they move to 1.7×10^4 and $7.9 \times 10^4 M_{\odot}$. Both RIAF thresholds now sit below the pulsar-timing point-mass cap, so the mass-dependence that drafts to v1.5 built the verdict on has largely gone: on the corrected efficiency law the constraint bites across the contested range rather than only at its top.

5.2 Which band binds which family

Figure 4 shows what the family definitions of Section 3.4 change. At each anchor the figure evaluates the forward model at that family’s own ε_{95} and plots the predicted νL_{ν} in every observed band against the published limit in that band, so the three instruments and the three families appear on one axis.

The radio leg binds every family that carries it, and by a wide margin. For the RIAF family the median predicted radio flux density at ε_{95} sits at 1.5 times the ATCA 5σ point at $8,200 M_{\odot}$ and 4.8 times it at $4 \times 10^4 M_{\odot}$, while the X-ray prediction sits at about 1 per cent of the Chandra bound and the tightest infrared filter at 0.2 per cent of its own. For the jet family, the one built to be infrared-bright, the radio ratios are 1.2 and 3.6 and the F770W ratio reaches 0.8 and 0.6 per cent: better than under the RIAF band fractions by a factor of five, and still short of the plane by more than two decades. Drafts to v1.5 described JWST as genuinely binding for this family. It does not, on their own numbers or these, and the correction of Section 2 to the infrared comparison widens the gap rather than closing it.

The thin disk carries no radio leg, and there the X-ray band binds alone, at 0.43 and 0.58 of the Chandra bound against F770W’s 0.09 and 0.12. Drafts to v1.5 reported those two legs sharing the constraint within 10 per cent of each other; that reading came from comparing a band-peak prediction against a band-integrated limit, and once the two are the same quantity the infrared sits a factor of five below the X-ray in every family. The honest summary is that this is a radio-plus-X-ray bound with an infrared consistency check, not a three-way joint constraint, and Section 7 prices deeper mid-infrared accordingly.




Figure 4: Predicted band luminosities at each family’s own ε_{95} , against the published limits (black bars, arrows marking the excluded direction). Points are medians over the 4×10^4 nuisance draws, whiskers the 5th to 95th percentiles. The thin-disk family has no radio point because the fundamental-plane leg is not applied to it (Section 3.4). The radio limit is drawn at 5σ ($5.5 \mu\text{Jy}$) so that it is comparable with the X-ray and infrared bounds, which are published as limits; the likelihood itself uses the $1.1 \mu\text{Jy}$ rms as a Gaussian σ and is unaffected. Infrared points are band-integrated luminosities, the same quantity as the published limits (§2); drafts to v1.5 plotted band-peak predictions against band-integrated limits here. Script `figs/fF_expand_v2.py`, numbers in `figs/fF_expand_sed_v2.json`.

Ratios above unity for the radio band are expected and are not a detection claim: ε_{95} is a 95 per cent posterior quantile marginalized over the plane’s 0.88-dex latent scatter, so at that ε_B the median draw predicts a flux the image would have seen while a substantial tail does not. The five-percentile whiskers, which reach two decades below the medians in the radio and a factor of ten in the other bands, are that scatter. The infrared points were flat across the four filters in drafts to v1.5, since the model predicts one band-peak νL_ν and the filters differed only in their limits. They are no longer flat: each is now converted to that filter’s band-integrated luminosity by its own $\Delta\lambda/\lambda_c$ (Table 1), which is the quantity the published limits are. F770W still sets the infrared constraint, by a factor of two over the next filter, because it combines the tightest limit with the narrowest fractional bandwidth.

5.3 Sensitivity budget

Table 6 prices the headline verdict rather than the limit behind it, and the two no longer move together in the way drafts to v1.5 described. Widening the n_e prior costs the RIAF exclusion fraction 0.09 at the fast-star anchor and 0.05 at the high anchor; shifting the prior centre down a decade costs 0.27 at the fast-star anchor while leaving the high anchor at 0.94, because at $4 \times 10^4 M_\odot$ the natural-flow band clears the limit by enough margin that a factor of ten in density does not close it. The prior-shift row is still the widest single move, and §4’s recount is an argument that it is the more likely direction rather than a worst case. Across every convention choice in this paper the RIAF exclusion fraction at $4 \times 10^4 M_\odot$ stays above 0.93 and at $8,200 M_\odot$ above 0.57; the radio-free column is the fragile one, spanning 0.02 to 0.40 at the fast-star anchor and 0.33 to 0.81 at the high anchor, and it should be read with that range attached. Two rows below the rule price the efficiency prescription itself rather than a convention: holding η at the lower

Table 5: Sensitivity of the RIAF-family ε_{95} to the dominant analysis choices. Each row changes one ingredient from the baseline.

Variant	$6,000 M_{\odot}$	$8,200 M_{\odot}$	$4 \times 10^4 M_{\odot}$
baseline	6.1×10^{-10}	2.9×10^{-10}	1.1×10^{-11}
n_e prior widened to 1.5 dex	3.5×10^{-9}	1.7×10^{-9}	7.4×10^{-11}
n_e prior median $\div 10$	3.7×10^{-9}	1.7×10^{-9}	4.4×10^{-11}
widened + wind truncation ($n_e \geq 0.02$)	1.0×10^{-9}	5.1×10^{-10}	1.8×10^{-11}
widened + retired floor ($n_e \geq 0.05$)	6.7×10^{-10}	3.3×10^{-10}	1.2×10^{-11}
ε_B -grid floor raised to 10^{-11}	5.5×10^{-9}	3.3×10^{-9}	4.4×10^{-10}
IR band correction removed (the v1.5 leg)	5.7×10^{-10}	2.8×10^{-10}	1.0×10^{-11}
k at $\Gamma = 2.3$, the source's own index	7.3×10^{-10}	3.5×10^{-10}	1.2×10^{-11}

Table 6: The posterior-predictive exclusion fraction under the same variants, for the RIAF family and for the radio-free curve. Table 5 moves only ε_{95} ; here the natural-flow draws move with the nuisances, since both are built from one set of draws. The ε_B -grid floor has no row: the statistic of §3.6 never reads ε_{95} , so the efficiency prior's lower cutoff cannot reach it, which is one practical gain of the reconstruction. Monte Carlo standard errors are below 0.003 throughout.

Variant	RIAF			RIAF, no radio		
	6,000	8,200	4×10^4	6,000	8,200	4×10^4
baseline	0.78	0.85	0.99	0.20	0.28	0.76
n_e prior widened to 1.5 dex	0.69	0.75	0.94	0.33	0.38	0.64
n_e prior median $\div 10$	0.47	0.58	0.94	0.02	0.04	0.33
widened + wind truncation ($n_e \geq 0.02$)	0.74	0.81	0.98	0.33	0.38	0.69
widened + retired floor ($n_e \geq 0.05$)	0.79	0.85	0.99	0.34	0.40	0.75
X-ray limit read two-sided 95 per cent	0.78	0.85	0.99	0.21	0.30	0.78
X-ray limit read as 3σ	0.78	0.85	0.99	0.25	0.34	0.81
k at $\Gamma = 2.3$	0.76	0.84	0.99	0.20	0.28	0.76
IR completeness row 99.7 per cent	0.78	0.85	0.99	0.20	0.28	0.76
IR completeness row 68 per cent	0.78	0.85	0.99	0.21	0.29	0.77
IR tightest single filter only	0.80	0.87	0.99	0.23	0.31	0.79
IR band correction removed (the v1.5 leg)	0.78	0.85	0.99	0.25	0.33	0.81
Bondi denominator, Paper E convention	0.73	0.82	0.99	0.15	0.22	0.70
η held at the XY12 fit's own lower edge	0.99	1.00	1.00	0.87	0.92	1.00
η linear in $\dot{m}$ (the v1.5 form)	0.36	0.48	0.92	0.01	0.02	0.27

edge of the range Xie & Yuan actually computed, instead of extrapolating their fit four to six decades below it, raises every entry towards unity, while the linear η of drafts to v1.5 lowers them to the values those drafts reported. The extrapolation this paper adopts is the weaker of the two readings of the cited fit, and the verdict is stated on that reading.

