Engineered Intermediate-Mass Black Hole Systems: Infrastructure Constraints, Observable Residue, and a Multi-Messenger Adjudication Framework

Tim Swanson

The Omega Centauri Society / Post Oak Labs

tim@postoaklabs.com

July 2026

Draft v0.93 (last revised 2026-07-30) — fifth paper of the set; companion to *The Macro Transcension Hypothesis* (Paper A),
the inward-migration review (Paper B), the ω Cen campaign (Paper C), and the migration economics (Paper D); prepared for omegacentauri.me

Abstract

The inward-migration resolutions of the Fermi paradox identify rapidly spinning intermediate-mass black holes (IMBHs) in dense, old stellar clusters as thermodynamically privileged destinations for computation-optimizing civilizations. Previous papers in this series argued the thermodynamic case (A), surveyed the hypothesis family (B), designed a multi-messenger campaign for the nearest candidate, Omega Centauri (C), and priced the migration decision (D). This paper addresses the two remaining questions. First, feasibility: whether large-scale computational infrastructure can persist around a Kerr IMBH embedded in a live cluster core (stellar density $\sim 3 \times 10^3\, M_\odot\, \mathrm{pc}^{-3}$, velocity dispersion $\sim 21\,\mathrm{km\,s^{-1}}$). Extending recent passive-stability results for stellar engines and Dyson bubbles (McInnes, 2026) to the combined Kerr-plus-cluster potential, and combining analytic tidal, thermal, and material limits with a Monte Carlo of gravitationally focused stellar flybys, we derive an allowed envelope for a fiducial $2 \times 10^4\, M_\odot$ hole in the ω Cen core: precession-tolerant swarms survive passively from $\sim 10^2$ gravitational radii out to the cluster stripping radius at $\sim 4 \times 10^3$ AU. Stellar flybys never set the boundary, and the margin is larger than an impulsive treatment shows. The Monte Carlo, run with a mass-segregated heavy-remnant perturber component ($31\, M_\odot$ black holes at 0.1–3 per cent number fraction), yields a conservative impulsive floor of $\sim 6 \times 10^7$ yr at the fiducial 1 per cent fraction (3×10^9 yr without remnants); but unbound cluster stars are adiabatically decoupled from every orbit in the envelope (the adiabatic parameter runs from 2.2 at the stripping radius to 4×10^3 at the ISCO), and applying the standard correction lifts the physical lifetime to $\gtrsim 3 \times 10^8$ yr at the envelope edge and to the cluster age within ~ 3 AU. The relaxed Bahcall–Wolf cusp around the hole is no longer a conditional: the González Prieto et al. (2025) realizations form one at 300 Myr with $\gamma \simeq 1.75$, steepening to $\gamma \simeq 2$ for the heavy remnants. Cusp members are bound, so the impulsive kernel applies to them where it does not to the unbound field, and the population is granular, roughly one black hole expected inside the stripping radius: a bound heavy remnant is present in about half of realizations at the envelope edge, giving a diffusion floor of a few times 10^7 yr, against the cluster-age cap when none is drawn. That bound-member diffusion channel, not the decoupled unbound field, is the binding clock at the envelope edge. The same remnant also forces the orbit secularly (Kozai–Lidov), but that forcing is coherent and quenched by relativistic apsidal precession inward of $\sim 4 \times 10^2$ AU, so it reads as a station-keeping cadence rather than a second survival limit. Tidal disruption events set the hazard-recurrence horizon ($\gtrsim 10^7$ yr), and the measured intracluster medium funds, via Bondi accretion at magnetically arrested efficiencies, a power budget of up to $\sim 8 \times 10^5 \, L_\odot$, provided the standard outflow suppression of hot low-Eddington flows is itself suppressed, an engineered capability; the natural suppressed supply is $\sim 10^3 \, L_\odot$, at or below the thermal-concealment ceiling. For the engineered case supply never binds; thermal concealment does. An abandoned deep swarm grinds to debris on an estimated $\sim 10^3$ yr timescale (Lacki, 2025) and is then drained by the hole, leaving spin as the only durable fossil of engineered history.

Second, residue and adjudication: the envelope implies three forward-modeled observables, a temperature-dependent waste-heat floor, a magnetically-arrested-disk (MAD) regulation signature, suppression of the flux-eruption variability characteristic of natural MAD accretion (Ripperda et al., 2022; Yang et al., 2026), and environmental dephasing of any extreme-mass-ratio inspiral, the only channel that constrains engineered mass (Barausse et al., 2014). We construct a hierarchical Bayesian framework, with per-messenger Bayes factors against an explicit menu of astrophysical nulls (quiescent IMBH; stellar-remnant subcluster) combined through coincidence likelihoods of the kind developed for gravitational-wave counterpart searches (Ashton et al., 2018; Breschi et al., 2024), and with pre-registered decision thresholds. Applied to the current ω Cen data, the framework yields a Bayes factor mildly favoring the astrophysical nulls ($\ln K = -0.47$ from the mid-infrared channel, a continuous per-filter likelihood against the real JWST/MIRI sensitivity curve, marginalized over swarm radius and temperature split); the value is a demonstration of the machinery, dominated by the dormancy prior rather than by the data, and moves toward -0.35 under an outflow-truncated fuel ceiling. That channel is one-sided and its evidence is capped at $\ln f_d = -0.69$ for any depth of photometry, of which 68 per cent is already spent, so the framework's present product is the information forecast: how much each planned observation from Paper C can move the odds, and how much headroom each channel has left. The adjudication machinery itself is validated, not only proposed: scored against three historical cases with an independently known resolution, including a positive control (the 1967 pulsar discovery, LGM-1), it registers a strong anomaly ($\ln K = +7.9$) ahead of the resolving hypothesis and collapses appropriately once that hypothesis is added to the menu, while returning no false positive on either mundane case tested (Appendix E). All results are conditional on the optimization premise shared by the hypothesis family; the contribution is the feasibility envelope, the forward-modeled residue, and the validated adjudication machinery.

Keywords: Fermi paradox; SETI; technosignatures; intermediate-mass black holes; Omega Centauri; megastructures; magnetically arrested disks; Bayesian model selection; multi-messenger astronomy

Contents

1 Introduction

1.1 The missing quadrant

Four papers in this series have examined the hypothesis that computation-optimizing civilizations migrate toward rapidly spinning massive black holes in dense old stellar systems. Paper A (Swanson, 2026c) argued why: the thermodynamic advantages of Kerr accretion efficiency, horizon entropy disposal, and Bekenstein-bounded information storage. Paper B (Swanson, 2026b) surveyed who: the six-member family of inward-migration proposals and their falsifiability grades. Paper C (Swanson, 2026d) planned how to look: eight instrument-matched programs targeting Omega Centauri (NGC 5139), the nearest strong IMBH candidate. Paper D (Swanson, 2026a) priced whether to go: closed-form crossover conditions between migration, densification, and seeding strategies.

Two questions remain unaddressed, in this series and, so far as we can determine, in the literature. First: could the infrastructure such a civilization requires actually persist at the destination? The destination is a demanding one. A cluster-core IMBH sits in a stellar environment three to four orders of magnitude denser than the solar neighborhood, threaded by gravitationally focused stellar traffic, and (if fueled) surrounded by a relativistic accretion flow. Black-hole computing proposals (Inoue and Yokoo, 2011; Opatrný et al., 2017; Dvali and Osmanov, 2023) treat the hole as an idealized resource and the environment as empty. Megastructure stability analyses (Wright, 2020; Raval and Srikanth, 2024; McInnes, 2026) treat single stars in isolation. The intersection, engineered structures bound to a Kerr hole inside a live cluster core, is unexamined.

Second: if the campaign of Paper C ever records an anomaly, by what procedure would the community decide what it had seen? Radio SETI has mature per-signal verification practice (Sheikh et al., 2021; Li et al., 2026) and Bayesian population inference (Grimaldi and Marcy, 2018), and multi-messenger astrophysics has quantitative coincidence formalisms (Ashton et al., 2018; Breschi et al., 2024), but no published framework adjudicates a candidate technosignature across multiple messengers against an explicit menu of astrophysical alternatives. Paper C's rule that "every anomaly is adjudicated by at least two independent messengers" was stated as policy; here we give it mathematical content.

The two questions are coupled, and that coupling is the reason they share a paper. A feasibility envelope is a prior: it tells the adjudicator where in parameter space an engineered system could sit, and therefore which anomalies deserve elevated scrutiny and which are excluded on engineering grounds regardless of how anomalous they appear. Forward-modeled residue channels are likelihoods: they specify what an engineered system would look like, channel by channel, so that Bayes factors can be computed rather than gestured at. The paper thus runs in one direction: constraints (§2) produce observables (§3) which feed an adjudication engine (§4) that we exercise on real data (§5).

1.2 Conditionality and register

Epistemic status: Everything downstream of §2 is conditional on the optimization premise examined critically in Papers A and B: that some long-lived technological lineages maximize discounted computation and act on the thermodynamic gradients described there. This paper does not defend that premise; it asks what follows from it observationally, and how a null or a detection would be established. The adjudication framework of §4 is independent of the premise and applicable to any technosignature search.

We work throughout with the fiducial system of Papers A and C: a black hole of mass $M = 2 \times 10^4\, M_\odot$ (bracketed by the $8.2 \times 10^3\, M_\odot$ kinematic lower bound of Häberle et al. 2024b and the $\sim 5 \times 10^4\, M_\odot$ N-body preferred value of González Prieto et al. 2025), spin left free, embedded in the ω Cen core: central density $\rho_0 \simeq 3 \times 10^3\, M_\odot\,\mathrm{pc}^{-3}$, core radius $r_c \simeq 3.6$ pc, line-of-sight velocity dispersion $\sigma \simeq 20$–$23\,\mathrm{km\,s}^{-1}$

near the center (Baumgardt and Vasiliev, 2021; Nitschai et al., 2023). Where a single value is needed we adopt $\sigma = 21\,\mathrm{km\,s^{-1}}$, the midpoint of that central range and consistent with the oMEGACat 3D kinematic analysis (Häberle et al., 2025), and use it consistently in the influence radius of §2.3, the flyby velocities of §2.4, and Appendix A. The series' public calculators default instead to the cluster-averaged $18.2\,\mathrm{km\,s^{-1}}$ of the shared measurement compilation, which is a global rather than central-region figure; the difference propagates as a factor $(21/18.2)^2 \simeq 1.3$ in r_{infl} and is within the parameter variations of Appendix A.

The gravitational radius is

$$r_{\mathrm{g}} \equiv \frac{GM}{c^2} \simeq 3.0 \times 10^9\,\mathrm{cm}\left(\frac{M}{2 \times 10^4\,M_\odot}\right) \simeq 2.0 \times 10^{-4}\,\mathrm{AU}\left(\frac{M}{2 \times 10^4\,M_\odot}\right), \tag{1}$$

so the dynamic range between horizon scale and cluster scale is nine orders of magnitude in radius. The feasibility question is where in that range infrastructure can live.

2 The feasibility envelope

2.1 Constraint inventory

We consider a one-parameter family of architectures indexed by compactness: at one end, a dense swarm of independent elements on near-circular orbits at $r \sim 10^2$–$10^4\,r_{\mathrm{g}}$ (the configuration favored by the thermodynamic argument of Paper A, since proximity to the horizon minimizes entropy-transport losses); at the other, an extended bubble or shell of elements at $r \sim 10$–10^3 AU supported partly by radiation pressure or tether stress, the black-hole analogue of the Dyson bubbles whose collective stability was recently established by McInnes (2026). Intermediate cases (rings, nested tori) inherit constraints from both ends (Raval and Srikanth, 2024). We deliberately do not commit to an architecture; the envelope is the set of $(r, \text{architecture})$ pairs surviving all constraints simultaneously.

Six constraints bound the envelope: (i) relativistic orbit stability (inner); (ii) tidal stress on extended elements (inner); (iii) accretion-flow radiation and magnetic environment, if the hole is fueled (inner); (iv) cluster tidal truncation of bound orbits (outer); (v) stellar flyby perturbation and collision hazard (outer); (vi) thermal rejection capacity, which couples to the residue analysis of §3.

2.2 Inner boundary

Orbit stability. Circular equatorial orbits around a Kerr hole are stable outside the innermost stable circular orbit, $r_{\mathrm{ISCO}} = 6\,r_{\mathrm{g}}$ (Schwarzschild) shrinking to r_{g} (prograde, extremal) (Bardeen et al., 1972). For $M = 2 \times 10^4\,M_\odot$ this is $\sim 10^{-3}$ AU: dynamically, matter can orbit extraordinarily deep. Orbital periods there are $P \simeq 2\pi(r^3/GM)^{1/2} \sim$ minutes, and communication latency across the swarm is milliseconds, which is part of the computational attraction (Paper A, §4).

Tidal stress. A rigid element of size ℓ at radius r experiences differential acceleration $\sim 2GM\ell/r^3$. Requiring internal stress below a material strength S for an element of density ρ_{el} gives

$$\ell \lesssim \left(\frac{S\,r^3}{2GM\rho_{\mathrm{el}}}\right)^{1/2} \simeq 2 \times 10^2\,\mathrm{km}\left(\frac{S}{10\,\mathrm{GPa}}\right)^{1/2}\left(\frac{\rho_{\mathrm{el}}}{10^3\,\mathrm{kg\,m^{-3}}}\right)^{-1/2}\left(\frac{r}{10^2\,r_{\mathrm{g}}}\right)^{3/2}\left(\frac{M}{2 \times 10^4\,M_\odot}\right)^{-1}, \tag{2}$$

where 10 GPa is the tensile strength of present-day carbon composites. Tides therefore never forbid the swarm architecture (elements of km scale are unconstrained down to $\sim 10\,r_{\mathrm{g}}$); they forbid monolithic structures below $\sim 10^3\,r_{\mathrm{g}}$ and thereby select swarms at small radii, a conclusion that parallels the single-star case (Wright, 2020).

Radiation and magnetic environment. If the hole is fueled at the level required for the Blandford–Znajek power budget of Paper A, the inner $\sim 10^2\, r_g$ contains a magnetically arrested flow with field strength $B \sim 10^2$–10^4 G near the horizon (Tchekhovskoy et al., 2011; Narayan et al., 2022) and a jet along the spin axis. Equatorial swarm orbits outside the disk body ($r \gtrsim 10^2\, r_g$, inclined to avoid both disk and jet) survive; structures inside $\sim 30\, r_g$ face erosion by the flow itself. We adopt $r_{\rm in} \simeq 10^2\, r_g$ as the practical inner boundary of a fueled configuration, and $r_{\rm ISCO}$ for a dormant one.

Formation coherence under relativistic precession. A constraint with no single-star analogue: swarm elements at different radii and inclinations precess differentially. Pericenter advance and Lense–Thirring nodal precession scale as (r_g/r) and $a_\star (r_g/r)^{3/2}$ per orbit. At $r = 10^2\, r_g$ the orbital period is $2\pi(r^3/GM)^{1/2} \simeq 634$ s $\simeq 11$ min, and a formation with a 2 per cent radial spread accumulates order-unity differential phase in $\sim 2 \times 10^2$ orbits, about 1.5 days. Because this exclusion recurs throughout the residue analysis, we state it as the paper's first numbered result:

> **Result 1 (precession decoherence).** For $r \lesssim 10^3\, r_g$, differential apsidal and Lense–Thirring precession destroys any rigid shaped formation on day-to-week timescales; the deep envelope admits only axisymmetric shells, which are precession-invariant by symmetry, or phase-agnostic swarms that route function through communication rather than geometry. (3)

Rigid geometric formations (the flat rings and discs whose stability McInnes (2026) and Raval and Srikanth (2024) analyze around stars) are therefore excluded deep in the envelope. This is the Kerr-specific amendment to the single-star stability literature: the deep envelope selects not only swarm over monolith (tides, above) but symmetric or phase-agnostic swarms over shaped formations.