Table 5 is the quantified version of Section 3.2’s caveat. Widening the n_e prior by a decade in each direction, or dividing its median by ten, each relax ε_{95} by a factor of about six: the curve’s normalization belongs to the gas-density assumption, in both its width and its centre, and no claim in this paper escapes that conditioning. The damage is bounded on one side and, since v1.5, no longer on the other. Truncating the widened prior at $n_e \geq 0.02 \text{ cm}^{-3}$, inside the range the wind recount of Section 4 supports, recovers about two thirds of the widening in the log: ε_{95} at the fast-star anchor moves from 1.7×10^{-9} back to 5.1×10^{-10} against a baseline of 2.9×10^{-10} . Drafts to v1.5 used a truncation at 0.05 cm^{-3} and reported that it returned the widened prior almost to baseline; that floor lies above the wind steady state and the claim is withdrawn with it. The ε_B -prior floor matters as it always does for non-detections, by a factor of 9 at the two lower anchors and 40 at $4 \times 10^4 M_\odot$ when the grid floor rises three decades; drafts to v1.5 quoted that as a single factor of about ten, which holds only at the lower anchors. Every ε_{95} here is quoted on the 10^{-14} grid, the machine-readable likelihood tables let any reader impose their own, and the exclusion fractions of Table 6 are now independent of that choice entirely.

5.3.1 Which nuisance carries the spread

Table 5 varies analysis choices. The nuisance parameters can be interrogated separately, by re-running the same posterior with one of them held at its prior’s central value while the others are drawn, and again with only that one drawn and the others held (Figure 5). The two views answer different questions: the first asks how much of the limit’s looseness a given parameter is responsible for, the second asks how much looseness that parameter would produce on its own. Seven parameters appear; the eighth, δ , is absent for the same structural reason as s , since it enters ε_{nat} and not the likelihood, so freezing it cannot move ε_{95} .

The answer is that the fundamental plane’s latent scatter carries the spread. Freezing δ_{FP} at zero tightens ε_{95} by a factor of 13 at $8,200 M_\odot$ and 29 at $4 \times 10^4 M_\odot$; drawing it alone, with everything else fixed, reproduces a limit 18 and 61 times looser than the fully frozen case. No other parameter comes within an order of magnitude. The gas density is second and is a distant second in this accounting: freezing it tightens ε_{95} by 21 to 25 per cent, and drawing it alone loosens the frozen limit by a factor 1.6 to 2.4. The X-ray band fraction contributes 8 to 10 per cent, the sound speed under 1 per cent, and the distance nothing measurable, as expected from a 1 per cent prior. The infrared band fraction moves the RIAF limit by 0.4 per cent in the opposite direction, which is the signature of a leg that never binds in this family, and the band correction of Section 2 makes it smaller still than the 2 per cent drafts to v1.5 reported.

The distance prior and the suppression index are inert here, for structural reasons worth naming. The distance prior is narrow enough that its 1 per cent width cannot register against the plane’s 0.88 dex. The suppression index s does not enter the likelihood at all: it sets the natural-flow expectation ε_{nat} against which the limit is compared, not the limit itself, so its entire influence in this paper runs through P_{excl} and the grey band of Figure 1 rather than through ε_{95} .

This does not contradict Section 3.2’s statement that the gas density owns the result. The two statements are about different things. Within the stated priors, the plane’s scatter is what widens the posterior; the density’s danger is in the choice of prior centre, which no freezing test can expose, and which Table 5’s prior-shift rows price at a factor ~ 6 . A reader who distrusts the plane has the radio-free curve; a reader who distrusts the density analogy has the shift rows and the wind recount of Section 4, which points the same way.




Figure 5: Per-nuisance contribution to the RIAF-family limit, at two mass anchors. Left: ε_{95} with that nuisance held at its prior’s central value, divided by the baseline ε_{95} ; a point far to the left means that parameter’s spread was inflating the limit by that factor. Right: ε_{95} with only that nuisance drawn, divided by the limit with all seven frozen. Neither panel is a variance decomposition, since ε_{95} is a posterior quantile rather than a sum of independent terms, and the two panels do not multiply back to unity. The distance and the suppression index sit on the reference line in both panels; s does so because it enters the natural-flow expectation and not the likelihood. Numbers in `figs/fF_expand_nuisance_v2.json`.

5.4 Where this sits against the survey population

The single-cluster depth used here is best read against the 50-cluster MAVERIC survey (Tremou et al., 2018), which is the reference population for radio IMBH limits and the source of the field’s null (Table 7). Two cautions travel with the comparison. First, each entry assumes its own distance: Tremou et al. (2018) adopt 4.9 kpc for ω Cen, Haggard et al. (2013) 5.2 kpc, Chen et al. (2026) and Mahida et al. (2026) 5.49 kpc, against this paper’s 5.43 kpc prior. The spread is ordinary paper-to-paper variation and it propagates as d^2 into any luminosity, so the flux densities compare directly while the derived quantities do not. Second, the survey’s mass limits are obtained by inverting the fundamental plane, the step Section 3.3 declines to take; they are quoted here as the literature’s own currency and are not ε_B and not convertible to it without the full nuisance treatment.

Table 7: Radio limits on central sources in globular clusters, as published. Mass limits are the sources’ own fundamental-plane inversions, quoted at 3σ ; they are not ε_B . Flux densities are directly comparable, derived quantities are not.

Dataset	Depth	Derived limit	Assumed d
MAVERIC ATCA, ω Cen (Tremou et al., 2018)	$< 8.8 \mu\text{Jy} (3\sigma)$	$M < 1000 M_\odot$	4.9 kpc
MAVERIC VLA stack, 24 clusters	$0.65 \mu\text{Jy beam}^{-1}$	$M < 800 M_\odot$	per cluster
MAVERIC ATCA stack, 14 clusters	$1.42 \mu\text{Jy beam}^{-1}$	$M < 970 M_\odot$	per cluster
This paper’s radio input (Mahida et al., 2026)	$1.1 \mu\text{Jy beam}^{-1}$ rms	see Fig. 1	5.49 kpc

The ATCA campaign of Mahida et al. (2026) is deeper on one cluster than the MAVERIC stack is on 24, which is what makes a per-object posterior worth building at all. Three ambiguities in the survey paper are carried rather than resolved: its ω Cen flux limit appears as $8.8 \mu\text{Jy}$ in Table 2 and $8.9 \mu\text{Jy}$ in Section V.2.1; its VLA stack mass limit is given as both < 800 and $< 730 M_\odot$ in a single sentence, and we quote the weaker; and whether ω Cen enters the 14-cluster ATCA stack is not stated, though the paper’s

explicit exclusion list omits it. Its comparison of the ω Cen radio limit against Haggard et al. (2013) also cites the X-ray limit as $1.7 \times 10^{30} \text{ erg s}^{-1}$ where Haggard et al. state 1.6×10^{30} , which looks like a transcription rounding. None of this affects the numbers of Section 5, which use no MAVERIC input; it affects only what the reader should do with the row above.