2.3 Outer boundary

Cluster truncation. The hole dominates the cluster potential inside its influence radius

$$r_{\rm infl} = \frac{GM}{\sigma^2} \simeq 0.2\,{\rm pc} \left(\frac{M}{2\times 10^4\, M_\odot}\right) \left(\frac{\sigma}{21\,{\rm km\,s^{-1}}}\right)^{-2} \simeq 4\times 10^4\ {\rm AU}. \tag{4}$$

Orbits bound to the hole are progressively stripped by the fluctuating cluster field as $r \to r_{\rm infl}$; by analogy with tidal truncation of binaries in clusters (Binney and Tremaine, 2008; Heggie and Hut, 2003), orbits are secure only for $r \lesssim 0.1\, r_{\rm infl} \sim 4 \times 10^3$ AU.

Stellar flybys. The operative outer constraint is sharper than truncation: stellar traffic. The rate at which cluster stars pass within impact parameter b of the hole, including gravitational focusing, is

$$\Gamma(b) = n_\star\, \pi b^2\, v_{\rm rel} \left(1 + \frac{2GM}{b\, v_{\rm rel}^2}\right), \tag{5}$$

with $n_\star \simeq 10^4$ pc^{-3} (for mean stellar mass $0.3\, M_\odot$ at the quoted mass density) and $v_{\rm rel} \simeq \sqrt{2}\,\sigma \simeq 30\,{\rm km\,s^{-1}}$, the velocity used in the focusing term throughout. At $b = 10^3$ AU the focusing term is ~ 40 (it would be ~ 80 if evaluated at $v = \sigma$; we use $v_{\rm rel}$) and Eq. 5 gives one passage per $\sim 10^3$ yr; at $b = 10^2$ AU, one per $\sim 10^4$ yr; at $b = 10$ AU, one per $\sim 10^5$ yr (focusing-dominated regime, $\Gamma \propto b$). Each passage tidally perturbs swarm orbits at the impulsive level $\Delta v/v_{\rm orb} \sim 2(m_\star/M)\,(a/b)^2\,(v_{\rm orb}/v_{\rm rel})$ for structure semi-major axis $a < b$, saturating at $\sim 2(m_\star/M)(v_{\rm orb}/v_{\rm rel})$ for penetrating passages, where $v_{\rm orb} = (GM/a)^{1/2}$. The mass ratio controls everything: at $m_\star/M \sim 10^{-5}$ even a penetrating passage delivers $\Delta v/v_{\rm orb} \lesssim 10^{-2}$, so no single flyby disrupts a hole-bound orbit, and the hazard is cumulative, a random walk in eccentricity. Two further protections apply. Direct star–element scattering requires approach within $\sim 2Gm_\star/(v_{\rm rel} v_{\rm orb})$, a fraction $\lesssim 10^{-7}$ of the orbital cross-section per penetrating passage: negligible. And deep in the envelope the encounter duration $b/v_{\rm rel}$ exceeds the orbital period by orders of magnitude, so the impulsive estimate above is an overestimate: adiabatic invariance suppresses the coupling of slow perturbations to fast orbits exponentially (Binney and Tremaine, 2008; Heggie and Hut, 2003).

Secular cluster-tide forcing. Flybys are the granular part of the cluster field; the smooth part forces eccentricity secularly, the cluster analogue of Lidov–Kozai oscillations analyzed for binaries in cluster tides by Hamilton and Rafikov (2019). The secular timescale for a hole-bound orbit of semi-major axis a is of order $t_{\rm sec} \sim P_{\rm orb}\, M/M_{\rm cl}(a)$, where $M_{\rm cl}(a) = (4\pi/3)\rho_0 a^3$ is the smooth cluster mass enclosed by the orbit. Because the hole outweighs the enclosed cluster mass by $\gtrsim 10^5$ everywhere in the envelope, the forcing is weak: at $a = 10^3$ AU, $M_{\rm cl} \sim 10^{-3}\, M_\odot$ against $2\times10^4\, M_\odot$ and $t_{\rm sec} \sim 3\times10^9$ yr, falling as $a^{-3/2}$ to $\sim 4\times10^8$ yr at the 4×10^3 AU stripping radius. The estimate assumes the tide's full quadrupole is available; ω Cen's measured near-spherical shape (global ellipticity ~ 0.1, rounder still in the core) suppresses the non-axisymmetric component that drives the largest eccentricity excursions, so these are upper limits on the forcing rate. Secular tides therefore do not bind interior to the stripping radius: they become comparable to the flyby-diffusion floor of §2.4 only at the envelope's outermost edge, where stripping already terminates it.

2.4 Monte Carlo of flyby histories

We quantify the cumulative hazard with a Monte Carlo over encounter statistics (Figure 1; method and parameters in Appendix A, code in the paper repository): impact parameters from the focused distribution of Eq. 5 within $b < 30a$, relative speeds drawn as described in Appendix A, impulsive kicks with the saturated form above, and dynamical survival defined as eccentricity random-walking to orbit-crossing ($e \sim 0.5$). The perturber mass function has two stellar components (main sequence plus white-dwarf tail; Häberle et al. 2024a) and, because kick variance scales as $m_\star^2$ and mass segregation concentrates heavy remnants inside the influence radius, a segregated remnant extension: neutron stars ($1.4\, M_\odot$, 2 per cent number fraction) and stellar-mass black holes at a fiducial local number fraction of 1 per cent, bracketed by 0.1 and 3 per cent. We take the black-hole mass to be $31\, M_\odot$, the mean mass of holes inspiraling into the IMBH in the González Prieto et al. (2025) ω Cen models, which identify $10\, M_\odot$ as the assumption they are correcting; the metal-poor progenitors of a cluster at $[\mathrm{Fe/H}] \sim -1.5$ leave heavier remnants than the solar-metallicity default. Two caveats attach to this unbound-field comparison bracket, which retains the stipulated 0.1–3 per cent fraction: the inspiraling set is segregation-biased toward the heavy end, so the population mean is lower than $31\, M_\odot$; and a bottom-heavy initial mass function of the kind those models require ($\alpha_3 = 2.5$) produces fewer black holes. The population that actually operates, the bound cusp channel below, is not assembled from separate figures; it is anchored jointly on the extended dark mass and the mean remnant mass directly (§2.4, Appendix A). At $31\, M_\odot$ per object, the 2–$3\times10^5\, M_\odot$ extended dark component favored by the pulsar timing (Bañares-Hernández et al., 2025) is $\sim 10^4$ black holes, a global number fraction of 0.1–0.15 per cent against 7×10^6 cluster stars, consistent with the same bottom-heavy IMF and reaching a central, segregation-enhanced value of 1.15 per cent, close to the unbound bracket's own fiducial. The dataset that drives the $H_{\rm q}$–$H_{\rm sub}$ contest of §5 therefore also pins the Monte Carlo's dominant parameter.

The result is structurally simpler than the constraint inventory led us to expect: flybys never set the outer boundary, but the remnant tail controls the lifetime normalization. The median impulsive-diffusion lifetime is roughly flat from the ISCO to 4×10^3 AU, because the kick variance per encounter and the encounter rate scale inversely ($\langle\delta^2\rangle \propto a^{-1}$ from the saturated penetrating passages that dominate it, $\Gamma \propto a$ in the focused regime), so their product is scale-free. Its value is $\sim 3\times10^9$ yr for the stars-plus-WDs function alone, dropping to $\sim 1\times10^9$, $\sim 6\times10^7$, and $\sim 2\times10^7$ yr at black-hole fractions of 0.1, 1, and 3 per cent. The perturber mass matters more than the fraction bracket: at fixed $f_{\rm BH} = 1$ per cent, moving from 10 to 31 $M_\odot$ raises $\langle m_\star^2 \rangle$ by a factor 8.0 and shortens the floor from $\sim 5 \times 10^8$ to $\sim 6 \times 10^7$ yr, a larger displacement than the whole 0.1–3 per cent bracket produces at fixed mass.

Adiabatic decoupling of the unbound channel. That floor is an impulsive estimate, and the impulsive approximation does not hold anywhere in the envelope. For the penetrating encounters that carry the kick variance, the adiabaticity parameter is $x = \omega_{\rm orb} \tau_{\rm enc} = (b/a)(v_{\rm orb}/v)$, which at $b \simeq a$ reduces to $v_{\rm orb}/v_{\rm rel}$: 2.2 at the cluster stripping radius, 4.5 at 10^3 AU, 14 at 10^2 AU, and 4×10^3 at the ISCO. Impulsive behaviour requires $x \ll 1$, which for this system means $a \gtrsim GM/v_{\rm rel}^2 \simeq 2 \times 10^4$ AU, five times outside the envelope. Unbound cluster stars are therefore adiabatically decoupled from every orbit an installation could occupy, and suppressing that correction is conservative in an unbounded way rather than by a stated factor. Applying the Gnedin and Ostriker (1999) kernel $A(x) = (1 + x^2)^{-5/2}$ per encounter, which is calibrated against N-body results and is the less suppressing of the standard choices, lifts the fiducial median to the 12-Gyr cluster-age cap within ~ 3 AU and to $\gtrsim 3 \times 10^8$ yr at the envelope edge, where the suppression is weakest. We carry both curves in Figure 1 and quote the impulsive one as the floor. The scale-free flatness is a property of the impulsive kernel, not of the physical problem: the corrected curve rises inward.

The Bahcall–Wolf cusp. The Monte Carlo uses the core-average perturber density, and a relaxed bound cusp around the hole raises it locally. Earlier drafts of this paper left cusp existence open and carried the correction as a bracket. It is not open. The González Prieto et al. (2025) realizations, which we cite elsewhere for growth history, report a density cusp of slope $\gamma \simeq 1.75$ inside the influence radius forming after 300 Myr, with the stellar population settling at $\gamma \simeq 1.3$ and the black-hole population steepening to $\gamma \simeq 2.0$: the multi-mass structure predicted by Bahcall and Wolf (1976, 1977), heavy species steeper than light. The timescale argument agrees. ω Cen's catalogued relaxation times are ~ 1.1 Gyr in the core and ~ 10 Gyr at the half-mass radius (Harris, 1996); the influence radius lies deep inside the core radius, so the core value governs, and it is roughly a tenth of the cluster age. (A local Spitzer relaxation time evaluated at ρ_0 and σ with $\langle m \rangle = 0.3\, M_\odot$ gives ~ 13 Gyr instead, the convention difference residing in $\langle m \rangle$ and $\ln \Lambda$; we use the catalogued core value and flag the discrepancy.) Either way the segregation clock settles it: cusp formation for a species of mass m proceeds on $\sim (\langle m \rangle / m)\, t_{\rm relax}$, which for a 31 $M_\odot$ remnant is 0.03–0.6 Gyr on either convention.

The cusp case is therefore the fiducial, and it interacts with the adiabatic correction in the direction that matters. Cusp members are *bound*, moving at the local Keplerian speed rather than at $v_{\rm rel}$, so $x \sim 1$ for them and they are not adiabatically protected: the impulsive treatment, with the same Gnedin–Ostriker kernel applied since x is order unity rather than negligible, is valid for the bound channel and invalid for the unbound one. Because $N(<r) \propto r^{3-\gamma} = r$ at $\gamma = 2$, only ~ 1 black hole is expected inside the stripping radius and ~ 10 inside $r_{\rm infl}$: the population is granular, not a smooth density, and we rebuild the Monte Carlo around it rather than estimate it analytically. Each species' cusp density follows $N_s(<r) \propto r^{3-\gamma_s}$ ($\gamma_\star = 1.3$, $\gamma_{\rm NS} = 1.5$, $\gamma_{\rm BH} = 2.0$; González Prieto et al. 2025), normalized against the core-average density at $r_{\rm infl}$; the perturber count for each history is Poisson-drawn from this profile rather than treated as a rate. At the envelope edge a bound black hole is present in 53 per cent of histories; when one is present the diffusion floor is $\sim 4 \times 10^7$ yr (16th percentile $\sim 2 \times 10^7$ yr), and when none is present the channel reaches the 1.2×10^{10}-yr cluster-age cap (Figure 1, panel b; Appendix A). That bimodal bound-member diffusion channel, not the earlier analytic estimate, is the binding clock at the envelope edge: flybys never bind anywhere in the envelope, and the deep envelope is limited by maintenance rather than by dynamics. The same remnant, held at $r_p = 4 \times 10^3$ AU, also forces the orbit secularly on $t_{\rm sec}(a) = (M/m_{\rm pert})(r_p/a)^3 P_{\rm orb}(a) \simeq 1.2 \times 10^7$ yr at $a \simeq 9 \times 10^2$ AU, but relativistic apsidal precession quenches the resonance ($t_{\rm sec}/t_{\rm GR} \propto a^{-4}$) inward of $a \simeq 4 \times 10^2$ AU, where $t_{\rm sec} \simeq t_{\rm GR} \simeq 3.7 \times 10^7$ yr;

a competing quench from apsidal precession forced by the light cusp population is comparable near 10^3 AU and evaluated per history. Folding in the 77.5 per cent mutual-inclination window for the resonance, this forcing is coherent and oscillatory, not diffusive, so it sets a station-keeping cadence over its 4×10^2–4×10^3 AU operating range rather than a survival limit, and is absent below it.

The survival criterion itself needs a second clock. The $e = 0.5$ threshold marks dynamical death, when swarm orbits cross and membership is lost; an unmaintained swarm dies operationally much earlier, because differential kicks pump internal velocity dispersion, and once neighboring orbits cross at their much smaller spacing ($\Delta a/a \sim 10^{-3}$ for a dense swarm) the collisional cascade of §3.3 begins. Scaling the random walk to that threshold gives a grinding-onset time of $(e_{\rm swarm}/e_{\rm cross})^2 \simeq 4 \times 10^{-6}$ of the diffusion floor, $\approx 3 \times 10^2$ yr against the impulsive floor near 1 AU and longer once the adiabatic correction is applied: of order 10^2–10^3 yr, consistent with the abandonment timescale derived independently in §3.3. The two clocks are reconciled by maintenance: the station-keeping needed to null the accumulated differential drift is a small trim budget (the per-encounter kicks are $\Delta v \lesssim$ cm s^{-1} against orbital velocities of 10^2–10^3 km s^{-1}), but it is strictly required. Passively bound is not passively functional, and every long-lived swarm in this paper is a maintained swarm. The operative outer boundary remains the cluster stripping radius of §2.3, $\sim 0.1\, r_{\rm infl} \sim 4 \times 10^3$ AU, with the flyby Monte Carlo setting the maintenance-free lifetime interior to it.

2.5 The envelope

Combining §2.2–§2.4: for the fiducial system, engineered swarms (symmetric or phase-agnostic, per the precession constraint) persist passively from $r \simeq 10^2\, r_{\rm g}$ (2×10^{-2} AU, fueled) or $r_{\rm ISCO}$ (dormant) out to the cluster stripping radius $\sim 4 \times 10^3$ AU. The unbound field is adiabatically decoupled everywhere in that range; the operative floor is the bound-member diffusion channel, $\sim 4 \times 10^7$ yr at the envelope edge in the 53 per cent of realizations that draw a bound heavy remnant, and the cluster-age cap otherwise. Beyond the stripping radius, orbits bound to the hole are not durable and only cluster-orbiting architectures (outside this paper's scope) remain.

The envelope is the paper's first result (Figure 1). Two corollaries matter downstream. First, the thermodynamically favored location (deep, near the flow) is also the dynamically safest, doubly so once adiabatic protection is counted: the environment does not penalize the architecture Paper A's physics prefers, which is a nontrivial consistency check the hypothesis could have failed. Second, the envelope concentrates any engineered mass within $\sim 4 \times 10^3$ AU $\sim 10^{-1}\, r_{\rm infl}$ of the hole, with the thermodynamic gradient pushing occupation deep into the arcsecond-unresolvable interior: observable only through its energetic and dynamical residue.

That is the subject of §3, after two remaining environmental questions: hazards other than flybys, and fuel.

2.6 Hazards beyond flybys

Three further environmental processes deserve quantitative treatment. Each is a candidate defeater; none defeats, though the first comes closest.

Tidal disruption events. A cluster-core IMBH tidally disrupts stars scattered into its loss cone. Rates for IMBHs in evolved globular clusters, computed with loss-cone methodology of the kind developed for the supermassive case (Stone and Metzger, 2016), are $\sim 10^{-8}$–10^{-7} yr^{-1} per cluster for main-sequence stars, with white-dwarf disruptions rarer by a factor of $\sim$ 30–100 (Fragione et al., 2018); the ω Cen-specific value from the González Prieto et al. (2025) models is $\sim 5 \times 10^{-8}$ yr^{-1}, and we use that number where a target-specific rate is needed. One reciprocal flag for Paper A: this is also the natural loss-cone supply

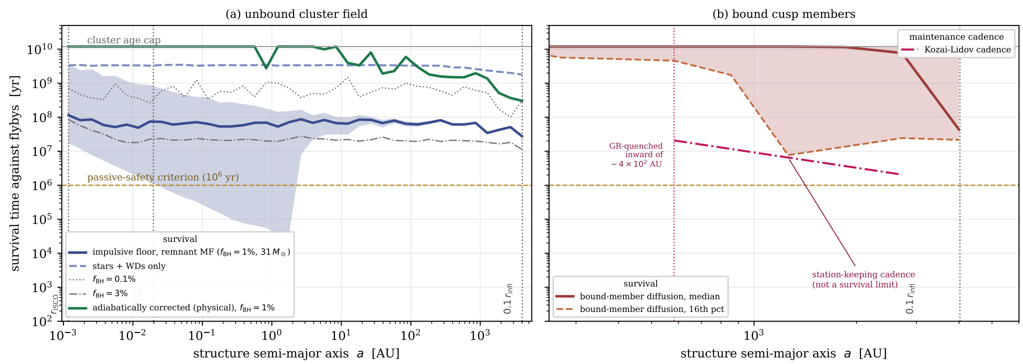

Figure 1: Survival time of hole-bound swarm structures against cumulative perturbation, from the rebuilt Monte Carlo of §2.4 (2×10^4 histories per radius bin; grid begins at the Schwarzschild ISCO). **(a) Unbound cluster field.** Blue curve and band: median and 16–84 per cent range of the impulsive estimate for the fiducial mass function with segregated remnants ($31\,M_\odot$ black holes at 1 per cent number fraction, $1.4\,M_\odot$ neutron stars at 2 per cent); dashed: the stars-plus-WDs baseline; thin curves: black-hole fractions of 0.1 and 3 per cent. Green curve: the same fiducial case with the Gnedin and Ostriker (1999) adiabatic correction applied per encounter, which shows this channel is decoupled from every orbit in the envelope; it does not operate and is retained as the measure of that decoupling. **(b) Bound cusp members.** Dark red curve and band: median and 16–84 per cent range of the bound-member diffusion channel, jointly anchored perturber population ($f_{BH,central} = 1.15$ per cent, $31\,M_\odot$); the population is Poisson-granular, so the 16th-percentile curve (dashed) is emphasized alongside the median because the distribution is bimodal by realization, not because the band is asymmetric noise. Magenta dash-dot: the Kozai–Lidov station-keeping cadence forced by the same bound remnant, drawn only over its GR-quenched-to-mass-quenched operating range and labeled on the curve as a maintenance requirement, not a third survival curve. This bound channel, not the decoupled field in panel (a), is the operative floor: flybys never bind anywhere in the envelope, and the deep envelope is limited by maintenance rather than by dynamics. Both panels: the band is a Monte Carlo precision interval, not a physical dispersion, and is narrow wherever the first-passage step count is large; vertical markers show the inner boundaries and the cluster stripping radius.

rate, and the staged architecture of Paper A §6 requires 10^{-4}–10^{-3} stars yr^{-1} during spin-up, so natural refilling falls short by three to four orders of magnitude and essentially the whole fuel budget of that phase must be delivered by engineered orbit-shaping. During a disruption the accretion luminosity approaches Eddington, $L \sim 3 \times 10^{42}$ erg s^{-1} for the fiducial mass, sustained for months to years of fallback: the flux at $r = 1$ AU is $\sim 10^9$ times the solar constant, and no plausible material hardening survives it in place. Feasibility therefore requires that a TDE be survivable by response rather than endurance: loss-cone stars are in principle identifiable and trackable long before disruption (the swarm sits at the bottom of the potential it would need to monitor), and the months-long fallback rise gives further warning, so temporary evacuation outward along the envelope, or shadowing at large inclination, is an engineering requirement we impose on any architecture rather than a reason to exclude one. The expected recurrence time of $\gtrsim 10^7$ yr then sets the natural amortization horizon of the installation: a 10^8–10^9-yr occupation must plan for several such events. Two observational corollaries. The TDE duty cycle also reconciles $H_{\rm eng}$ with the present silence at no extra cost: post-TDE fallback outshines any steady configuration for $\sim 10^2$–10^3 yr per $\sim 10^7$-yr recurrence, a duty cycle of $\lesssim 10^{-4}$, so catching ω Cen mid-flare was never likely under any hypothesis. And TDEs cut the other way, as an observational gift: any future TDE flare in ω Cen both confirms the hole and activates the variability channel of §3.2 at high signal-to-noise, and the framework of §4 treats a TDE with anomalous light-curve regulation as one of the few single-epoch events that can move $\ln K$ substantially.

Black-hole wander. The hole is a Brownian particle in the stellar bath: N-body characterizations give r.m.s. displacements of order 10^{-2}–10^{-1} pc for 10^3–$10^4\,M_\odot$ holes in ω Cen-like cores, with wander amplitude scaling inversely with hole mass (de Vita et al., 2018). For the fiducial $2 \times 10^4\,M_\odot$ hole the expected wander is $\lesssim 10^{-2}$ pc $\sim 2 \times 10^3$ AU: comparable to the 4×10^3 AU envelope. That figure is a lower bound on two counts. The equilibrium amplitude in a multi-mass bath is set by the mass-weighted mean $m_{\rm eff} = \langle m^2\rangle/\langle m\rangle$ rather than by $\langle m\rangle$, since the velocity diffusion scales with $n\langle m^2\rangle$ while the drag scales with $n\langle m\rangle$ (Chatterjee et al., 2002; Merritt, 2005); for the segregated mass function of §2.4 that is $m_{\rm eff} \simeq 2.3\,M_\odot$ against $\langle m\rangle = 0.54\,M_\odot$, so an equal-mass estimate understates the amplitude by a factor of a few. And the quoted N-body results predate the retained black-hole populations now inferred for ω Cen, so the mass function they carried is lighter than the one §2.4 adopts. Paper C's wander program (§4.2 there) inherits both corrections, in the direction that makes the measurement easier. Structures bound to the hole simply ride along (the swarm orbits the hole, and hole plus swarm wander together through the cluster), so wander does not threaten the installation; what it threatens is the observer's astrometry, since the kinematic center and the hole need not coincide at the 10^{-2}-pc level. Paper C's astrometric programs already marginalize over center position; we note here that the wander amplitude is itself mass-dependent and therefore carries independent information about M in long-baseline proper-motion data.

Gas drag and erosion. ω Cen retains a measurable intracluster medium: sodium-absorption mapping and pulsar dispersion in comparable clusters give central ionized densities $n_e \sim 0.1$–$0.3\,{\rm cm}^{-3}$ (Freire et al., 2001; Abbate et al., 2018; Wang et al., 2025). Ram-pressure drag on a swarm element of areal density $\Sigma \sim 1\,{\rm g\,cm}^{-2}$ at orbital velocity of a few $\times 10^3$ km s^{-1} (the deep envelope; 4×10^3 km s^{-1} at the 1 AU fiducial) removes a fractional momentum $\sim \rho_{\rm gas} v t/\Sigma$ per time t: at the measured densities this is $\lesssim 10^{-6}$ per kyr, negligible. Sputtering and dust impacts are similarly small in an old cluster with no star formation and depleted debris populations. Gas matters for fuel, however, which is the next subsection.

2.7 The fuel budget

The Blandford–Znajek architecture of Paper A needs mass supply, and the measured medium bounds it. Bondi accretion from gas of density $n_e \simeq 0.23\,\mathrm{cm}^{-3}$ (Wang et al., 2025) at relative velocity $\sim \sigma$ gives

$$\dot{M}_\mathrm{B} = \frac{4\pi(GM)^2\rho_\mathrm{gas}}{(\sigma^2 + c_s^2)^{3/2}} \simeq 3\times10^{18}\,\mathrm{g\,s^{-1}}\left(\frac{M}{2\times10^4\,M_\odot}\right)^2\left(\frac{n_e}{0.23\,\mathrm{cm}^{-3}}\right) \simeq 5\times10^{-8}\,M_\odot\,\mathrm{yr}^{-1}, \tag{6}$$

about 10^{-4} of the Eddington rate. Two caveats bound this from both sides. It is an upper bound insofar as the density at the hole's actual location may be below the cluster mean: the accretion non-detections are consistent with the hole occupying a locally evacuated region (Mahida et al., 2026), and the Bondi radius ($\sim GM/c_s^2 \sim 2\times10^5$ AU) samples gas the surveys average over. It is a lower bound insofar as it ignores harvesting: stellar winds from the $\sim 10^3$ giants inside the influence radius, or deliberately imported mass, can raise supply by orders of magnitude at the cost of visibility. At MAD-plus-spin effective efficiencies of order unity, realized only at high spin ($\eta_\mathrm{jet} \simeq 1.4$ at $a_\star \simeq 0.99$, falling roughly as $a_\star^2$ to a factor ~ 5 lower at $a_\star = 0.5$; Paper A; Tchekhovskoy et al. 2011), the ambient rate gives extractable power $P \sim \dot{M}_\mathrm{B}c^2 \simeq 3\times10^{32}$ W $\simeq 8\times10^5\,L_\odot$, with the fuel ceiling carrying the same $\eta(a_\star)$ dependence: the ambient medium alone, with no harvesting of stellar winds or imported fuel, funds a computational budget two orders of magnitude above the waste-heat ceiling derived in §3.1.

That headroom assumes the Bondi rate reaches the horizon, and in natural hot, low-Eddington flows it does not: the inflow declines inward, $\dot{M}(R) \propto R^s$ with $s \simeq 0.3$, as most of the captured gas is unbound into outflows (Blandford and Begelman, 1999; Yuan and Narayan, 2014). Over the nine decades from the Bondi radius ($r_\mathrm{B} \simeq 1.8\times10^5$ AU) to r_g, the suppression is $(r_\mathrm{B}/r_\mathrm{g})^{0.3} \approx 500$, giving $P_\mathrm{fuel} \simeq 1.6\times10^3\,L_\odot$; adopting instead the empirical Sgr A* figure of ~ 1 per cent of the Bondi-radius rate gives $\simeq 8\times10^3\,L_\odot$. Either value lands at or below the warm-swarm waste-heat ceiling of §3.1, so for a natural flow the supply and concealment constraints bind together, and the "supply never binds" ordering holds only for the engineered case. The engineered case is available, and it unifies two of this paper's results: a system drawing more than the ADIOS-suppressed rate is a system that has suppressed its own outflow, and that is the same control action, regulating horizon magnetic flux and mass supply, that produces the low-R variability deficit of §3.2. The fuel budget and the R statistic are two faces of one control action. We therefore carry both cases through the paper: natural-suppressed ($P_\mathrm{fuel} \sim 10^3$–$10^4\,L_\odot$) and engineered-unsuppressed ($P_\mathrm{fuel} \simeq 8\times10^5\,L_\odot$), with the Appendix B prior ceiling stated for the pair.

For an installation with outflow control, then, the binding constraint on a present-day ω Cen system is thermal (getting rid of entropy without detection), never supply; without it, the two constraints meet. The unsuppressed ordering is a nontrivial output: for stellar-mass holes in the same environment the two constraints reverse even in the engineered case, which is an independent reason the hypothesis family selects IMBHs. The same numbers close the loop on the accretion non-detection: the deep ATCA and JWST limits require the natural radiative efficiency of any Bondi-fed flow to be $\lesssim 4\times10^{-3}$ (Mahida et al., 2026; Chen et al., 2025), which is uncomfortable for H_q at the upper end of the mass range and what H_eng predicts for a flow whose output is being extracted as work rather than radiation. The adjudication in §5 prices this observation for both hypotheses rather than letting either claim it informally.

3 Observable residue

3.1 Waste-heat floor

Paper A argued that horizon entropy disposal permits computation with radiated waste heat far below the Dyson-sphere expectation. Here we make the floor quantitative, and we first fix the bookkeeping, because the mainstream position is an energy-conservation argument: Curtis et al. (2026) state that Dyson-scale computation must reradiate nearly all of the energy it absorbs, making mid-infrared thermal emission the robust technosignature. The horizon-sink architecture disputes that position at the level of energy conservation itself, not merely entropy accounting. Let P_{comp} be the power processed by the swarm and f_{sink} the fraction of the waste *energy* delivered across the horizon rather than radiated; the entropy rides with its carriers, so a single fraction serves for both. Energy beamed across the horizon is not reradiated at all: it adds to M, and since $T_H \propto M^{-1}$ the disposal channel deepens as it is used, a weakly self-improving property. What must be radiated is only the missed fraction, $L_{\mathrm{waste}} = (1 - f_{\mathrm{sink}})\, P_{\mathrm{comp}}$, re-emitted thermally at the temperature set by the swarm radius. For single-sided radiators at unit covering fraction, $T_{\mathrm{eff}} \simeq (L_{\mathrm{waste}}/4\pi r^2 \sigma_{\mathrm{SB}})^{1/4}$: $1\, L_\odot$ at $r = 1$ AU gives 394 K, $1\, L_\odot$ at 10^3 AU gives 12.5 K, and a 50 K swarm at 10^3 AU corresponds to $L_{\mathrm{waste}} \approx 260\, L_\odot$. This $T_{\mathrm{eff}}(r)$ assumes unit covering fraction with single-sided radiators; a sparse swarm of covering fraction f_{cov} runs hotter by $f_{\mathrm{cov}}^{-1/4}$, moving mass into the JWST wedge and shifting the composite $\ln K$ of §5 to -0.34 at $f_{\mathrm{cov}} = 10^{-2}$, inside the quoted band.

The engineering limit on f_{sink} is transport: waste energy must be carried inward, against the swarm's own power draw, by mass flux or directed radiation into the horizon. The capture-cone and étendue accounting of Appendix A.3 bounds $1 - f_{\mathrm{sink}} \gtrsim 10^{-4}$ for swarm architectures within the envelope, set by the fraction of beamed power unavoidably intercepted and re-thermalized by swarm elements along the transport path. The sign of its influence is worth stating: a lower true floor (better beaming than the A.3 accounting allows) weakens the mid-infrared charge against H_{eng} and moves $\ln K$ toward zero, while a higher floor tightens the P_{comp} ceiling and strengthens it (§6). The floor is therefore

$$L_{\mathrm{waste}} \gtrsim 10^{-4}\, P_{\mathrm{comp}}. \tag{7}$$

Converting that floor into a P_{comp} bound requires a temperature axis, because which instrument limits L_{waste} depends on where the thermal emission peaks. The deep JWST limits at the kinematic center (Chen et al., 2025) constrain warm sources: a swarm at $r \lesssim$ a few AU re-radiates at $T_{\mathrm{eff}} \gtrsim 150$ K, peaks within MIRI coverage, and inherits the sub-$L_\odot$ point-source limit, giving $P_{\mathrm{comp}} \lesssim L_{\mathrm{lim}}/(1 - f_{\mathrm{sink}}) \lesssim 10^3$–$10^4\, L_\odot$. A cool outer swarm evades it: at $r \sim 10^2$–10^3 AU the emission peaks at 20–60 μm and beyond, past the coverage of JWST/MIRI's nine imaging filters (F560W through F2550W, the 25.5 μm cutoff); WISE W3/W4 and Spitzer/MIPS, whose 6–18 arcsec beams are confusion- and crowding-limited in the ω Cen core, are the instruments that would bound such a cool swarm observationally, though Appendix B does not draw on them directly. The single-number bound of an earlier draft implicitly assumed the warm case; the defensible statement is radius- and temperature-dependent, and the continuous per-filter Appendix B likelihood (§C) integrates the two-temperature swarm SED against the real MIRI sensitivity curve rather than marginalizing over instrument-specific placeholder limits, so a sufficiently cool, extended swarm falls below every MIRI filter and registers as undetected rather than excluded. An engineered system operating at the radius-appropriate ceiling is consistent with all current data; deeper MIR photometry tightens the warm wedge linearly, while the cool wedge waits on far-infrared sensitivity. Figure 2 assembles the warm-swarm plane and the temperature-resolved plane, with the transport floor and the fuel ceiling of §2.7, into the constraint space that defines the surviving H_{eng} parameter volume.