6 The duty-cycle loophole

The three inputs sample time very differently, and Mahida et al. (2026)’s Table 1 makes the radio sampling explicit: 25 observing blocks across three ATCA projects, 2010 January 22 to 2024 December 27, summing to 177.19 hr on source (CX556, 20 blocks, 147.09 hr; C2877, 3 blocks, 11.79 hr; C2158, 2 blocks, 18.31 hr). Those rows exceed the ~ 170 hr of the paper’s abstract and the 172 hr of its Section II.3; we report the sum as printed and leave the reconciliation to that paper, noting that prose totals quoted after flagging would run below a raw table sum rather than above it. Chandra contributes four exposures in two epochs twelve years apart, JWST a single epoch.

Two consequences follow on different timescales. For variability slower than the sampled span, the limits above apply to the time-average $\delta \varepsilon_{\text{on}}$ of a flow active a fraction δ of the time at ε_{on} , so $\varepsilon_{\text{on}} \lesssim \varepsilon_{95}/\delta$; the radio image is a weighted combination over the full fourteen-year span, so this is the operative statement there. For variability faster than a block, the source must be off during every epoch to escape, which for $N \simeq 30$ independent epochs (25 radio blocks, four Chandra exposures, one JWST visit) gives a miss probability $(1 - \delta)^N$ reaching 50 per cent only near $\delta \approx 0.02$. This order-of-magnitude estimate counts every block or exposure as one independent trial regardless of its duration relative to the assumed flare timescale τ ; Section 6.1’s block-by-block treatment does not carry that assumption and supersedes this estimate in rigor. The excluded region is therefore the band $\varepsilon_{\text{on}} > \varepsilon_{95}/\delta$ down to a few per cent duty cycle, and effectively nothing below that: a flow that spends 99 per cent of its time off is constrained only by the luck of the sampling. An earlier draft put the epoch count at $N \sim 6$ and the sampling break at $\delta \approx 0.1$, having treated the radio campaign as one epoch. This is the standard loophole of every quiescent-limit paper, here made explicit because Paper E’s tidal-disruption arithmetic predicts exactly such intermittency on 10^7 -yr recurrence with 10^{-4} duty, far below the reach of any current sampling; nothing in this paper’s nulls bears on that channel, in either direction.

6.1 What the block list actually excludes

The block list supports a stronger statement than a miss probability, because each block is an image in its own right. Scaling the combined image’s sensitivity by on-source time, $\sigma_i = 1.1 \mu\text{Jy} \sqrt{177.19 \text{ hr}/t_i}$, the 25 blocks reach 4.5 to $12.0 \mu\text{Jy}$ rms with a median of $5.3 \mu\text{Jy}$ over a median 7.6 hr on source, so a single block detects a source at 5σ near $27 \mu\text{Jy}$ while the combined image reaches $5.5 \mu\text{Jy}$. That scaling assumes each on-source hour is equally sensitive across the three projects and both observing frequencies, which Mahida et al. (2026) do not state block by block; it is the only scaling their published table supports. The block statistics in this paragraph regenerate from `figs/ff_epochs.json`, the export of that table: 25 blocks spanning 1.48 to 10.70 hr on source, rms 4.48 to $12.04 \mu\text{Jy}$, median $5.31 \mu\text{Jy}$.

A flare train is then a two-parameter object: amplitude S and duration τ , recurring with period T_{rec} at an unknown phase. Laying such a train over the real block timeline (2010 January 22 to 2024 December 27, a 14.9-yr span containing 177.19 hr of on-source time) and asking whether any single block’s τ -averaged flux exceeds its own 5σ point, or the campaign mean exceeds $5.5 \mu\text{Jy}$, gives a detection probability over phase. Table 8 reports, for each (S, τ) , the longest recurrence interval that would still have been caught in at least 95 per cent of phases. Anything recurring faster than that is excluded; anything slower is not.

Table 8: Longest flare recurrence interval excluded at 95 per cent of phases, from the 25-block ATCA timeline (4000 fixed-seed realizations per cell; script `figs/ff_expand.v1.py`, seed 20260820). Entries are hours; a dash means no recurrence interval on the grid is excluded for that amplitude and duration. The implied duty cycle at each threshold, τ/T_{rec} computed at the unrounded threshold, is in parentheses; recomputing it from the rounded hours in this table gives a value differing by up to 2 per cent. Thresholds are located on a recurrence grid of 45 logarithmic points spanning 6 hr to 2×10^5 hr, a 27 per cent step, which sets the precision of every entry.

Flare amplitude	$\tau = 0.5$ hr	2 hr	8 hr	24 hr
10 μJy	–	–	12 (0.655)	31 (0.763)
30 μJy	–	8 (0.263)	25 (0.322)	130 (0.184)
100 μJy	8 (0.066)	31 (0.064)	130 (0.061)	209 (0.115)
300 μJy	15 (0.032)	81 (0.025)	130 (0.061)	209 (0.115)
1 mJy	81 (0.006)	81 (0.025)	130 (0.061)	209 (0.115)

Low-amplitude trains are excluded only when they are almost always on: a 10 μJy source is invisible in any single block and can be caught only through the campaign mean, which requires a duty cycle above 0.7. Bright, brief trains are bounded by the sampling rather than by the sensitivity: at 300 μJy and above, for durations of 2 hr and longer, the exclusion saturates near a recurrence interval of 3 to 9 days, because 177 hr of on-source time spread over 14.9 yr covers 0.14 per cent of the span, and a flare that recurs more slowly than about a week simply misses every block most of the time no matter how bright it is. Between those regimes the campaign excludes real parameter space: a 100 μJy , 2-hr flare is excluded for recurrence faster than 31 hr, a duty cycle of 6×10^{-2} .

The scaling’s own uncertainty can be priced. Drawing per-block sensitivity log-normally about it, at fixed combined campaign sensitivity, since the 1.1 μJy combined rms is measured rather than assumed, moves the table where the single-block channel sets the threshold and leaves it alone where the campaign mean does. At 0.2 dex of scatter, two of the fourteen well-determined cells move by one grid step: the 30 μJy , 24-hr threshold falls from 130 to 81 hr and the 100 μJy , 2-hr threshold from 31 to 25 hr. Both 10 μJy rows hold to the digit at every scatter level tested, up to 0.3 dex, because a 10 μJy source is caught only through the campaign mean, whose depth is measured. The floor of the excluded region is likewise unaffected at ± 30 per cent per-block rms, and reaches a duty cycle of 1×10^{-2} only at 0.3 dex. The caveat this paragraph opens with therefore attaches to the mid-amplitude, long-duration cells rather than to the faint end, which the source paper’s combined image already constrains directly. Numbers in `figs/ff_calcF1_duty.band.json`.

The floor of the excluded region is therefore a duty cycle of 6×10^{-3} , reached at 1 mJy and half an hour, and it is set by sampling coverage rather than by image depth. This tightens the order-of-magnitude estimate above by a factor of a few and it changes the instrument implication: for the flare channel, more epochs at the present depth would buy more than the same time added to a single deep image, which is the opposite of the priority ordering that Section 7 gives for the steady channel. Paper E’s predicted tidal-disruption intermittency, at 10^{-4} duty on 10^7 -yr recurrence, remains two decades below the floor and 12 decades outside the recurrence range, so the conclusion of the preceding paragraph is unchanged: this campaign says nothing about that channel.