3.2 The MAD-regulation signature

The second channel is new to this paper. General-relativistic magnetohydrodynamic simulations establish that natural accretion in the magnetically arrested state is intrinsically episodic: magnetic flux accumulates at the horizon, chokes the inflow, and erupts in quasi-regular flux-expulsion events, producing characteristic variability in jet power and radiative output with a MAD > intermediate > SANE hierarchy in both luminosity and its variance (Tchekhovskoy et al., 2011; Narayan et al., 2022; Yang et al., 2026). The

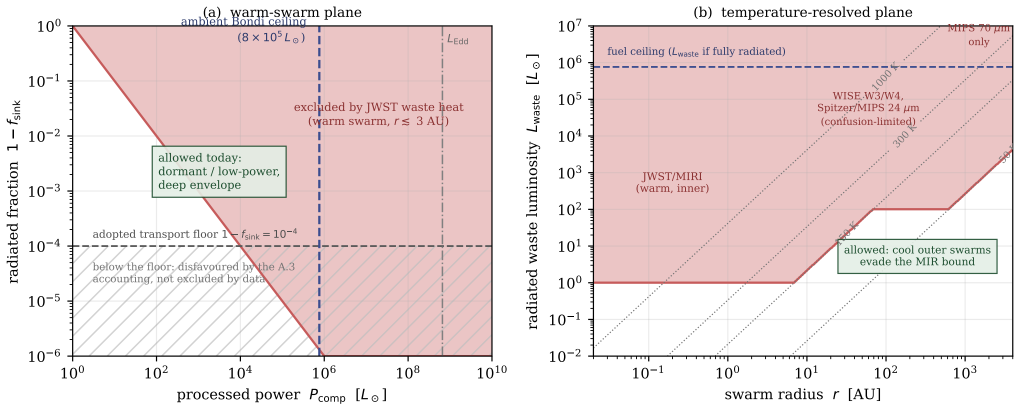

Figure 2: Constraint planes for an ω Cen installation. **(a)** The $(P_{\rm comp},\ 1 - f_{\rm sink})$ plane for a warm (inner-envelope) swarm. Red: excluded by the JWST mid-infrared point-source limit. Grey hatching below the dashed grey line: the adopted entropy-transport floor $1 - f_{\rm sink} = 10^{-4}$, a parameter carried over from the Appendix A.3 accounting rather than a data-driven boundary (§6). Blue dashed: the ambient Bondi fuel ceiling ($8 \times 10^5\ L_\odot$, the engineered-unsuppressed case; the ADIOS-suppressed natural supply is $\sim 10^3\ L_\odot$, §2.7; imported mass moves it right at the cost of visibility). **(b)** The temperature-resolved $(r,\ L_{\rm waste})$ plane: dotted lines are $T_{\rm eff}$ contours, and the excluded wedge is the radius-dependent instrument limit (JWST/MIRI for warm inner swarms; WISE W3/W4 and Spitzer/MIPS, confusion-limited in the core, for cool outer ones). The surviving region is dormant, low-power, or cool and outer.

eruption mechanism is plasmoid-mediated reconnection of the accumulated horizon flux, resolved ab initio by Ripperda et al. (2022), whose flare cycle of $\sim 10^2$–$10^3\ r_{\rm g}/c$ between flux accumulation and expulsion is the recurrence we adopt; the state-resolved variability tables of Yang et al. (2026) place the MAD eruption power in the same interval at matched mean luminosity. For $M = 2 \times 10^4\ M_\odot$ that cycle is ~ 10–10^2 s, comfortably within the cadence of X-ray monitoring.

A system engineered for steady power extraction has an incentive absent in nature: eruptions are interruptions. Regulating the delivered mass and flux (the control variable identified by the simulations is horizon magnetic flux itself) suppresses the eruption cycle and its variability signature. We therefore forward-model the residue as a *variability deficit*: an accreting IMBH whose X-ray (or optical) power spectrum lacks the flux-eruption band that GRMHD baselines predict for its luminosity and inferred state. Quantitatively, we define the regulation statistic R as the ratio of integrated power in the eruption band (10^{-2}–10^{-1} Hz for the fiducial mass) to the GRMHD-calibrated expectation at matched mean luminosity; natural MADs populate $R \sim 1$ with factor-of-few scatter, and $R \ll 10^{-1}$ sustained over many cycle times has no identified natural mechanism (Figure 3). The band edges follow directly from scaling the Ripperda et al. (2022) flare cycle of 10^2–$10^3\ r_{\rm g}/c$ to $M = 2 \times 10^4\ M_\odot$ ($r_{\rm g}/c = 0.1$ s, hence 10–100 s), with the Yang et al. (2026) state tables placing the eruption power in the same interval; an earlier draft extended the band a decade lower, which nothing in the cited cycle supports. The main confusion risk in applying the statistic is contamination from the intermediate magnetic state, whose eruption power is partially suppressed relative to MAD, so the baseline must be conditioned on the inferred state, never on luminosity alone.

Two instrument realities shape where R can be measured. First, at the fiducial mass R is an X-ray/optical statistic; radio is excluded twice over. Synchrotron self-absorption in any compact GHz-emitting flow puts the $\tau = 1$ photosphere at $\sim 10^2$–$10^4\ r_{\rm g}$, low-pass filtering intrinsic variability at $\Delta t \gtrsim 10$–10^3 s, directly on top of the eruption band, so a radio non-detection of the band is expected under $H_{\rm q}$ and $H_{\rm eng}$ alike and carries no discriminating power. Interstellar scintillation compounds the problem: a microarcsecond-scale source at 5.4 kpc scintillates at GHz frequencies on minutes-to-hours timescales,

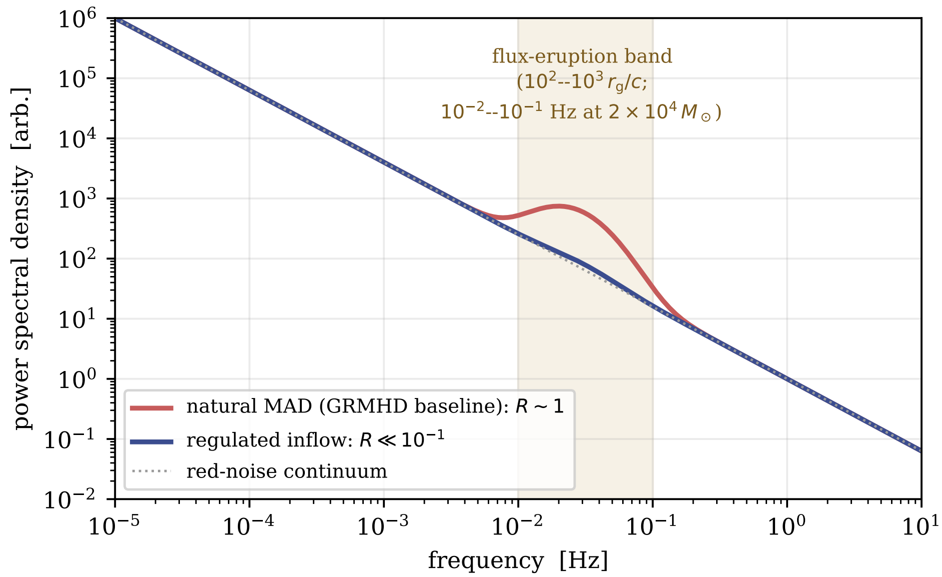

Figure 3: Schematic definition of the regulation statistic R: power spectra of an accreting IMBH at matched mean luminosity under natural MAD accretion (flux-eruption band populated, $R\sim1$) and regulated inflow ($R\ll10^{-1}$). Band location scales inversely with black-hole mass; values shown for $M=2\times10^4\,M_\odot$.

adding modulation uncorrelated with accretion in and near the band, so even a genuinely regulated (low-R) source would show a measured R biased upward unless the scintillation is modeled out. Second, the statistic has a detectability threshold. Distinguishing $R\sim1$ from $R<0.1$ requires several counts per ~10 s cycle sustained over many cycles, i.e., $F_X\gtrsim$ a few $\times10^{-12}$ erg cm^{-2} s^{-1} for m^2-class effective area, corresponding to $L_X\gtrsim10^{34}$–10^{35} erg s$^{-1}\approx10^{-8}$–$10^{-7}\,L_{\rm Edd}$ at the fiducial mass. The channel therefore activates in flare and TDE states rather than in deep quiescence, which is consistent with the TDE-as-gift framing of §2.6: the events that confirm the hole are the events bright enough to measure R.

The signature has three properties valuable for adjudication. It is conditional: it activates only if accretion is ever detected, so it presently constrains nothing (ω Cen's hole is electromagnetically silent; Mahida et al. 2026). It is differential: it compares a source against a physics baseline at matched parameters, canceling many systematics. And it is disprovable in place: detection of normal flux-eruption variability from a future ω Cen accretion flare would count against $H_{\rm eng}$ in the framework below, making the channel one of the few that can move evidence in both directions.

3.3 Relic residue: the abandoned swarm

Persistence is conditional on maintenance, and Lacki (2025) has shown what maintenance failure means for megastructures: once guidance fails, swarm elements collide, and a collisional cascade grinds the population to dust on a timescale of roughly the orbital period divided by the covering fraction. The deep envelope makes this dramatic. At $r\sim1$ AU around the fiducial hole the orbital period is ~3 days; for covering fractions 10^{-4}–10^{-2} the first-collision timescale $\sim P/f_{\rm cov}$ is years to decades, and the full cascade completes on a timescale of order 10^3 yr, orders of magnitude faster than for the stellar-orbit megaswarms Lacki considered and consistent with the grinding-onset clock of §2.4. The end state differs too, and the difference is observationally decisive: cascade debris around a star settles into a long-lived warm dust population, a passive relic technosignature, whereas debris around a black hole is progressively drained, its periapsis distribution fed toward the loss cone by continuing collisions and flyby perturbations and its finest grindings coupled to whatever gas flow exists; the drainage rate remains an estimate, qualified in §6. The hole removes its own debris. An abandoned engineered-IMBH system therefore passes through a brief bright phase (cascade grinding plus enhanced accretion luminosity as debris drains) and then reverts to a naked quiescent hole, indistinguishable from H_q except possibly through spin.

Three consequences follow. First, a fourth hypothesis joins the menu at zero structural cost: H_{eng}^{relic}, a formerly engineered now-abandoned system, whose observables today equal H_q plus high spin, and whose prior couples to the goal-stability open problem flagged in Paper B. Second, the searchable relic window is short, so population-level searches for dead installations (the analogue of the dust-relic searches now proposed for stellar systems) have low yield around IMBHs specifically; non-detection of relic dust in ω Cen carries almost no evidence either way, and the framework scores it accordingly. Third, spin becomes the only durable fossil, and an objection must be met before relying on it: Blandford–Znajek extraction spins the hole down, and MAD jets remove angular momentum faster than accretion supplies it (Narayan et al., 2022), so one might expect long use to erase the very fossil we propose reading. The rates rescue the argument: in the MAD spin-down calculus of Narayan et al. (2022), the spin-down per unit accreted mass is of order $|\Delta a_\star| \sim \Delta M/M$, so erasing a high spin requires cycling a mass comparable to the hole's own through the flow, and at the ambient fueling of §2.7 that takes $M/\dot{M}_B \sim 4 \times 10^{11}$ yr. Only a civilization importing mass at far above ambient rates for $\gg 10^9$ yr would measurably despin the hole, and that regime is separately excluded by the waste-heat ceiling.

What spin discriminates deserves a more careful statement than our earlier draft gave it, including an assumption the earlier draft carried without stating it. The natural expectation is a channel mixture, set by growth history rather than by any single attractor: repeated comparable-mass mergers drive χ toward the $\simeq 0.7$ attractor (Fishbach et al., 2017; Gerosa and Berti, 2017); gas-poor growth by minor mergers and inspirals of compact objects random-walks the spin to low values under isotropic capture; and prolonged coherent disk accretion spins the hole up toward the radiation-limited maximum $a_\star = 0.998$ (Bardeen, 1970; Thorne, 1974). Which of the three operated at ω Cen is a question its growth models answer, conditional on the isotropy assumption stated next.

For the minor-merger regime, isotropic capture ($\langle \cos\iota \rangle = 0$ for the captured objects' orbital planes) gives $a_{\star,\rm rms} = (\tilde{L}/\sqrt{3})\sqrt{\mu/M_f}$ (Appendix B); the assumption is load-bearing for everything that follows and we state it here rather than leave it implicit. A net alignment fraction $\epsilon_{\rm rot} = \langle \cos\iota \rangle$ above isotropic adds a coherent term $a_\star^{\rm coh} \approx (1.7\text{–}2.0)\,\epsilon_{\rm rot}$ that carries no $\sqrt{N}$ suppression, so it matches the random walk at $\epsilon_{\rm rot} \simeq 0.035$ and dominates above it, erasing the low-spin prediction entirely once $\epsilon_{\rm rot} \gtrsim 0.4\text{–}0.5$. The prediction below holds only for $\epsilon_{\rm rot} \lesssim 0.035$. ω Cen is a rotating cluster and the inspiraling population is the mass-segregated heavy-remnant tail, so isotropy is a premise to check rather than a safe default; $\langle \cos\iota \rangle$ for the inspiraling set is a direct output of the González Prieto et al. (2025) realizations and has not yet been extracted from them. §3.3 carries the tolerance as an explicit condition on the green branch.

The González Prieto et al. (2025) realizations (cited here for growth history only; they do not track spin) grow a 500–5000 $M_\odot$ seed to $\sim 5 \times 10^4\,M_\odot$ primarily by mergers with 30–40 $M_\odot$ black holes, roughly 10^3 events at seed-epoch mass ratio $q \simeq 0.006$–0.08; by the end of the growth history the ratio the final spin depends on is the smaller $\mu/M_f \simeq 7 \times 10^{-4}$. That is the isotropic minor-merger regime, not the comparable-mass one: the walk reaches the $\chi \simeq 0.7$ attractor only at $q \approx 0.12$ (Appendix B, a continuous scaling rather than a literature-sourced cutoff), and ω Cen sits a factor ~ 175 below that boundary on the final-mass ratio. So the $\chi \simeq 0.7$ attractor of Fishbach et al. (2017) and Gerosa and Berti (2017) does not apply to this target; under isotropic capture the expected outcome, derived rather than imported from a comparable-mass literature value (Appendix B), is $a_\star \sim 0.05$–0.10 (central ≈ 0.06). An earlier draft imported the generic hierarchical-merger caution and reduced the discriminating contrast to $\simeq 0.7$ versus

$\gtrsim 0.9$ on that basis. For ω Cen the contrast is the sharper low-versus-high one, and the attractor caution belongs to the class-level statement rather than to this system. The information forecast of §5 restores the spin channel accordingly, with the caveat that the class-level application to the LISA IMBH population must carry the mixture, since comparable-mass growth histories elsewhere populate the attractor.