7 Forecasts

Three measurements would harden or move the curves, in descending order of leverage per unit effort.

A measured core density. The 19-pulsar timing set (Colom i Bernadich et al., 2026) makes an ω Cen

dispersion-measure gas detection feasible for the first time, by the method that measured 47 Tuc (Freire et al., 2001; Abbate et al., 2018); the caveat, that DM gradients constrain the column through the pulsar volume rather than the density at the hole, is real and enters as geometry. The forecast gain on ε_{95} itself is modest (a ± 10 per cent density measurement tightens the baseline by ~ 22 per cent, since the 0.5-dex prior is already informative). A measurement converts every curve from analogy-conditional to measured, removes the largest single objection to the enterprise, and, if the density comes in low, legitimately weakens the bounds when the conditions are tested.

Table 9 is the requirement that forecast implies, computed from the published timing set. All 19 pulsars have measured dispersion measures; eight have them from full timing solutions with formal uncertainties of 5×10^{-5} to 1.3×10^{-3} pc cm $^{-3}$, and five of those eight also have a measured first DM derivative. Timing precision is not the obstacle.

Geometry sets what a dispersion-measure gradient can reach. Gas filling the cluster core, a uniform sphere of the catalogue core radius $2.37'$ (Harris, 1996), or 3.74 pc at 5.43 kpc, imprints a central-chord excess of 0.75 pc cm $^{-3}$ at $n_e = 0.1$ cm $^{-3}$, and 13 of the 19 sight lines pass inside it. Gas confined to the influence radius of the hole, the 0.4 pc region the Bondi rate is actually about, imprints 0.08 pc cm $^{-3}$ and no published sight line passes through it: the innermost pulsar, H at $0.56'$, projects to 0.88 pc. A DM determination therefore measures the core-filling component and reaches the Bondi-scale density only through an assumed profile, which is the caveat named above, now with a number attached.

The sight-line scatter sets how well it can be measured. The 19 published DMs span 94.3 to 102.6 pc cm $^{-3}$ with an r.m.s. of 2.9 pc cm $^{-3}$, three orders of magnitude above the per-pulsar measurement precision and four times the signal a core-filling $n_e = 0.1$ cm $^{-3}$ would produce. Fitting the chord template plus a free foreground constant to the 19 measurements by least squares returns an amplitude of -0.33 ± 0.24 cm $^{-3}$, consistent with zero and formally negative, with a residual r.m.s. of 2.8 pc cm $^{-3}$. We do not quote the implied upper limit as a constraint on the cluster's gas: a model whose residuals exceed its measurement errors by three orders of magnitude is falsified by its own fit, and the formal error bar of a falsified model is not a credible interval. What the fit does establish is the requirement. At the present per-sight-line scatter, a 3σ determination of $n_e = 0.1$ cm $^{-3}$ needs the scatter modelled down by a factor of 7, or of order 10^3 sight lines at the current scatter, which no foreseeable timing programme will deliver. The tractable path is the former: the excess scatter is structure, in the Galactic foreground or in the cluster, and it is measurable in its own right through the DM derivatives that five of these pulsars already have.

Deeper radio. The radio-leg limit falls with rms as $\sigma_S^{1/0.60}$ for a single draw, which would give a factor of 46 for a tenfold improvement. The realized marginal gain is much smaller, because the plane's scatter and the other two legs flatten the response: at SKA-Mid-era depth (~ 0.1 μ Jy) the RIAF ε_{95} tightens by 12.8, 11.8 and 5.9 at the three anchors ($2.9 \times 10^{-10} \rightarrow 2.5 \times 10^{-11}$ at $8,200 M_\odot$), and the exclusion fraction moves from 78 to 94 per cent at the pulsar-timing cap, 85 to 97 per cent at the fast-star anchor and 99 to over 99.9 per cent at $4 \times 10^4 M_\odot$. Drafts to v1.5 quoted the single-draw exponent as though it were the realized one while printing a realized factor of eleven beside it; the two are reconciled here.

Deeper mid-infrared. Under the RIAF family, deeper MIRI photometry changes nothing: the X-ray leg dominates it in every draw, and JWST's role there is the independent-systematics cross-check. Under the jet-dominated family, where the infrared comes closest to mattering, a threefold depth gain moves ε_{95} by 14 per cent at the fast-star anchor (17 per cent at $6,000 M_\odot$, 9 per cent at $4 \times 10^4 M_\odot$). Drafts to v1.5 reported 29, 31 and 21 per cent for the same three numbers, on the uncorrected infrared comparison of Section 2; correcting it roughly halves the forecast gain, and the earlier draft's improvement over the single-constant infrared proxy it replaced is correspondingly smaller. The instrument-priority implication for Paper C is unchanged in order and sharper in degree: radio first, density second, mid-infrared as discrimination between families rather than as depth.

Table 9: The 19 timed pulsars as a gas probe. θ is the angular offset from the cluster centre (Colom i Bernadich et al., 2026), b the projected radius at 5.43 kpc, and ΔDM the excess a uniform core-filling $n_e = 0.1 \text{ cm}^{-3}$ would imprint on that sight line. Pulsars without a quoted DM uncertainty have search-level or partial solutions. Numbers in `figs/ff_expand_dm.json`.

PSR J1326–4728	θ (')	b (pc)	DM (pc cm^{-3})	σ_{DM}	$\Delta\text{DM at } 0.1 \text{ cm}^{-3}$
H	0.56	0.88	98.1716	0.0005	0.73
B	0.76	1.20	100.2806	0.0008	0.71
F	1.00	1.58	98.29	–	0.68
P	1.00	1.58	102.17	–	0.68
O	1.50	2.37	94.305	–	0.58
E	1.58	2.50	94.33967	5×10^{-5}	0.56
J	1.80	2.84	97.28	–	0.49
K	1.89	2.99	94.7836	0.0007	0.45
A	1.93	3.05	100.3267	0.0004	0.43
G	1.96	3.10	99.7445	0.001	0.42
C	1.98	3.13	100.6643	0.0013	0.41
Q	2.30	3.63	95.923	–	0.18
S	2.32	3.66	96.24	–	0.15
M	2.40	3.79	101.47	–	0
D	2.50	3.95	96.5462	0.0009	0
N	2.66	4.20	102.1	–	0
L	3.32	5.24	101.4759	–	0
I	3.53	5.58	102.555	–	0
R	3.90	6.16	102.2	–	0

8 Scope of the adjudication

Epistemic status: Sections 2–7 are instrument-level results and standard accretion arithmetic; nothing in them depends on any hypothesis of this paper set. This section prices the result against the two live readings and is interpretive.

Under the null reading, the center is gas-starved or its flow is suppressed. Drafts to v1.5 could say that this reading paid a price only at the high-mass anchor and that natural-flow parameter space stayed comfortably open below it. On the corrected efficiency law it pays a price everywhere in the contested range: at each of the three anchors it must give up one of its own ingredients, the mass, the fiducial density, or the standard suppression bracket. That is a constraint on conventional astrophysics, publishable and testable on its own terms, and it no longer sharpens the mass tension from one side only, since electromagnetic silence is now hard to buy naturally across the whole range the kinematics argue over. The density is the ingredient this paper’s own Section 4 makes it easiest to give up, and doing so is the cheapest way to restore the null reading. Under the engineered reading of Papers A and E, a system extracting accretion power as work rather than radiation presents a deep multi-band null as its expected signature, at any mass; the data cannot separate that reading from the null, and this paper does not claim otherwise. What it contributes to that adjudication is the currency: a likelihood over ε_B with stated conditioning, which Paper E’s framework ingests in place of the single-band, single-convention numbers it previously had to translate. The duty-cycle loophole of Section 6 is charged to both readings equally.