3.4 EMRI dephasing: the mass channel

Every residue channel above constrains power, variability, or spin. One channel constrains engineered *mass*, and it comes for free with the gravitational-wave observation the series already relies on. Matter in the vicinity of an inspiral imprints secular phase shifts on the waveform; the environmental-effects literature for extreme- and intermediate-mass-ratio inspirals, founded by Barausse et al. (2014), quantifies how accretion disks, dark-matter spikes, and other mass distributions dephase an inspiral over an observation. An engineered swarm is a mass distribution like any other. If a stellar-mass compact object is ever caught inspiraling into the ω Cen hole, the waveform is a scale: it weighs whatever hardware shares the inspiral's neighborhood, independent of that hardware's luminosity, temperature, or duty cycle.

The channel's reach must be stated carefully, because it is short. The secular dephasing from an exterior axisymmetric mass falls as $(r_{\rm orb}/r_{\rm torus})^3$ per orbit, so a debris torus or swarm at 10^2–10^3 AU accumulates negligible phase against an inspiral at $\sim 10^{-3}$–10^{-2} AU; we make no claim of debris-torus reach, and the channel does not repair the relic-searchability gap of §3.3. What it constrains is mass at or inside the inspiral track, $r \lesssim 10^{-2}$ AU, and that limitation is well matched to the hypothesis: Paper A's thermodynamic gradient concentrates hardware at exactly those depths, and the envelope of §2 permits it there. A LISA-band inspiral through an occupied deep envelope would traverse the swarm itself. An order-of-magnitude scope for the fiducial system: an inspiral accumulates $\sim 10^5$–10^6 radians of orbital phase over a multi-year observation, and matched filtering resolves phase drifts of order unity, so fractional perturbations at the 10^{-6}–10^{-5} level in the enclosed mass along the track are in principle measurable, a sensitivity to swarm masses far below any other channel's floor. Turning that scope into a forecast requires real waveform modeling, including the degeneracies with the astrophysical environmental effects already cataloged in the literature; we defer that calculation, and any figure it would support, rather than publish an unmodeled curve.

Two properties carry into the adjudication. The channel is two-sided, like R: a measured *vacuum* inspiral, dephasing consistent with zero environmental mass, counts against $H_{\rm eng}$ at the depths the hypothesis most values, and the framework of §4 scores it accordingly. And it inherits the spin channel's schedulability problem: the generic per-IMBH IMRI rate is $\sim 10^{-9}$–$10^{-6}\,{\rm yr}^{-1}$ (Mandel et al., 2008; Amaro-Seoane, 2018; Fragione et al., 2018), with Arca Sedda et al. (2021) the current superseding compilation for dense-cluster rates (LISA detection 0.02–60 ${\rm yr}^{-1}$, in-cluster IMRI formation probability 5–50 per cent rising with IMBH mass), and ω Cen's own rate is narrower still, $(4–8) \times 10^{-8}\,{\rm yr}^{-1}$ (§5), so the channel is population-level, a constraint LISA accumulates across the Galactic IMBH census rather than a scheduled ω Cen observation. It enters the information forecast of §5 as the fourth channel on those terms.

3.5 Predicted-signature table

Table 1 summarizes what each channel looks like under the engineered hypothesis and the two astrophysical nulls defined in §4.1.

The table exposes the central inferential difficulty, stated plainly in Paper A and now quantified: silence is predicted by everything. $H_{\rm eng}$ and $H_{\rm q}$ are near-degenerate in every electromagnetic channel while the system is dormant. The channels that separate them are spin (engineered selection favors $a_\star \gtrsim 0.9$; the

Table 1: Forward-modeled observables under the three hypotheses of §4.1. "Silent" means below current limits. GRMHD baseline from Yang et al. (2026).

Channel	$H_{\rm eng}$ (engineered)	$H_{\rm q}$ (quiescent IMBH)	$H_{\rm sub}$ (remnant subcluster)
MIR point source	floor at $10^{-4} P_{\rm comp}$; may sit near current limits	silent; $L \propto \dot{M}$ of ambient gas only	silent; no compact source
Radio continuum	silent while dormant; self-absorbed if fueled (no R information; §3.2)	silent or canonical fundamental-plane track	silent
Variability (R statistic, X-ray/optical)	$R \ll 1$ if accreting above the §3.2 threshold	$R \sim 1$ if accreting	n/a
EMRI/IMRI dephasing	environmental phase drift if hardware sits at or inside the inspiral track (§3.4)	vacuum inspiral	no inspiral (no massive central object)
LISA EMRI spin	high-tail spin, $a_\star \gtrsim 0.9$ (selection at arrival; §3.3)	low spin, $a_\star \sim 0.05$–0.10, from the minor-merger growth history the ω Cen models find under isotropic capture (González Prieto et al. 2025, Appendix B); the $\chi \simeq 0.7$ attractor applies to comparable-mass histories elsewhere in the class (Fishbach et al. 2017; Gerosa and Berti 2017)	no EMRI, but a comparably dense stellar-mass black-hole population instead produces resolvable mHz binaries (SNR $\sim$ few$\times 10^4$ at 2 mHz over a 4-yr mission), a distinct LISA signature in its own right
Fast-star kinematics	IMBH-like point mass	IMBH point mass	extended mass profile
MSP timing	point-mass potential	point-mass potential	extended potential (current data lean this way; Bañares-Hernández et al. 2025; Colom i Bernadich et al. 2026)

growth history modeled for ω Cen favors $a_\star \sim 0.05$–0.10 under isotropic capture, §3.3), the R statistic (active only during accretion episodes above the §3.2 threshold), EMRI dephasing (active only if an inspiral occurs; §3.4), and the waste-heat floor (which separates them only near the sensitivity ceiling). This degeneracy structure, rather than any single measurement, is what the adjudication framework must manage.

4 The adjudication framework

4.1 Hypothesis space

We adjudicate among three hypotheses about the ω Cen center:

- H_q: a quiescent IMBH of mass 10^4–$5 \times 10^4\, M_\odot$, electromagnetically silent for want of fuel, spin drawn from an uninformative prior. The parsimonious reading of the fast stars (Häberle et al., 2024b).
- H_sub: no IMBH; the central mass is an extended subcluster of stellar remnants of $\simeq 2$–$3 \times 10^5\, M_\odot$ within the core, as favored by the pulsar-timing analysis (Bañares-Hernández et al., 2025).
- H_eng: an IMBH hosting engineered infrastructure within the envelope of §2, with observables forward-modeled in §3 and parameters $(P_\mathrm{comp}, f_\mathrm{sink}, R, a_\star)$ assigned the priors stated in Appendix B.

The menu is explicit and extensible; adding hypotheses (an exotic compact object, instrumental systematics on the fast stars, the abandoned-system variant $H_\mathrm{eng}^\mathrm{relic}$ of §3.3) changes bookkeeping, never structure. $H_\mathrm{eng}^\mathrm{relic}$ is degenerate with H_q on every current channel (§3.3) and is therefore not separately scored until a spin measurement exists. We follow the convention that H_eng is reported only as odds against the best-performing null, never against a strawman.

4.2 Per-messenger Bayes factors

For each observational channel i with data d_i, the Bayes factor between hypotheses H_a, H_b is

$$K_i^{ab} = \frac{\int \mathcal{L}(d_i \,|\, \theta_a)\, \pi(\theta_a)\, d\theta_a}{\int \mathcal{L}(d_i \,|\, \theta_b)\, \pi(\theta_b)\, d\theta_b}, \tag{8}$$

with channel likelihoods built from the forward models of §3 and, following radio-SETI practice, an explicit interference/systematics component mixed into every likelihood (Sheikh et al., 2021; Li et al., 2026) so that "instrumental artifact" is priced inside each channel rather than adjudicated informally afterward. Priors on H_eng parameters are pre-registered (Appendix B) and varied over stated ranges in sensitivity analysis; the framework's outputs are always reported as (Bayes factor, prior-sensitivity band) pairs, never as bare numbers.

One evaluation rule governs the interface with Paper C's campaign, and it matters more than it looks. Paper C promotes a candidate to adjudication when a pre-registered frequentist trigger fires, and that promotion is data-dependent. The likelihoods above are therefore always evaluated on the complete campaign record, every epoch and every channel including the ones that recorded nothing, never on the promoted subset alone. The trigger layer allocates follow-up effort and carries no inferential weight. Scoring the selected data as though it were the whole dataset would require the conditional form $K \mid \text{trigger} = K \times P(\text{trigger} \mid H_b)/P(\text{trigger} \mid H_a)$, and at Paper C's global threshold of $p < 5 \times 10^{-7}$ that correction is ~ -14 nats against H_eng: large enough to move anything from the candidate band of §4.4

to null-favored. The worked example of §5 follows the full-record rule, scoring six channels of which five are nulls.

One-sided channels have a capacity. Where a channel's data are a non-detection and the null predicts that non-detection with probability one, as the mid-infrared channel does today, the Bayes factor is bounded below by the dormancy prior: $\ln K_i = \ln[f_d + (1 - f_d)P(\text{quiet} \mid \text{active})] \geq \ln f_d$, with equality in the limit of infinite sensitivity. At the fiducial $f_d = 0.5$ the whole channel is worth 0.69 nats however deep the photometry goes, and it can never produce evidence in favour of H_{eng}. We report each such channel's capacity and spent fraction alongside its value, because the quantity a program should be ranked on is the headroom, not the potential.

4.3 Cross-messenger combination

Channels are combined multiplicatively where independent, $K_{\text{tot}} = \prod_i K_i$, with three corrections. Two are imported from gravitational-wave counterpart methodology (Ashton et al., 2018; Breschi et al., 2024). First, a *coincidence term*: hypotheses that predict correlated anomalies across channels (as H_{eng} does for MIR excess and low R during an accretion episode) earn a likelihood contribution from the observed coincidence structure itself, computed from the joint forward model rather than the product of marginals. Second, a *data-side independence audit*: channels sharing calibrators, atmospheric paths, or reduction pipelines are grouped and their shared systematics marginalized jointly before multiplication.

The third is hypothesis-side and is easy to miss, because it is a property of the hypothesis rather than of the instruments. A dormant installation is quiet in every electromagnetic channel at once, so the dormancy fraction f_d is a single nuisance parameter shared across channels, and the product rule marginalizes it once per channel instead of once in total. Any parameter shared this way must be marginalized jointly:

$$K_{\text{tot}} = \frac{f_d \prod_i P(d_i \mid \text{dormant}) \; + \; (1 - f_d) \displaystyle\int \prod_i P(d_i \mid \text{active}, \theta)\, \pi(\theta)\, d\theta}{\prod_i P(d_i \mid H_b)}. \tag{9}$$

The difference is not cosmetic. With n non-detection channels each returning $P(\text{quiet} \mid \text{active}) = 0.5$ and $f_d = 0.5$, Eq. (9) gives $\ln K = -0.47, -0.63, -0.68$ at $n = 2, 4, 6$, saturating at $\ln f_d = -0.69$ as the capacity argument of §4.2 requires, while the product of marginals gives $-0.58, -1.15, -1.73$ and diverges. The product rule would let a menu grow evidence against H_{eng} without bound by adding degenerate silent channels, which is the mirror image of the mis-scoring that §6 warns about for omitted hypotheses. The correction is inactive in the worked example below, where only one channel is live, and becomes necessary on the day a second one activates.

All three corrections are conservative in effect: the coincidence term rewards only pre-registered correlation patterns, the data-side audit only ever weakens evidence, and the joint marginalization only ever moves $|\ln K|$ down.

4.4 Decision thresholds

We adopt the logarithmic odds scale of Kass and Raftery (1995) with pre-registered action bands, stated here for $\ln K$ of H_{eng} over the best null:

Band	$\ln K$	Action
Null-favored	< 0	routine monitoring; H_{eng} parameter space shrinks
Uninformative	0 to 1	no action; report in campaign updates
Anomaly	1 to 3	targeted follow-up on the discriminating channels
Strong anomaly	3 to 5	independent-team replication; data release
Candidate	> 5 sustained across ≥ 2 messengers	community adjudication per post-detection protocols

These bands are on the Bayes factor, not on posterior odds, and the distinction carries a consequence we state rather than bury. Posterior odds are prior odds times K, and this paper sets no prior on H_{eng}: the epistemic-status box of §1 makes everything conditional on the optimization premise, and that premise is the quantity a prior would have to price. The bands are therefore resource-allocation thresholds, specifying what the campaign does next, and they carry no belief content on their own. A reader who wants a belief statement must supply prior odds; at a plausible 10^{-6} per target, $\ln K = 5$ still leaves posterior odds near 2×10^{-4}. The same gap governs class-level use: applied across the candidate hosts of Paper A's Table 1, prior odds must scale as $1/N$ or the "candidate" band will be crossed by chance somewhere in the sample. Bayes factors supply that multiplicity protection automatically once the prior odds are stated, which is a further reason to state them.

Two design points. The "candidate" band requires multi-messenger support by construction, honoring Paper C's two-messenger rule; no single channel, at any significance, can reach it, because single-channel anomalies are where every historical false alarm has lived (Sheikh et al., 2021). And the kill conditions of Paper A map onto the same scale in the other direction: the pre-registered falsifiers (e.g., LISA measurement of low spin; resolution of the mass tension in favor of H_{sub}) enter as ordinary likelihood terms and drive $\ln K$ negative, so confirmation and falsification run through one pipeline rather than two standards. The framework is deliberately instrument-agnostic: nothing in §4.2–§4.4 references ω Cen, and the machinery applies unchanged to any technosignature target with a defined null menu, complementing the qualitative Rio scale (Forgan et al., 2019) with a quantitative interior.

5 Worked example: Omega Centauri today

We now run the current ω Cen data through the framework. Inputs: the fast-star kinematics (Häberle et al., 2024b), the N-body modeling (González Prieto et al., 2025), the MSP timing bound (Bañares-Hernández et al., 2025) with the TRAPUM 2026 mass limit (Colom i Bernadich et al., 2026), the deep radio silence (Mahida et al., 2026; Tremou et al., 2018), and the JWST infrared limits (Chen et al., 2025). The computation implements the Appendix-B model (deliberately minimal: radius-marginalized hard-threshold MIR likelihood, log-flat priors, 2×10^6 prior draws); code in the paper repository, results in Figure 4.

- **H_q vs H_{sub}:** the interesting contest, and the framework's finding is that it is currently a stand-off dominated by the kinematics-versus-timing tension: the fast stars pull toward a point mass, the pulsars toward an extended one, and the combined $|\ln K|$ is $\lesssim 1$ with sign depending on how the two datasets' systematics are weighted. As an independent check, the `imbh-constraints` aggregation engine of the series' public toolchain, run on its nine-constraint curated compilation, returns an empty jointly-allowed mass window (formal tension verdict: the kinematic floor exceeds the tightest model-dependent ceiling), confirming that no point-mass value satisfies all published constraints at face value. The framework adds discipline, and no verdict, to a tension the community already knows; per series policy the tension is recorded, never collapsed.

- H_{eng} **vs best null:** $\ln K$ sits in the null-favored band, as it should. The only active channel today is waste heat, and, with the continuous per-filter likelihood of Appendix B marginalized over swarm radius and temperature split against the real JWST/MIRI sensitivity curve, it yields $\ln K = -0.47$ at fiducial priors (prior-sensitivity band -1.14 to -0.08 as the dormancy prior runs from 0.1 to 0.9): the electromagnetic silence sometimes informally cited as consistent with engineering is, in the accounting, mildly against H_{eng}, because the mid-infrared limits now exclude 76 per cent of the active-installation prior volume, up from 50 per cent under the retired hard-threshold model, while every other channel is inactive or degenerate and contributes nothing. The radius prior matters: narrowing its upper bound from the fiducial 4×10^3 to 10^3 AU, excluding the coolest outer configurations, raises detectability to $\ln K = -0.52$, confirming that cool, extended swarms are what preserve H_{eng}'s viability. The leak prior's lower bound is an engineering estimate (§3.1, Appendix A.3): repeating the computation with the prior extended down to $1 - f_{\text{sink}} = 10^{-6}$ gives $\ln K = -0.37$ instead of -0.47, so a softer transport bound softens the conclusion without reversing it. A second sensitivity is the fuel ceiling: the fiducial prior extends P_{comp} to the unsuppressed $7.6 \times 10^5\, L_\odot$, and truncating it at the ADIOS-suppressed natural ceiling of §2.7 ($\sim 10^3\, L_\odot$) removes most of the detectable prior volume, moving the mid-infrared $\ln K$ to -0.35 (Appendix B). A third, and now the largest of the three, is the lower edge of the same prior: P_{comp} is drawn from $1\, L_\odot$ upward, a boundary that coincides numerically with the warm instrument limit and has no physical warrant, and extending it to $10^{-3}\, L_\odot$ moves $\ln K$ to -0.32. Since $\ln K$ here is a ratio of prior volumes, every boundary is a lever and all are now tabulated in Appendix B rather than only the two that were varied in earlier drafts. Two things dominate the total, and both are structural rather than observational: the present evidence budget is set by the dormancy prior (the data enter only through the excluded active-prior fraction $\xi = 1 - P(\text{quiet} \mid \text{active}) = 0.755$, which is the part the data determine), and the channel is capped at $\ln f_d = -0.69$ in total, of which 68 per cent is spent. The framework's present value is the forecast, not the number. The surviving configurations are dormant or low-power, deep-envelope systems (Figure 2), and they are also post-spin-up. Paper A's Phase 4 drives the hole toward high spin by sustained near-Eddington feeding, which at the fiducial mass is $6.6 \times 10^8\, L_\odot$: four to five orders of magnitude above the concealment ceiling of §3.1, and not concealable by any architecture in the envelope. The episode's duration is set by the mass it must deliver, not by a round number: reaching Thorne's radiation-limited $a_\star = 0.998$ from $a_\star \approx 0$ requires $M_f/M_i \approx 2.20$ (Bardeen, 1970; Thorne, 1974), and at the Salpeter e-folding time $t_S \approx 4.5 \times 10^7$ yr for $\eta = 0.1$ accretion, that mass gain takes $\ln(2.20)\, t_S \approx 3.5 \times 10^7$ yr, roughly a third of the round duration used in earlier drafts. H_{eng} at ω Cen therefore requires that the spin-up episode has already ended, which is a statable constraint rather than a loose reading of "dormant", and it carries a second observable, secular core depletion, that §6 flags as testable against existing surface-density data. It also carries a mass prediction the spin-up requirement makes on its own: the $\approx 2.20\times$ growth factor applies on top of whatever mass the merger history alone delivers, so H_{eng} at ω Cen predicts $M \approx 2.20 \times M_{\text{merger}} \approx 2.20 \times 5 \times 10^4 \approx 1.1 \times 10^5\, M_\odot$ against the merger-only $M(H_{\text{q}}) \approx 5 \times 10^4\, M_\odot$ (González Prieto et al., 2025), a second, already-observable discriminant between H_{eng} and H_{q} that does not wait for LISA, and one that sits above the TRAPUM 2026 ceiling of $10^5\, M_\odot$ (Colom i Bernadich et al., 2026) rather than merely near it, so the late-spin-up branch predicts a mass nominally excluded by existing timing data if that ceiling holds. Per series policy we do not collapse the mass tension in either direction; we note only that the spin-up requirement is a prediction H_{eng} makes, not merely a constraint it must survive. Quantifying the core-depletion shrinkage is the point; the hypothesis pays for its unfalsified survival in measured parameter volume.

- **Information forecast:** channels are ranked by expected $|\Delta \ln K|$ per unit cost against their remaining headroom, not their nominal potential, since the mid-infrared channel has 0.41 nats left in total (§4.2) and further depth cannot buy more. The stated metric is an expectation, conditional value times probability of arrival, and the spin channel's two levels of that expectation diverge sharply enough that they need separate rankings rather than one merged list.

Target-level, ω Cen itself. The ω Cen-specific IMBH–compact-object capture rate from the González Prieto et al. (2025) models is $(4–8) \times 10^{-8}\ \mathrm{yr}^{-1}$, giving $\sim 3 \times 10^{-7}$ probability of ω Cen yielding its own inspiral over a four-year mission. Even against the large conditional value of a decisive spin measurement (a null at $a_\star \approx 0.06 \pm 0.02$ makes an engineered $a_\star \gtrsim 0.9$ a many-σ excursion), the expected $|\Delta \ln K|$ this channel delivers for ω Cen itself is $\sim 10^{-6}$ nats, negligible against the other target-level channels. The target-level leaders are therefore (1) resolution of the mass tension by continued MSP timing plus Gaia DR4 astrometry, which arbitrates H_q vs H_sub and thereby moves H_eng's denominator; (2) any detection of accretion at any level, which activates the R statistic, the only cheap channel with large discriminating power in both directions; (3) EMRI/IMRI environmental dephasing (§3.4), the only channel constraining engineered mass, which shares the spin channel's LISA-era timing and schedulability and, like R, moves evidence in both directions, since a vacuum inspiral counts against H_eng at the depths the hypothesis most values; spin itself ranks last at the target level, its 3×10^{-7} arrival probability attached rather than left implicit.

Class-level, the LISA IMBH population. Here spin is the leading channel, because the constraint accumulates across every ω Cen-like system LISA catches rather than waiting on this one target, and probability of arrival is ≈ 1 for the population even though it is $\sim 10^{-7}$ for any single target. The class-level version must carry the spin mixture, since comparable-mass growth histories elsewhere populate the $\chi \simeq 0.7$ attractor, and it inherits the isotropy caveat of §3.3 (Appendix B): the per-target inference only sharpens the population inference if the alignment fraction is itself characterized. The relic analysis of §3.3 raises the stakes on this channel: spin is the one fossil that survives abandonment, so the LISA-era population measurement adjudicates the engineered-history hypotheses for the whole Galactic IMBH population at once, and the per-target frameworks of this paper compose naturally into that population-level test.

Deeper radio silence, by contrast, now moves $\ln K$ weakly for every pairing: for this hypothesis menu the existing limits are already deep enough that further depth buys little discrimination.

6 Discussion

6.1 Limitations

The feasibility envelope depends on cluster-center parameters that carry real uncertainty (mass and number fraction of the remnant tail; the bound-cusp population inside the influence radius) and on a material-strength scale taken from present technology; both are varied in Appendix A. The largest lever is the remnant tail, and within it the perturber mass rather than the number fraction: for the unbound comparison bracket, the $10 \rightarrow 31\ M_\odot$ correction is worth a factor 8.0 in $\langle m_\star^2 \rangle$, against a factor ~ 10 across the whole 0.1–3 per cent fraction bracket; for the bound channel that now sets the operative floor, jointly anchoring f_BH and m_BH on the extended dark mass narrows the equivalent bracket to a factor 3.1 in kick variance (Appendix A). The cusp existence question, which earlier drafts carried as a second lever, is closed by the González Prieto et al. (2025) realizations, and the modeling gap it left behind, that the Monte Carlo modeled only the decoupled unbound flyby channel while the channel that survives, perturbation by bound cusp members, was treated analytically, is now closed: the bound-member diffusion channel and its

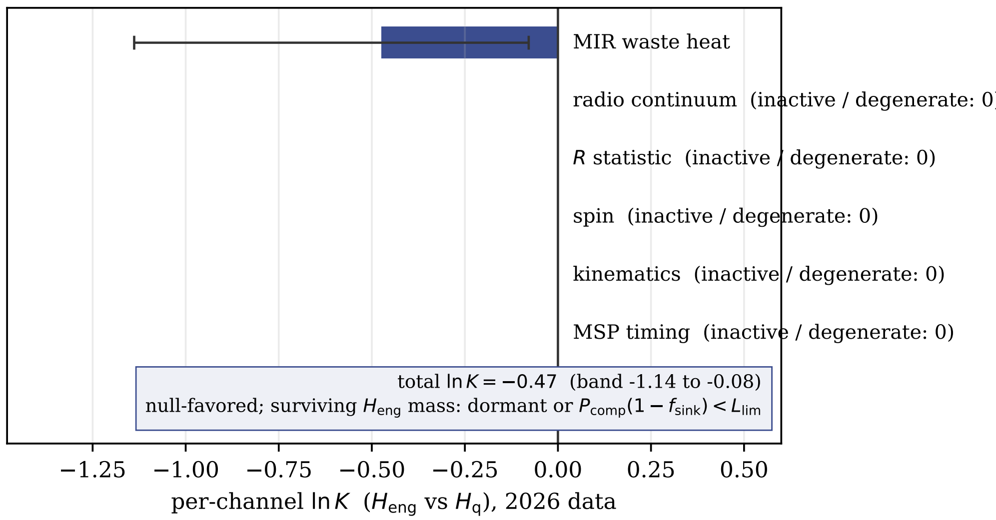

Figure 4: Per-channel $\ln K$ (H_{eng} vs H_{q}) on 2026 data, from the Appendix-B computation with the mid-infrared likelihood marginalized over swarm radius. Waste heat is the only live channel; its −0.474 (whiskers: dormancy-prior sensitivity; extending the leak prior down to 10^{-6} moves the central value to −0.371) is the entire present evidence budget. The empty bars are the framework's forecastable headroom: each activates with the observations ranked in the information forecast.

secular cadence are both Monte Carlo outputs (§2.4, Appendix A). What remains a modeling limitation rather than a parameter is the quasi-static cusp assumption, perturber positions are drawn once per history although the cusp itself relaxes on the same 0.03–0.6 Gyr segregation timescale that builds it, which biases toward longer lifetimes than a resampled population would give; the remaining parameters move the boundaries by factors of a few. The MAD-regulation channel inherits the systematic uncertainties of the GRMHD baseline literature, which is simulation-calibrated rather than observationally calibrated at IMBH masses. The adjudication framework shares the standard vulnerability of Bayesian model selection to unmodeled hypotheses: it ranks the menu it is given, and a surprise outside the menu (an astrophysical mechanism not yet imagined) would be mis-scored until added. The mitigation is procedural rather than mathematical: the menu is public, extensible, and versioned. That procedural claim is no longer unvalidated: Appendix E scores the machinery itself, not the ω Cen models, against three historical cases with an independently known resolution, including a positive control (the 1967 pulsar discovery) that the machinery registers at $\ln K = +7.9$ before correctly collapsing to −0.5 once the rotating-neutron-star hypothesis enters the menu. Two mundane cases (Project Hephaistos, Boyajian's star) are checked for false positives and pass; one data-driven collapse test falls short of its pass threshold as constructed, reported rather than tuned away, and detailed in Appendix E.

Two numbers used above deserve explicit qualification at the claim level. **The transport floor.** The geometric content of the bound $1 - f_{\mathrm{sink}} \gtrsim 10^{-4}$ is now a calculation (capture cone, beaming gain, étendue, and the plasma-cutoff carrier condition; Appendix A.3); what remains estimated is the re-interception term, the covering-fraction floor on power re-thermalized within the swarm, which a full radiative-transfer treatment of the swarm interior would replace. Its influence on the adjudication is bounded and signed: extending the leak prior down to 10^{-6} moves $\ln K$ from −0.474 to −0.371, well inside the dormancy-prior band of [−1.138, −0.079], so no conclusion of §5 turns on the exact value. **Core depletion is a weaker signature than the series has been claiming.** Paper A and Paper C both list secular core depletion as a Phase-4 observable, and neither had quantified it. Doing so changes its standing. Spinning the fiducial hole up to Thorne's radiation-limited $a_\star = 0.998$ (§5; $M_f/M_i \approx 2.20$) requires a mass gain of $\Delta M = 1.20 \times 2 \times 10^4 = 2.4 \times 10^4\, M_\odot$, and the delivered stellar mass is larger still, $(2.7–3.4) \times 10^4\, M_\odot$ once 10–30 per cent radiated and jetted loss is counted. Against the adopted core ($\rho_0 = 3 \times 10^3\, M_\odot\, \mathrm{pc}^{-3}$ inside $r_c = 3.6$ pc, total cluster mass $3.55 \times 10^6\, M_\odot$), that is 191 per cent of the mass now inside 1 pc, 24 per cent inside 2 pc, 4.1 per cent inside the core radius, and 1.0 per cent inside 10 pc. The fractional deficit is therefore set entirely by the delivery radius, and Paper A's own fuel census puts it at ~ 10 pc, where a sub-per-cent deficit sits below what a cluster surface-density profile can isolate from mass-function and distance systematics. Only fuel drawn from within ~ 2 pc would leave a deficit large enough to see, and that reservoir does not hold enough stars to supply the phase.

One further point closes the channel while a second reopens it as a discriminant rather than a dead end. The removed stellar mass is not lost from the cluster: it becomes central point mass, so the *total* enclosed-mass profile outside the feeding region is preserved and only the luminous-to-dark split changes, which is the quantity the mass-tension programme of §5 already measures. The endpoint is not degenerate with ordinary growth: $2 \times 10^4 \times 2.2 \approx 4.4 \times 10^4\, M_\odot$ is below the González Prieto et al. (2025) merger-only mass of $\sim 5 \times 10^4\, M_\odot$, so a spin-up episode starting from this fiducial hole does not reproduce the mass the models reach without one, and the final mass carries information about which history produced it. Applied instead to the merger-only endpoint itself as the pre-spin-up mass, the same $\approx 2.20\times$ factor gives the $M \approx 1.1 \times 10^5\, M_\odot$ prediction of §5; the two calculations bracket the endpoint mass depending on when in the growth history the spin-up episode runs, and both sit above the TRAPUM 2026 timing ceiling's neighborhood, which is why §5 reads the mass channel as a live discriminant. What remains of the depletion signature is a lookback limit worth stating on its own: a depleted region refills on the local relaxation time, ~ 1 Gyr in this core, and the episode itself lasts only a few per cent of that (3.5×10^7 yr; §5), so any depletion signature is erased within about a gigayear of the episode ending. Combined with the post-spin-up requirement of §5, the observable and the requirement are in tension: a system old enough to be quiet is old enough to have erased its own depletion.

The drainage timescale. The $\sim 10^3$ yr figure of §3.3 is an estimate. The asymmetry it rests on is robust, since relic dust persists around a star while debris around a hole is fed toward the loss cone and removed, but the rate is not. Closing that gap requires a loss-cone refill calculation for the debris population, tracking collisional periapsis diffusion and the coupling of the finest grindings to the ambient flow; we state the gap rather than supply a number that calculation has not yet produced.

6.2 Relation to the series and beyond

Within the series, this paper closes the constructive quadrant: A argued the destination is attractive, and §2 finds the destination is habitable by infrastructure, with the thermodynamically preferred region also the dynamically safest. It also arms Paper C's campaign with the scoring machinery its two-messenger policy presupposed, and returns to Paper D a refined survival input (p_s now decomposable into transit and residence terms, the latter bounded here). One dependency flag on that export: the residence-survival figure Paper D imports (≈ 0.8 over 10^8 yr) was derived from the impulsive diffusion floor. It is now supported by the adiabatically corrected curve, whose envelope minimum is 3×10^8 yr and whose interior values sit at the cluster-age cap (§2.4), so the figure stands with margin rather than on a bracket; Paper D should cite it against the corrected curve and note that the impulsive floor is the conservative alternative. Beyond the series, the two exportable products are the envelope method, applicable to any proposed megastructure environment with a stated perturber population, and the adjudication framework, applicable to any technosignature program willing to pre-register its nulls.

6.3 Conclusion

The engineered-IMBH hypothesis survives its first engineering audit and its first quantitative adjudication, in both cases by narrowing: infrastructure persists only within $\sim 4 \times 10^3$ AU of the hole, and the surviving parameter space after current data is dormant, low-power, and deep. Every forthcoming measurement listed in §5 shrinks it further or breaks it open. Either outcome is progress that a hypothesis without an envelope, a residue model, and a scoring rule could not deliver.