9 Conclusion

The deepest radio, infrared, and X-ray observations of ω Cen’s center jointly admit a posterior on the Bondi radiative efficiency of any central accretion flow: $\varepsilon_{95} = 2.9 \times 10^{-10}$ at the fast-star mass anchor for a RIAF-like flow under the stated priors, 1.4×10^{-8} with no radio information, and a further factor of 27 tighter at the top of the contested mass range. Scored against a natural outflow-suppressed hot flow at the fiducial density, the data reject 85, 78 and 99 per cent of that flow’s parameter space at the fast-star anchor, the pulsar-timing cap and $4 \times 10^4 M_{\odot}$, and above 0.57, 0.47 and 0.93 across every density prior and convention of Table 6. The curve’s normalization is owned by the unmeasured central density, and the cluster’s own stellar winds now bound that density from above rather than below: a recount of the giants inside the influence radius gives eight of them and a no-removal steady state near 0.03 cm^{-3} , a factor of eight under the 47 Tucanae value the prior is centred on. That recount points the density evidence towards the weaker end of the exclusion range even as the efficiency correction points the verdict towards the stronger, and the current pulsar set can settle it. The single observation that converts these limits from analogy-conditional to measured is a DM-gradient fit to the 19-pulsar set at the ± 10 per cent level (§8.1). Until it exists, Figure 1 is the quantitative meaning of ω Cen’s silence, and the engineered-silence hypothesis of this paper set lives inside it on equal terms with a gas-starvation null that must now give up one of its own ingredients at every anchor rather than only at the highest.

Data availability

The posterior script (`figs/ff_posterior_v4.py`, fixed seed), its full output (`figs/ff_v4_results.json`), its seed study (`figs/ff_v4_seedsread.json`) and realized-scaling export (`figs/ff_v4_scalings.json`), the standalone recomputes behind the corrections of this revision (`figs/ff_r9_checks.py` and its output), the superseded v0.3 script and output retained unchanged as the previous computation of record, the exclusion-fraction tables, the measured-input provenance file that supplies every observational number used above with its source location and verbatim quotation (`figs/ff_measured_inputs.json`) together with the derived epoch and survey-context exports (`figs/ff_epochs.json`, `figs/ff_maveric_context.json`), the superseded v0.2 script and output retained for comparison, the appendix cross-check script (`figs/ff_joint_bound.py`) and its output (`figs/ff_results.json`), the expansion-analysis scripts behind Figures 3, 4 and 5 and Tables 8 and 9 (`figs/ff_expand_v1.py` for the duty-cycle and dispersion-measure analyses, which read no likelihood and are unchanged, and `figs/ff_expand_v2.py` for the three products that do, with outputs `ff_expand_nuisance_v2.json`, `ff_expand_sed_v2.json`, `ff_expand_exclcurve_v2.json`, `ff_expand_duty.json` and `ff_expand_dm.json`), and the figure coordinate tables are distributed with the paper source at omegacentauri.me. The pulsar dispersion measures of Table 9 are read from the timing-set extraction `h/data/pulsars.json`, which records each value’s source location in [Colom i Bernadich et al. \(2026\)](#). The provenance file also records two internal inconsistencies in the sources that a reader checking our inputs will meet: [Mahida et al. \(2026\)](#)’s conservative efficiency limit is stated as 4×10^{-3} in their abstract and Section IV.1 and as 4×10^{-6} in their Section V, and their Table 1 observing hours exceed their own prose totals. Neither enters any number in this paper, and both are flagged here for that paper’s authors. The analysis would improve with, and the machinery directly ingests, the measured central-pixel flux density and its uncertainty from the ATCA image, the per-epoch Chandra source counts and background at the center list, and the per-filter MIRI/NIRCam limiting fluxes; we request them from the respective teams.

A Conventions, unit conversions, and cross-paper reconciliation

Bondi rate. Eq. 1 uses the $\gamma = 5/3$ eigenvalue $\lambda = 0.25$ (Bondi, 1952). A constant copied from the 1952 numerics belongs to the same error species as the ξ correction recorded in Paper H, so we state how it verifies. Bondi’s general polytropic eigenvalue, valid for $1 < \gamma < 5/3$, is

$$\lambda(\gamma) = \frac{1}{4} \left[\frac{2}{5 - 3\gamma} \right]^{(5-3\gamma)/2(\gamma-1)}, \quad (10)$$

which evaluates to 0.500 at $\gamma = 1.5$, 0.625 at $\gamma = 1.4$ and 1.120 as $\gamma \rightarrow 1$, reproducing the tabulated values, while as $\gamma \rightarrow 5/3$ the bracketed factor tends to unity and leaves the prefactor $1/4$ (see also Frank et al., 2002). Drafts to v1.5 printed Eq. 10 with a second factor $[(\gamma + 1)/2]^{-(\gamma+1)/2(\gamma-1)}$, which does not belong to it; that factor is what produced the 0.286 and 0.362 those drafts quoted as the tabulated values, and it sends the $\gamma \rightarrow 5/3$ limit to $9/64 = 0.141$ rather than to $1/4$. The spurious factor is removed here and the values are recomputed in `figs/fF_r9_checks.py`. No number elsewhere in this paper used it: $\dot{M}_B$ was always evaluated at $\lambda = 0.25$. At $\gamma = 5/3$ itself the regular sonic point of the general solution merges with the origin, so the value is read as this limit rather than from a critical-point evaluation. The rate uses a gas-rest-frame denominator c_s^3 , with c_s drawn between the $\mu = 0.6$ isothermal value (11.7 km s^{-1}) and the pure-hydrogen adiabatic value (16.6 km s^{-1}) at 10^4 K . Paper E’s fuel budget (its Eq. 2) uses $\lambda = 1$ with a $(\sigma^2 + c_s^2)^{3/2}$ denominator as a turbulent-medium convention; the net ratio between the two conventions is $0.25 (\sigma^2 + c_s^2)^{3/2} / c_s^3$, which is 1.43 at the median sound speed $c_s = 14.15 \text{ km s}^{-1}$ and reaches 1.9 only at $c_s \approx 12.3 \text{ km s}^{-1}$, the 13th percentile of the prior. Drafts to v1.5 quoted the 1.9 figure as the median value and paired it with $\dot{M}_B = 5.3 \times 10^{18} \text{ g s}^{-1}$ at $2 \times 10^4 M_\odot$; both belong to that same 13th-percentile sound speed. At the median the rate is $\dot{M}_B = 3.5 \times 10^{18} \text{ g s}^{-1}$ against Paper E’s 3×10^{18} , so this paper’s $\dot{M}_B$ is $1.4\times$ Paper E’s at equal (M, n_e) under median nuisances and its ε_B limits are correspondingly $1.4\times$ tighter than under E’s convention. Both statements are correct within their stated conventions; the correction is to the exchange rate, not to either paper’s own numbers.

Electron mean molecular weight. $\mu_e = 2/(1 + X) = 1.17$ at $X = 0.71$. An earlier draft used 1.5, which corresponds to no physical composition and tightened all limits by 27 per cent; corrected here.

Radio. $L_R \equiv \nu L_\nu$ at 5 GHz under a flat spectrum: $L_R = 4\pi d^2 (5 \text{ GHz}) S_\nu$, with S_ν the 7.25 GHz flux density (L_ν flat). At $d = 5.43 \text{ kpc}$, $1.1 \mu\text{Jy}$ corresponds to $L_R = 1.9 \times 10^{26} \text{ erg s}^{-1}$ (5.8×10^{26} at 3σ).

X-ray bands. The fundamental plane is calibrated on 2–10 keV; the Chandra limit is 0.5–7 keV; $F(2\text{--}10)/F(0.5\text{--}7) = \ln(10/2)/\ln(7/0.5) = 0.610$ for $\Gamma = 2$, which is the value used throughout. Haggard et al. (2013) derive their own limit under $\Gamma = 2.3$, and at that index the ratio of power-law integrals is 0.462, lower by 24 per cent. The radio leg enters as $L_R \propto (k f_X L_{\text{bol}})^{0.60}$, so the change moves the predicted radio luminosity by 0.85 and loosens the RIAF anchors by 15 to 20 per cent (Table 10), the largest single convention effect in this appendix. It is carried as a ladder row rather than adopted, because the plane’s own calibration sample is not restricted to $\Gamma = 2.3$ sources and changing k alone would mix the two calibrations.

X-ray reference distance. Haggard et al. (2013) adopt 5.2 kpc (their abstract and Table 1, from the Harris catalogue). An earlier draft of this paper carried 4.8 kpc for that reference and warned that the mismatch against the 5.43 kpc prior shifts the X-ray leg by 28 per cent. Both halves of that warning were wrong. The distance is 5.2 kpc, and the shift is zero regardless: the X-ray likelihood is evaluated on flux, which is what Chandra measured and is distance-independent, so the drawn distance enters the predicted flux and never the limit. The reference distance now serves as a consistency check on the source’s own numbers: $4\pi D^2 f_{X,\text{lim}} = 1.62 \times 10^{30} \text{ erg s}^{-1}$ at 5.2 kpc against the 1.6×10^{30} Haggard et al. (2013) state, which holds at their distance and fails at 4.8 kpc (1.38×10^{30}). The script asserts it at run time. The

infrared limits, by contrast, are published as luminosities and do carry their source distance, so they are referenced at [Chen et al.](#)'s 5.49 kpc and rescaled draw by draw.

Run-time self-checks. Three consistency assertions accompany the shipped code and print their result rather than assuming it, the discipline of Paper H's self-audit. The expansion script (`figs/ff_expand.v2.py`) reproduces the RIAF anchors of Table 4 against `ff_v4_results.json` and asserts the match before proceeding. The Haggard luminosity-anchor check above asserts $4\pi D^2 f_{X,\text{lim}}$ against [Haggard et al.](#)'s stated $1.6 \times 10^{30} \text{ erg s}^{-1}$ at their own distance. And the v0.4 posterior asserts, before computing anything of its own, that with its four corrections disabled it reproduces the v0.3 output file to floating-point equality, so the refactor behind this revision cannot have moved a number by itself; every difference between v1.5 and v2.0 is attributable to a named correction. All three halt the build on failure, so a compiled paper carries their pass by construction.

Confidence conventions. Each leg uses the confidence its source states, mapped through the normal quantile. The X-ray limit is one-sided 95 per cent ($z = 1.645$); the infrared limits are taken at the 95 per cent completeness row on the same convention. Neither choice is load-bearing:

Table 10: Convention ladders for the two limit legs, plus the X-ray band-conversion index. RIAF-family ε_{95} at the three mass anchors; the primary rows are those used throughout the paper.

Leg	Convention	$6,000 M_{\odot}$	$8,200 M_{\odot}$	$4 \times 10^4 M_{\odot}$
X-ray	one-sided 95 % ($z=1.645$, primary)	6.10×10^{-10}	2.94×10^{-10}	1.08×10^{-11}
X-ray	two-sided 95 % ($z=1.960$)	6.01×10^{-10}	2.90×10^{-10}	1.07×10^{-11}
X-ray	earlier draft's $\text{lim}/3$	5.72×10^{-10}	2.77×10^{-10}	1.04×10^{-11}
IR	99.7 % completeness row	6.13×10^{-10}	2.95×10^{-10}	1.08×10^{-11}
IR	95 % completeness row (primary)	6.10×10^{-10}	2.94×10^{-10}	1.08×10^{-11}
IR	68 % completeness row	6.05×10^{-10}	2.92×10^{-10}	1.07×10^{-11}
IR	tightest single filter, 95 %	6.12×10^{-10}	2.95×10^{-10}	1.08×10^{-11}
X-ray band	$\Gamma = 2$ ($k = 0.610$, primary)	6.10×10^{-10}	2.94×10^{-10}	1.08×10^{-11}
X-ray band	$\Gamma = 2.3$ ($k = 0.462$, the source's)	7.34×10^{-10}	3.52×10^{-10}	1.24×10^{-11}

The X-ray ladder spans 3 to 7 per cent on these anchors and the infrared ladder 1 per cent, the latter having narrowed with the band correction of §2 because the infrared terms now sit further from binding. On the jet family, where the infrared leg comes closest to mattering, the tightest-single-filter reading gives 1.39×10^{-9} , 6.57×10^{-10} and 2.08×10^{-11} against the combination's 1.36×10^{-9} , 6.45×10^{-10} and 2.06×10^{-11} , a 2 per cent effect, because F770W dominates the combination in any case. Holding the infrared limits at this paper's 5.43 kpc instead of [Chen et al.](#)'s 5.49 kpc moves the jet anchors by 0.3 per cent.

Reconstruction of the limit legs. A published bound states $P(\text{observed} < f_{\text{lim}})$, which is a survival statement, while Section 3.3 evaluates a Gaussian centred at zero with the same σ . Substituting the survival form on the X-ray and infrared legs, the two published as limits with no central value, moves the radio-free and thin-disk anchors by 1 per cent and leaves every exclusion fraction unchanged to three decimals; those two curves carry only those legs, so that is the size of the choice wherever it applies alone. Drafts to v1.5 put the same figure at 29 to 30 per cent. The difference is the band-integration correction of §2: with the infrared σ_j a factor of four to six larger, the infrared terms sit far enough from their limits that the two likelihood shapes agree, which is the regime in which they are supposed to. The radio leg is excluded from the substitution because [Mahida et al. \(2026\)](#) report a measured central flux density rather than a bound, and a survival form there would discard the measurement; doing so anyway loosens the RIAF anchors by a factor 3.2 to 3.9 and the jet anchors by 2.8 to 3.0, which is the value of the measurement rather

than a property of the reconstruction. An injection study of 200 realizations per cell puts the frequentist coverage of the 95 per cent credible upper limit at 0.92 at $8,200 M_\odot$ and 0.96 at $4 \times 10^4 M_\odot$ when the true ε_B sits at the published limit, against a binomial uncertainty of 0.02. A credible interval carries no coverage guarantee, and the check is run at one point in (ε_B, M) rather than over the parameter space, so it is a calibration spot-check of the reconstruction convention rather than a coverage statement about the method. Numbers in `figs/ff_calcF1_v2_partc.json`, recomputed on the v0.4 likelihood; the v0.3 values are retained in `figs/ff_calcF1_hybrid.json` and `figs/ff_calcF1_coverage.json`.