A Flyby Monte Carlo: method and parameters

Method (implemented in `figs/fig1_envelope.py`; fixed seed 20260717, offset per configuration so that no curve depends on the order in which the configurations are run). For each of 40 logarithmic radius bins $a \in [6\,r_g, 4 \times 10^3]$ AU (the grid begins at the Schwarzschild ISCO; earlier drafts began at $5\,r_g$, inside it, where circular orbits do not exist and the Newtonian orbital speed reaches $0.44c$): (1) 2×10^5 encounters are drawn with impact parameter from the gravitationally focused cumulative distribution $Q(b) \propto v_{\rm rel}^2 b^2 + 2GMb$ truncated at $b_{\max} = 30a$ (contributions beyond carry $< 10^{-3}$ of the kick variance), relative speed drawn as $v = 0.3\sigma + x$ with x a Rayleigh variate of scale $\sigma = 21$ km s^{-1}, i.e., a Rayleigh distribution shifted upward by 0.3σ (mean $1.25\sigma + 0.3\sigma = 1.55\sigma$; the JSON output of Appendix D records the exact parameterization), below the Maxwellian relative-speed mean of $\sim 2.26\sigma$; since slower encounters deliver larger kicks, the choice is conservative for swarm survival, and mass from the perturber mass function. The baseline function has two stellar components ($0.35\,M_\odot$ with weight 0.7; $0.6\,M_\odot$ white-dwarf tail with weight 0.3); the fiducial function adds segregated remnants, $1.4\,M_\odot$ neutron stars at number fraction 0.02 and $31\,M_\odot$ black holes (§2.4) at number fraction $f_{\rm BH} \in \{0.001, 0.01, 0.03\}$ with the stellar weights renormalized, at fixed total number density; a $10\,M_\odot$ configuration is retained in the output so that the change of perturber mass is auditable. (2) Each encounter contributes an impulsive eccentricity kick $\delta = 2(m_\star/M)\min[1, (a/b)^2](v_{\rm orb}/v)$, i.e., $\delta = \delta v_{\rm tid}/(2v_{\rm orb})$ with the $O(1)$ geometric coefficient set to unity; that choice sits on the optimistic side by a factor of ~ 2 in the δ^2 timescale, offset by the conservative velocity sampling above and by the neglect of adiabatic suppression, so the net sits within the stated factor-few accounting. The baseline curves apply no adiabatic suppression; a parallel run applies the Gnedin and Ostriker (1999) kernel $A(x) = (1 + x^2)^{-5/2}$ to δ^2 per encounter, with $x = (b/a)(v_{\rm orb}/v)$, and is the physical estimate (§2.4). We use the power-law kernel rather than the $\exp(-x^2)$ form of Spitzer (1987) because the latter over-suppresses relative to N-body results; the power law is the conservative choice. (3) Per history (2×10^4 per bin), survival time is the random-walk first-passage to $e = 0.5$: encounters to threshold $n_\star = e^2/\langle\delta^2\rangle$ with CLT scatter, divided by the total encounter rate $\Gamma(< b_{\max})$, capped at 12 Gyr. The bookkeeping is thus: the 2×10^5 encounter draws per bin characterize the per-encounter kick-variance statistics; each of the 2×10^4 histories then draws its encounter count analytically from that first-passage distribution rather than re-simulating individual encounters, so a history at the median spans $\sim n_\star \sim 10^4$–10^6 encounters depending on bin. Single-passage disruption is impossible over the part of the grid where the impulsive kernel is meaningful: at $a \gtrsim 1$ AU the saturated kick is $\delta \lesssim 5 \times 10^{-3}$ for stellar perturbers and $\lesssim 1.4 \times 10^{-1}$ for the rare black-hole perturbers. Inside ~ 1 AU the saturated form returns $\delta > 1$ for a heavy penetrating passage ($\delta = 1.4$ at 10^{-2} AU), and the first-passage step count $n_\star = e^2/\langle\delta^2\rangle$ falls below 10^2; neither the small-kick premise nor the random-walk premise holds there, which is a further reason the impulsive curves are a floor rather than an estimate in the deep envelope, and is why the adiabatic run is carried alongside them. Direct star–element scattering requires approach within $\sim 2Gm_\star/(v_{\rm rel}v_{\rm orb})$ and contributes negligibly at any realistic covering fraction. This unbound-field calculation uses a core-average perturber density (10^4 pc^{-3}) with gravitational focusing capturing only the unbound flux; it is the curve retained in Figure 1(a) as the measure of that channel's adiabatic decoupling, not as an operative floor.

Bound-cusp channel (panel b). Implemented in the same script; species density profiles and joint anchoring in `common.py`. The relaxed Bahcall–Wolf cusp that the González Prieto et al. (2025) realizations form raises the local density of *bound* perturbers, four species (main-sequence and white-dwarf stars at $\gamma_\star = 1.3$; neutron stars, $1.4\,M_\odot$, interpolated $\gamma_{\rm NS} = 1.5$; black holes, $\gamma_{\rm BH} = 2.0$), each following $N_s(<r) = [4\pi n_{s,0} r_{\rm infl}^3/(3-\gamma_s)](r/r_{\rm infl})^{3-\gamma_s}$ normalized against the core-average density $n_{s,0} = f_s N_{\star,\rm core}$ at $r_{\rm infl}$. The black-hole number fraction is anchored jointly rather than assembled from separate sources: $f_{\rm BH,global} = (M_{\rm dark}/m_{\rm BH})/N_{\star,\rm cluster}$ with $N_{\star,\rm cluster} = 7\times10^6$, and $f_{\rm BH,central} = 10\times f_{\rm BH,global}$ with the $\times10$ segregation enhancement stipulated from the independent-anchor estimate of the referee round that specified this model, not derived here. At the fiducial $M_{\rm dark} = 2.5\times10^5\,M_\odot$, $m_{\rm BH} = 31\,M_\odot$: $f_{\rm BH,central} = 1.15$ per cent. Because kick variance scales as $n_{\rm BH}\langle m_{\rm BH}^2\rangle \propto M_{\rm dark} m_{\rm BH}$, a lighter perturber is linearly, not quadratically, less damaging at fixed dark mass; the bracket $\{M_{\rm dark} = 2.0\times10^5\,M_\odot, m_{\rm BH} = 15\,M_\odot\}$ to $\{3.0\times10^5\,M_\odot, 31\,M_\odot\}$ spans a factor 3.1 in kick variance, against the factor 30 the unbound comparison bracket spans. Since $\gamma_{\rm BH} = 2$ gives a shell count linear in radius, the number of bound black holes inside the stripping radius has expectation $\lesssim 1$ everywhere in the envelope; each history draws its perturber count per species as an independent Poisson variate from the expected shell count $\lambda_s(a) = N_s(<\sqrt{2}\,a) - N_s(<a/\sqrt{2})$ rather than using a smooth rate, which is what makes the resulting lifetime distribution bimodal rather than unimodal. Relative velocities are drawn as Maxwellian with the local Keplerian dispersion $s_{{\rm rel},s} = \sqrt{2GM/[(1+\gamma_s)a]}$ (species-dependent, since $x\sim1$ for bound members and σ does not apply); impact parameters are uniform in b^2 out to $b_{\max} = 3a$, set by where the adiabatic kernel's contribution to the kick variance drops below 1 per cent; gravitational focusing is dropped (it uses the perturber's own mass for an already-orbiting body and is a $\sim10^{-3}$ correction at $b = a$). The same Gnedin–Ostriker kernel applies, since $x\sim1$ rather than $\ll1$ for these perturbers; an un-kerneled run is retained in the JSON (`bound_fid_nokernel`) as a diagnostic only. Survival is the same random-walk first-passage criterion as the unbound channel, with the total rate summed over the four species' Poisson-drawn counts for that history.

Secular cadence. $t_{\rm sec}(a) = (M/m_{\rm pert})(r_p/a)^3 P_{\rm orb}(a)$ for the outermost bound black hole, fixed at $r_p = 4\times10^3$ AU (where $N_{\rm BH}(<r) = 1$); quenched by relativistic apsidal precession, $t_{\rm GR}(a) = [a/(3r_g)]P_{\rm orb}(a)$, and by mass precession from the light cusp population enclosed within a (drawn per history from the same $N_s(<a)$ profile, excluding the perturbing black hole), $t_{\rm mass}(a) = [M/M_{\rm enc}(<a)]P_{\rm orb}(a)$; active only where $t_{\rm sec} < \min(t_{\rm GR}, t_{\rm mass})$ and within the isotropic-cusp resonance window, probability 0.775 for mutual inclination in $(39.2°, 140.8°)$.

Results. For the stars-plus-WDs baseline, median impulsive lifetimes are 3.4×10^9, 3.3×10^9, and 2.5×10^9 yr at $a = 10, 10^2, 10^3$ AU: flat, by the scale cancellation $\langle\delta^2\rangle \propto a^{-1}$, $\Gamma\propto a$ discussed in §2.4. Adding the remnant components lowers the floor in proportion to the added $\langle m_\star^2\rangle$: medians at the same radii are $(1.0, 0.93, 0.58)\times10^9$ yr at $f_{\rm BH} = 0.001$, $(6.6, 6.3, 5.2)\times10^7$ yr at the fiducial $f_{\rm BH} = 0.01$, and $(2.2, 2.0, 1.9)\times10^7$ yr at $f_{\rm BH} = 0.03$, with envelope minima of 1.0×10^8, 2.7×10^7, and 1.1×10^7 yr respectively. The perturber mass dominates the fraction bracket: the same fiducial run at the earlier $10\,M_\odot$ assumption gives $(5.4, 5.1, 3.5)\times10^8$ yr with an envelope minimum of 2.0×10^8 yr, so the $10\to31\,M_\odot$ substitution is worth a factor 8.0 in $\langle m_\star^2\rangle$ and displaces the floor further than the whole 0.1–3 per cent bracket does at fixed mass. Varying the mean stellar mass by $\times2$, the density by $\times3$, and $e_{\rm cross}$ over 0.3–0.7 moves the floor by additional factors of a few. The quoted minima are min-of-noisy-bins statistics; the script reports the spread over five seeds directly, giving ±20–40 per cent (for example the $f_{\rm BH} = 0.01$ minimum spans 2.8–3.7×10^7 yr), so no hard floor is claimed at better than a factor of two.

The adiabatically corrected run, which is the physical estimate for the unbound channel, sits far above all of these: at the fiducial mass function its median reaches the 12-Gyr cap for $a\lesssim3$ AU and returns

7.3×10^9, 4.5×10^9, and 1.7×10^9 yr at $a = 10$, 10^2, 10^3 AU, with an envelope minimum of 3.0×10^8 yr at the outer edge where the suppression is weakest. The adiabatic parameter at $b = a$ runs from 4.1×10^3 at the ISCO to 2.2 at the stripping radius. The operational reading is that the unbound-flyby channel is removed from the problem across the envelope, and that any process that does bind must involve perturbers bound to the hole, which panel (b) now supplies directly rather than as an analytic estimate.

Bound-cusp results. At the envelope edge ($a = 4 \times 10^3$ AU), a bound black hole is drawn in 53 per cent of 2×10^4 histories ($p_{\mathrm{BH,present}}$; the analytic expectation from $1 - e^{-\lambda_{\mathrm{BH}}}$ at the joint-anchored fiducial). The diffusion floor conditional on presence is bimodal by construction, not merely broad: median 4.3×10^7 yr, 16th percentile 2.2×10^7 yr, 27 per cent of histories still reaching the 12-Gyr cap because no black hole (or a distant one) was drawn. Interior to the edge the same channel strengthens rapidly: at $a = 4 \times 10^2$ AU, $p_{\mathrm{BH,present}} = 7.2$ per cent and 93 per cent of histories survive to 10^8 yr, against 47 per cent at the edge; this is the trade the analytic single-perturber estimate could only gesture at, rare but severe inward, common but mild at the edge. The perturber-mass bracket of the joint anchoring moves the edge floor from 1.0×10^8 yr at $\{2.0 \times 10^5\, M_\odot, 15\, M_\odot\}$ to 4.3×10^7 yr at fiducial to 4.4×10^7 yr at $\{3.0 \times 10^5\, M_\odot, 31\, M_\odot\}$, the factor-3.1 variance bracket stated above; the un-kerneled diagnostic run drops the edge floor to 6.7×10^6 yr, confirming the adiabatic suppression matters by a factor of a few for the bound channel too, as expected since $x \sim 1$ rather than $\ll 1$ there. The secular clock reaches $t_{\mathrm{sec}} = t_{\mathrm{GR}} \simeq 3.7 \times 10^7$ yr at $a \simeq 4.0 \times 10^2$ AU, its GR-quenched inner edge, and $t_{\mathrm{sec}} \simeq 1.2 \times 10^7$ yr at $a \simeq 9 \times 10^2$ AU against $t_{\mathrm{GR}} \simeq 2.6 \times 10^8$ yr there, unquenched; it is drawn in Figure 1(b) only over 4×10^2–4×10^3 AU and is not a survival curve at any point on that range.

A.3: transport bound. The floor $1 - f_{\mathrm{sink}} \gtrsim 10^{-4}$ is a calculation in three steps: delivery geometry, carrier physics, and re-interception, with only the last carrying an estimated coefficient.

Delivery. The photon-capture cross-section of the hole is $\sigma = 27\pi r_{\mathrm{g}}^2$ (critical impact parameter $3\sqrt{3}\, r_{\mathrm{g}}$), so from a platform at $r = 10^2\, r_{\mathrm{g}}$ the disposal channel subtends $\Omega_c/4\pi = 6.75 \times 10^{-4}$: isotropic emission delivers less than a thousandth of its power to the horizon, and $f_{\mathrm{sink}} \to 1$ requires a beaming gain $\gtrsim 1.5 \times 10^3$. That gain is optically trivial. The capture cone has half-angle $\theta_c \simeq 0.052$ rad, so diffraction demands only an aperture $D \gtrsim 23\lambda$, and the étendue (concentration) limit $1/\sin^2\theta_c \approx 3.7 \times 10^2$ puts the optic area at a few hundred radiator areas. Delivery is not the binding term.

Carriers. The minimum-energy vacuum carrier, wavelength $\sim r_{\mathrm{g}}$, energy $hc/r_{\mathrm{g}} \simeq 7 \times 10^{-33}$ J, is a ~ 10 Hz wave and cannot propagate: it lies below the plasma frequency of any realistic environment (Paper A, this revision). Inside a fed envelope with $n_e \sim 10^4$–10^8 cm^{-3}, $\nu_p \approx 1$–90 MHz, so the minimum propagating carrier energy is $\gtrsim 10^{-27}$ J, three to six decades above the vacuum figure. The delivered per-bit disposal cost is therefore set by the plasma cutoff together with dense (multi-bit-per-photon) coding up to channel capacity, and matter carriers, cold mass dropped down the capture cone, are the fallback that evades the cutoff entirely. None of this changes the geometric floor below; it fixes the cost per bit, not the intercepted fraction.

Re-interception. Beamed flux traverses the swarm on its way to the cone and is intercepted with probability of order the swarm covering fraction seen from a typical element, $\gtrsim 10^{-4}$ for the covering fractions that make the swarm computationally worthwhile; the intercepted power is re-thermalized and radiated at swarm temperature. This is the binding term, and its coefficient is the one number a full radiative-transfer treatment of the swarm interior would refine (§6); the sensitivity analysis of Appendix B shows the adjudication does not turn on it.