Cross-paper exchange rates. At the common reference point ($M = 2 \times 10^4 M_\odot$, $n_e = 0.23 \text{ cm}^{-3}$, median nuisances): this paper’s ε_B relates to the horizon-side efficiency as $\eta = \varepsilon_B/f_B$ with $f_B \in [10^{-4.4}, 10^{-2.6}]$, the bracket of Section 3.1 under $r_g = GM/c^2$ (drafts to v1.5 printed $[10^{-4.1}, 10^{-2.5}]$ here, which follows from no (c_s, s) corner under either r_g convention); Mahida et al. (2026)’s bounded quantity is an accretion fraction under their assumed radiative model and is neither ε_B nor η ; Haggard et al. (2013)’s efficiency is an ε_B -like quantity evaluated at their own inputs, and those inputs are now stated here rather than left implicit: $M = 1.8 \times 10^4 M_\odot$, $n_e = 0.038 \text{ cm}^{-3}$ from a Pfahl–Rappaport wind budget, and $T = 10^4 \text{ K}$, giving the $\eta = 4 \times 10^{-9}$ to 1.3×10^{-8} they compute for a simple Bondi rate. Because $\varepsilon_B \propto 1/(n_e \lambda)$ at fixed luminosity, the exchange factor into this paper’s convention is $(0.038 \lambda_H)/(0.23 \times 0.25)$: at $\lambda_H = 0.25$ it is 0.17, and it reaches 0.66 only if their rate is read with $\lambda_H = 1$. Drafts to v1.5 printed a bracket of 0.5–1.5, which no reading of the two densities supports; the factor is 0.17–0.66 depending on the λ convention attributed to that paper, and the residual ambiguity is theirs to resolve. The density difference is the substance of it: their 0.038 cm^{-3} is a factor 6.1 below this paper’s prior median, and §4 now reaches a comparable figure from an independent route. The same caution applies to Mahida et al.’s accretion fraction, whose model normalization is not stated in convertible form. Until both are re-evaluated in one convention, the qualitative ordering (radio strongest, X-ray next, IR family-dependent) is what transfers.

A.1 Shared fiducial parameters across the paper set

This table is a cross-reference index, not a mandate: it records the value each paper in the set actually uses for seven recurring cluster parameters, and where a paper’s choice differs from the majority it names the deviation rather than silently harmonizing it. No paper’s own derivation is changed by this table; where two papers differ, both retain their own number, and the table exists so a reader moving between papers can see the exchange rate rather than assume equality (R9-XPAPER-1 X-R9-08/09/10/11, folding R9-SSOT-1 and the R8 leftover X-R8-13 on core-radius conventions).

σ fiducial disclosure. This paper’s likelihood does not use σ directly, but the flyby and dynamical calculations elsewhere in the set do, at the 21 km s^{-1} paper-set fiducial rather than the 18.2 km s^{-1} the public MCP calculators default to; the two are not harmonized, and the $(21/18.2)^2 \approx 1.3$ factor in r_{infl} is the size of the resulting offset (X-R9-12, the same disclosure Paper E already carries at its own flyby-envelope figures).

B Threshold-inversion cross-check

An earlier construction of this bound (script `figs/ff_joint_bound.py`, retained unchanged with its own seed and output) inverted the per-band thresholds as that draft labelled them, 3σ in every band, through the same nuisance draws and reported percentiles of the per-draw minimum. It is transparent and reproducible, and its 95th-percentile “conservative” values (1.5×10^{-8} at $8,200 M_\odot$ with the radio leg; 3.0×10^{-7} without) sit one to two decades above the credible limits of Table 4, for two reasons with

Quantity	Value used, by paper (deviation noted where it exists)
Distance d	A/C/E/G/H: 5.49 kpc (Häberle et al., 2025) (H uses 5.494 kpc to three decimals; G’s angular catalogue does not carry a distance). B/D follow the same 5.49 kpc set fiducial. F: 5.43 kpc (Baumgardt and Vasiliev, 2021), disclosed as this paper’s own likelihood input and a 1.2σ offset from the set fiducial (§3.2), not harmonized away.
Velocity dispersion σ	Paper-set fiducial $\sigma = 21 \text{ km s}^{-1}$ (e.g. Paper E’s flyby Monte Carlo, its Appendix A). The public MCP calculators default to 18.2 km s^{-1} instead; $(21/18.2)^2 \approx 1.33$ in $r_{\text{infl}} \propto \sigma^{-2}$, so a reader moving between a paper figure and a calculator output should expect a ~ 33 per cent offset in influence radius from this choice alone (§A.1 below; the same disclosure Paper E already carries).
Core radius r_c	Not one quantity: three different conventions appear under the same symbol and are not interchangeable. A/F: $2.37'$, the Harris (1996)/King catalogue convention. G: $155'' (= 2.58')$, its own catalogue core radius (Anderson & van der Marel 2010), used as the reference scale for source offsets and cluster-membership weighting. E: 3.6 pc physical ($\approx 2.25'$ at 5.49 kpc), the adopted flat-core radius of its enclosed-mass model. The three differ by up to ~ 15 per cent in angular terms and are stated here with their conventions attached rather than forced to one number.
Cluster mass M_{cl}	A: $\approx 4 \times 10^6 M_{\odot}$ (Harris, 1996; Baumgardt and Vasiliev, 2021), its adopted catalogue total (its own parameters table). E: $3.55 \times 10^6 M_{\odot}$, the Plummer-profile total mass used in its enclosed-mass model. H’s own visible-mass fit gives $M_{\text{vis}} = 3.54 \times 10^6 M_{\odot}$, consistent with E’s figure to < 1 per cent and not with A’s catalogue total, which is not itself a dynamical-model input in A.
Central density ρ_c	A: catalogue value $\sim 3 \times 10^3 M_{\odot} \text{ pc}^{-3}$ (Baumgardt and Vasiliev, 2021), alongside its own adopted King-model value $4.0 \times 10^3 M_{\odot} \text{ pc}^{-3}$ (34 per cent above the catalogue figure; both are disclosed side by side in its own parameters table). C/E: adopted flat-core value $3 \times 10^3 M_{\odot} \text{ pc}^{-3}$, matching A’s catalogue row rather than its King-model row.
Mean/effective stellar mass $\langle m \rangle$	A: $0.43 M_{\odot}$ per star, the conversion used for its stellar-count enclosed-mass table. E: $0.425 M_{\odot}$, the mean of its two-component baseline mass function with no remnants (its flyby-Monte-Carlo appendix); the two differ by < 2 per cent and are not reconciled further here.
Electron density n_e	F’s own wind-recount work (§4): prior median 0.23 cm^{-3} (log-normal, 0.5 dex, 47 Tuc analogy; Freire et al. 2001; Abbate et al. 2018), independently bounded from above by F’s stellar-wind budget at $\sim 0.030 \text{ cm}^{-3}$ (a ceiling, not a floor, after the R9 wind-count correction of §4); no other paper in the set adopts an independent n_e value.

Table 11: Shared fiducial parameters, by paper, as each paper’s own text currently states them. Not a harmonization: differences are cross-referenced, not resolved, except where a paper’s own revision (outside this WU’s scope) already brought it into line.

opposite signs: the threshold construction discards the information in the measured values (weakening), and its percentile is a nuisance-prior quantile of a limit rather than a posterior statement about ε_B (not comparable in coverage). The disagreement is largest where the fundamental-plane tail dominates, as expected. Its X-ray threshold also predates the confidence relabelling of Section 2, which is a further reason the two constructions are not expected to agree numerically. The posterior supersedes it; it remains in the distribution as a check that no single modeling choice in Section 3.3 drives the headline numbers by itself.

References

- Abbate, F., Possenti, A., Ridolfi, A., Freire, P. C. C., Camilo, F., Manchester, R. N., and D’Amico, N. (2018). Internal gas models and central black hole in 47 Tucanae using millisecond pulsars. *Monthly Notices of the Royal Astronomical Society*, 481(1):627–637.
- Bañares-Hernández, A., Calore, F., Martín Camalich, J., and Read, J. I. (2025). New constraints on the central mass contents of Omega Centauri from combined stellar kinematics and pulsar timing. *Astronomy & Astrophysics*, 693:A104.
- Baumgardt, H. and Vasiliev, E. (2021). Accurate distances to Galactic globular clusters through a combination of Gaia EDR3, HST, and literature data. *Monthly Notices of the Royal Astronomical Society*, 505:5957–5977.
- Blandford, R. D. and Begelman, M. C. (1999). On the fate of gas accreting at a low rate on to a black hole. *Monthly Notices of the Royal Astronomical Society*, 303:L1–L5.
- Bondi, H. (1952). On spherically symmetrical accretion. *Monthly Notices of the Royal Astronomical Society*, 112(2):195–204.
- Chatterjee, P., Hernquist, L., and Loeb, A. (2002). Dynamics of a massive black hole at the center of a dense stellar system. *The Astrophysical Journal*, 572:371–381.
- Chen, S., Hare, J., Kargaltsev, O., Yang, H., Cioffi, D., Häberle, M., and Seth, A. (2026). The intermediate mass black hole in Omega Centauri: Constraints on accretion from JWST. *The Astrophysical Journal*.
- Colomo Bernadich, M., Dai, S., Abbate, F., Kerr, M., Bachetti, M., Bhargava, Y., Buchner, S., Johnston, S., Burgay, M., Possenti, A., Nag, R., Ridolfi, A., Carleo, A., Corongiu, A., Freire, P. C. C., Camilo, F., Chen, W., Cadelano, M., Risbud, D., Padmanabh, P. V., Champion, D. J., Kramer, M., Stappers, B., Serylak, M., Balakrishnan, V., Bailes, M., Dutta, A., Vleeschower Calas, L., Venkatraman Krishnan, V., and Men, Y. (2026). A joint MeerKAT and Parkes view of Omega Centauri: New TRAPUM searches and pulsar timing. *arXiv e-prints*, page arXiv:2603.21845.
- Di Cintio, P., Pasquato, M., Barbieri, L., Trani, A. A., and Di Carlo, U. N. (2023). Dynamics of intermediate mass black holes in globular clusters. Wander radius and anisotropy profiles. *Astronomy & Astrophysics*, 673:A8.
- Frank, J., King, A., and Raine, D. (2002). *Accretion Power in Astrophysics*. Cambridge University Press, Cambridge, 3rd edition.
- Freire, P. C., Kramer, M., Lyne, A. G., Camilo, F., Manchester, R. N., and D’Amico, N. (2001). Detection of ionized gas in the globular cluster 47 Tucanae. *The Astrophysical Journal Letters*, 557:L105–L108.

- Häberle, M., Neumayer, N., Clontz, C., Seth, A. C., Bellini, A., et al. (2025). oMEGACat. VI. Analysis of the overall kinematics of Omega Centauri in 3D: Velocity dispersion, kinematic distance, anisotropy, and energy equipartition. *The Astrophysical Journal*, 983:95.
- Häberle, M., Neumayer, N., Seth, A., Bellini, A., Libralato, M., Baumgardt, H., Whitaker, M., Dumont, A., Alfaro-Cuello, M., Anderson, J., Clontz, C., et al. (2024). Fast-moving stars around an intermediate-mass black hole in ω Centauri. *Nature*, 631(8020):285–288.
- Haggard, D., Cool, A. M., Heinke, C. O., van der Marel, R., Cohn, H. N., Lugger, P. M., and Anderson, J. (2013). A deep Chandra X-ray limit on the putative IMBH in Omega Centauri. *The Astrophysical Journal Letters*, 773:L31.
- Harris, W. E. (1996). A catalog of parameters for globular clusters in the Milky Way. *The Astronomical Journal*, 112:1487.
- Lu, T.-N. and Kong, A. K. H. (2011). Radio continuum observations of 47 Tucanae and ω Centauri: Hints for intermediate-mass black holes? *The Astrophysical Journal Letters*, 729:L25.
- Maccarone, T. J., Fender, R. P., and Tzioumis, A. K. (2005). Upper limits on central black hole masses of globular clusters from radio emission and a possible black hole detection in the Ursa Minor dwarf galaxy. *Monthly Notices of the Royal Astronomical Society*, 356:L17–L22.
- Mahida, A. D. et al. (2026). No evidence for accretion around the intermediate-mass black hole in Omega Centauri. *The Astrophysical Journal*, 996:122.
- Merloni, A., Heinz, S., and Di Matteo, T. (2003). A fundamental plane of black hole activity. *Monthly Notices of the Royal Astronomical Society*, 345:1057–1076.
- Noyola, E., Gebhardt, K., and Bergmann, M. (2008). Gemini and Hubble Space Telescope evidence for an intermediate-mass black hole in ω Centauri. *The Astrophysical Journal*, 676:1008–1015.
- Noyola, E., Gebhardt, K., Kissler-Patig, M., Lützgendorf, N., Jalali, B., de Zeeuw, P. T., and Baumgardt, H. (2010). Very large telescope kinematics for omega Centauri: Further support for a central black hole. *The Astrophysical Journal Letters*, 719:L60–L64.
- Plotkin, R. M., Markoff, S., Kelly, B. C., Körding, E., and Anderson, S. F. (2012). Using the fundamental plane of black hole activity to distinguish X-ray processes from weakly accreting black holes. *Monthly Notices of the Royal Astronomical Society*, 419:267–286.
- Strader, J., Chomiuk, L., Maccarone, T. J., Miller-Jones, J. C. A., Seth, A. C., Heinke, C. O., and Sivakoff, G. R. (2012). No evidence for intermediate-mass black holes in globular clusters: Strong constraints from the JVL. *The Astrophysical Journal Letters*, 750:L27.
- Tremou, E., Strader, J., Chomiuk, L., Shishkovsky, L., Maccarone, T. J., Miller-Jones, J. C. A., et al. (2018). The MAVERIC survey: Still no evidence for accreting intermediate-mass black holes in globular clusters. *The Astrophysical Journal*, 862:16.
- van der Marel, R. P. and Anderson, J. (2010). New limits on an intermediate-mass black hole in Omega Centauri. II. Dynamical models. *The Astrophysical Journal*, 710:1063–1088.

- Wang, Z., Seth, A. C., Latour, M., Strader, J., Häberle, M., et al. (2025). oMEGACat. VII. tracing interstellar and intracluster medium of ω Centauri using sodium absorptions. *The Astrophysical Journal*, 994:143.
- Xie, F.-G. and Yuan, F. (2012). Radiative efficiency of hot accretion flows. *Monthly Notices of the Royal Astronomical Society*, 427(2):1580–1586.
- Yuan, F. and Narayan, R. (2014). Hot accretion flows around black holes. *Annual Review of Astronomy and Astrophysics*, 52:529–588.