B Spin channel: random-walk derivation and net-alignment sensitivity

Derivation. For a capture of mass $\mu \ll M$ onto a hole of spin parameter $a_\star = Jc/(GM^2)$, the change in $a_\star$ per event is $\delta a_\star = (\mu/M)[\tilde{L}\cos\iota - 2a_\star\tilde{E}]$, with $\tilde{L}$ the specific angular momentum deposited (units GM/c) and ι the inclination of the captured orbit to the existing spin axis. Under isotropic arrival, $\langle\cos\iota\rangle = 0$, the mean drift comes only from the mass term and $a_\star$ is set by the vector random walk of J against a denominator M^2 that grows with each capture. At $a_\star \ll 1$ the walk is free of the Hughes–Blandford damping term, which does not act until $a_\star$ approaches its ~ 0.3 equilibrium, not reached here. Each capture deposits $|\delta J| = \mu\tilde{L}(GM_i/c)$ in a random direction, with $M_i = M_0 + i\mu$; summing in quadrature over $N = (M_f - M_0)/\mu$ events,

$$\langle J^2\rangle = \left(\frac{\tilde{L}\mu G}{c}\right)^2 \sum_i M_i^2 \approx \left(\frac{\tilde{L}\mu G}{c}\right)^2 \frac{M_f^3}{3\mu} \qquad (M_f \gg M_0), \tag{10}$$

so that

$$a_{\star,\mathrm{rms}} = \frac{J_{\mathrm{rms}}c}{GM_f^2} = \frac{\tilde{L}}{\sqrt{3}}\sqrt{\frac{\mu}{M_f}}, \tag{11}$$

which for a quasi-circular capture from the Schwarzschild ISCO ($\tilde{L} = 2\sqrt{3}$) reduces to $a_{\star,\mathrm{rms}} = 2\sqrt{\mu/M_f}$. The M_i^2 weighting means the walk is controlled by the final mass: the last captures are both the largest angular-momentum depositors and the ones normalized against the largest M^2.

Validation against Mandel (2008). That result is the isotropic random walk for a hole gaining half its mass from neutron-star captures, $M_f = 2M_0$, for which $\sum M_i^2 = 0.2917\, M_f^3/\mu$ and $a_{\star,\mathrm{rms}} = 1.871\sqrt{\mu/M_f}$. Setting $a_\star = 0.2$ with $\mu = 1.4\, M_\odot$ gives $M_f = 123\, M_\odot$, matching Mandel's stated IMRI setup. The spread checks too: a 3-D random walk gives a Maxwellian $|a_\star|$ distribution with $\sigma/\mathrm{mean} = \sqrt{3 - 8/\pi}/(2\sqrt{2/\pi}) = 0.42$, against Mandel's quoted $0.08/0.2 = 0.40$. The formula reproduces both the central value and the dispersion of the one quantitative anchor the literature supplies, and we adopt it rather than importing Mandel's number directly, since his M_f/M_0 and mass ratio differ from ω Cen's by more than an order of magnitude each.

Applied to ω Cen. With $\mu = 30$–$40\, M_\odot$ (González Prieto et al., 2025), $M_f = 2\times10^4$ (this paper's fiducial mass) to $5\times10^4\, M_\odot$ (the González Prieto et al. (2025) endpoint), and $\tilde{L} = 2\sqrt{3}$ (circularized capture) to 4 (direct plunge on a marginally bound orbit):

Corner	μ ($M_\odot$)	M_f ($M_\odot$)	$\tilde{L}$	$a_{\star,\mathrm{rms}}$
Minimum	30	5×10^4	$2\sqrt{3}$	0.049
Central	35	5×10^4	$2\sqrt{3}$	0.053
This paper's fiducial mass	35	2×10^4	$2\sqrt{3}$	0.084
Maximum	40	2×10^4	4	0.103

Converting the central rms value to the Maxwellian mean and dispersion gives $a_\star = 0.049 \pm 0.021$ (90th percentile 0.077); the full derived range across all corners is $a_\star \sim 0.05$–0.10, quoted in the main text as central ≈ 0.06. This replaces the imported literature value $a_\star \sim 0.1$–0.3 (Mandel et al., 2008), which is the outcome of a different growth factor and mass ratio never derived for this target.

The mass-ratio boundary. No numeric cutoff for the random-walk regime exists in Hughes and Blandford (2003) or its successors; the boundary is not sharp because the underlying physics is continuous in mass ratio. Setting $a_{\star,\mathrm{rms}} = 0.7$, the numerical-relativity attractor value, in the scaling above with $q \equiv \mu/M_f$ gives $2\sqrt{q} = 0.7$, so $q \approx 0.12$: the random walk alone reaches attractor-level spin at $q \approx 0.12$

and produces less below that, as $\sqrt{q}$. This is consistent with the only quantitative statements in the literature, Fishbach et al. (2017)'s tested $q \gtrsim 0.7$ range and $q < 0.32$ anti-aligned bound. ω Cen sits at $q = \mu/M_f \approx 7 \times 10^{-4}$, a factor ~ 175 below the boundary.

Net-alignment sensitivity. Relaxing $\langle \cos \iota \rangle = 0$ to a net alignment fraction $\epsilon_{\rm rot} = \langle \cos \iota \rangle$, the coherent deposit is $J_{\rm coh} = \epsilon_{\rm rot}(\tilde{L}\mu G/c) \sum_i M_i$, with $\sum_i M_i \approx M_f^2/(2\mu)$, giving

$$a_\star^{\rm coh} = \epsilon_{\rm rot} \frac{\tilde{L}}{2} = (1.73\text{–}2.0)\, \epsilon_{\rm rot} \tag{12}$$

for $\tilde{L} = 2\sqrt{3}$ to 4. The mass dependence cancels: unlike the random walk, the coherent channel carries no $\sqrt{N}$ suppression, so it dominates as soon as it exists. It equals the random walk's 0.053 at $\epsilon_{\rm rot} \approx 0.035$, reaches $3\times$ the random walk at $\epsilon_{\rm rot} = 0.10$, sits inside the imported 0.1–0.3 band at $\epsilon_{\rm rot} = 0.20$, and reproduces the attractor value by this separate route at $\epsilon_{\rm rot} \approx 0.40$. Above $\epsilon_{\rm rot} \approx 0.5$ the low-spin prediction is gone. ω Cen is a rotating cluster and the inspiraling population is mass-segregated by dynamical friction, a process with no obvious reason to randomize the angular momentum a captured object arrived with, so isotropy is a premise to check against the González Prieto et al. (2025) realizations (which can supply $\langle \cos \iota \rangle$ for the inspiraling set directly) rather than a default to assume.

C Priors, likelihoods, and computation

Implemented in `figs/fig3_lnk.py` (2×10^6 prior draws; fixed seed). $H_{\rm eng}$ parameters: $P_{\rm comp}$ log-uniform on $[1, P_{\rm fuel}]\, L_\odot$ with $P_{\rm fuel} = 7.6\times10^5\, L_\odot$ (the ambient Bondi ceiling of §2.7, quoted as 8×10^5 in the text). That ceiling is the engineered-unsuppressed, near-extremal-spin value and is not cosmetic: it falls $\sim 5\times$ at $a_\star = 0.5$ through the $\eta(a_\star)$ dependence of §2.7, and $\sim 10^2$–$500\times$ under natural outflow suppression, so a fully stated prior is a mixture over ($a_\star$, suppressed/unsuppressed). We instead quote the sensitivity band, since the arithmetic is analytic in $\ln K = \ln[f_d + (1 - f_d)\, P(\text{quiet} \,|\, \text{active})]$: at the fiducial unsuppressed ceiling, $P(\text{quiet} \,|\, \text{active}) = 0.245$ and $\ln K_{\rm MIR} = -0.47$; truncating the ceiling at the ADIOS-suppressed $\sim 10^3\, L_\odot$ removes most of the detectable prior volume, raising $P(\text{quiet} \,|\, \text{active})$ to 0.405 and moving $\ln K_{\rm MIR}$ to -0.35; the $a_\star = 0.5$ ceiling sits between. Other parameters: $1 - f_{\rm sink}$ log-uniform on $[10^{-4}, 1]$; swarm radius r log-uniform over the envelope $[2\times10^{-2}, 4\times10^3]$ AU (an uninformative choice within the allowed region; the thermodynamic gradient of Paper A would weight it deeper, which is the conservative direction for $H_{\rm eng}$ since deep swarms are warm and more detectable); dormancy probability $f_d = 0.5$, varied over $[0.1, 0.9]$ for the sensitivity band. MIR channel likelihood (v1.0, continuous, OCS-MIRI-1): the waste luminosity $L_{\rm waste} = (1 - f_{\rm sink})P_{\rm comp}$ splits into a warm component $f_{\rm warm}L_{\rm waste}$ and a cool component $(1 - f_{\rm warm})L_{\rm waste}$, each re-radiated as a blackbody from the same swarm radius r via $T_{\rm eff}(L, r)$, giving a two-temperature SED with $f_{\rm warm}$ log-uniform on $[10^{-3}, 0.999]$ (the one new free parameter this build introduces) as the only lever beyond $(P_{\rm comp}, f_{\rm sink}, r)$. Each component's flux density at every MIRI imaging filter's pivot wavelength (F560W through F2550W, the 25.5 μm imager cutoff, Dicken et al. 2024) is compared against the JDox ETC-6.0 point-source sensitivity (STScI, 2026), penalized by the core crowding factor of $3\times$ that Paper C already applies; a configuration is detected iff the summed flux exceeds the crowded limit in *any* filter, combined as a logical OR, never a sum, over the nine filters. This replaces the piecewise step in $T_{\rm eff}$ and retracts the "Wien-peak band" and "cross-band leakage" language of earlier drafts, which the shipped script never implemented (R5-V7). Filter bandpass width is not sourced and is not folded in; each filter is treated as a delta function at its pivot wavelength, a stated approximation. The crowding factor is inherited from Paper C's NIRCam-calibrated recovery and has not been verified for the MIRI point-spread function, but the adjudication does not turn on it: removing the penalty altogether ($3\times \to 1\times$) moves $\ln K_{\rm MIR}$ by 0.014 nats, an order of magnitude below every prior boundary in the sensitivity table below. The far-infrared corner (swarm cold enough that its Wien tail is negligible at every MIRI wavelength) is not clamped by an assumed placeholder limit, as earlier drafts did; it falls out of the physics as simply undetected by any filter in the array, which is the honest state, since no far-infrared instrument is currently operating (OCS-MIRI-DATA). Data = non-detection in all nine filters; $P(\text{quiet}\,|\,\text{active}) = 0.245$, giving $\ln K_{\text{MIR}} = \ln[f_d + (1 - f_d) \times 0.245] = -0.47$ at fiducial f_d, band $[-1.14, -0.08]$. A logistic softening of the per-filter threshold, width 0.5 dex in $\log_{10}(F_\nu/F_{\text{lim}})$, combined across filters as $P(\text{nondetect}) = \prod_i (1 - p_{\text{det},i})$, moves $\ln K$ by -0.012 nats, so the threshold shape remains immaterial.

Two statistics are reported alongside that number because it is not, on its own, a measure of what the data said. The first is the data-only quantity $\xi \equiv 1 - P(\text{quiet}\,|\,\text{active}) = 0.755$, the fraction of the *active*-installation prior volume the current limits exclude; no dormancy prior enters it. The second is the channel capacity: because $P(\text{no detection}\,|\,H_q) = 1$, this channel is one-sided and bounded by $\ln K \geq \ln f_d = -0.69$ (§4.2), of which 68 per cent is now spent, against 41 per cent under the hard-threshold model this build replaces. Further mid-infrared depth can buy at most 0.22 nats more, and can never buy evidence for H_{eng}.

Prior-edge sensitivity. $\ln K$ here is a ratio of prior volumes, so every prior boundary is a lever and all of them are reported rather than the two that earlier drafts varied. Each row re-runs the same computation with one boundary moved; each is a command-line switch on the script.

Boundary moved	Value	ξ	$\ln K_{\text{MIR}}$	Switch
— (fiducial)	—	0.755	−0.474	—
Transport floor $1 - f_{\text{sink}}$	10^{-6}	0.620	−0.371	`--leak-floor-dex -6`
P_{comp} floor	$10^{-3}\ L_\odot$	0.549	−0.321	`--pcomp-floor-dex -3`
P_{comp} ceiling (ADIOS)	$10^{3}\ L_\odot$	0.595	−0.353	`--pcomp-ceiling 1e3`
Radius prior	$[0.02, 10^3]$ AU	0.812	−0.521	`--r-range 0.02 1000`
Temperature split floor	$f_{\text{warm}} \geq 10^{-1}$	0.747	−0.468	`--split-floor-dex -1`
Threshold shape	logistic, 0.5 dex	0.770	−0.486	`--soft-threshold 0.5`

The largest lever is the one earlier drafts did not vary. The P_{comp} prior runs upward from $1\ L_\odot$, a boundary that coincides numerically with the warm instrument limit and has no physical warrant; extending it three decades downward is worth +0.153 nats, more than the transport-floor sensitivity (+0.103) and now the fuel-ceiling one as well (+0.121). The temperature-split floor is the weakest lever measured, +0.006 nats, because the total waste luminosity and the swarm radius, not its partition between the two components, set most of the detectability. Stated plainly: the evidence budget is still dominated by the dormancy prior, the data entering only through ξ, but the continuous per-filter likelihood raises ξ from 0.50 to 0.76 against the same 2026 non-detections, because the real JWST/MIRI depths (sub-μJy at the warm filters) exclude far more of the active-installation prior volume than the old order-of-magnitude piecewise limits did; this moves $\ln K_{\text{MIR}}$ from −0.29 to −0.47 and the spent capacity fraction from 41 to 68 per cent, without touching the capacity bound $\ln f_d$ itself. The same forward model, a two-temperature swarm SED integrated against the real per-filter MIRI sensitivities, supplies the point-source SED discriminant Paper C's mid-infrared program needs (Program 1, Swanson 2026d): the warm/cool flux ratio across F560W through F2550W is the model's falsifiable prediction for any point-source excess the program detects. Radio, R, spin, kinematic, and timing channels are inactive or degenerate between H_{eng} and H_q on current data and contribute $\ln K = 0$ each, as Table 1 requires. The H_q–H_{sub} contest is checked against the `imbh-constraints` v1.0.0 aggregation engine (nine curated constraints; empty joint window, tension verdict), whose mass-window math is anchored by the series' replayed golden and CI parity gate.

D Reproducibility

Figure scripts and outputs live beside this manuscript (`paper/figs/`: `common.py` for shared constants, one script per figure, `fig1_results.json` for machine-readable Monte Carlo output). The measurement compilation and aggregation library are the public series toolchain (`imbh-constraints` v1.0.0, Zenodo DOI 10.5281/zenodo.20689279, ASCL submitted); interactive calculator versions of the Figure 1 and §2.7 computations are planned for omegacentauri.me. All random draws use fixed seeds; all quoted numbers regenerate from the scripts.

E E6 validation of the adjudication machinery

§4's machinery is tested here against three historical cases where the resolution is independently known, per the protocol committed in advance (OCS-E6-PROTOCOL) and implemented in `figs/appendix_d_validation.py`. The object under test is the scoring machinery itself, the explicit-menu requirement, the per-channel Bayes factor, the priced catch-all null, and the action bands, never the ω Cen forward models; what varies per case is the channel likelihood the machinery is fed. LGM-1, the discovery of the first pulsar, is the positive control: the only case that can drive $\ln K$ strongly positive, and the only one that carries a menu-completion test. Project Hephaistos and Boyajian's star are the mundane cases, scored for false positives against a menu that already contained the resolving alternative before the resolving datum arrived. 'Oumuamua is excluded per the protocol: its lightsail proposal has no settled community verdict to calibrate against. Priors and menus for each case were sealed in a dedicated commit containing no scoring code before any $\ln K$ was computed (`paper/e6-priors/priors.json`, commit `f82e11d`); the scoring commit (`96e0d1e`) follows it. Every channel construction is stated in that file and in the script's comments; several are explicit modeling choices rather than re-measurements of the primary data, flagged as such below.

LGM-1. At Epoch 1a (late November 1967, after the fast-recorder capture, before the orbital-Doppler test), the menu is terrestrial interference, artificial signal, and a natural-source catch-all. The catch-all is priced against the fastest coherent celestial periodicity catalogued at that date, of order one hour; CP 1919's coherent 1.337-s period sits 3.4 dex beyond that boundary, and at the fiducial inflation $\lambda = 1$ this channel alone gives $\ln K = +7.90$ against the natural catch-all, comfortably inside the candidate band and satisfying P1 ($\ln K \geq 3$). The framework tolerates catch-all generosity up to $\lambda = 2.63$ before P1 fails; the $(\pi_{\rm nat}, \lambda)$ grid confirms $\ln K$ is invariant in $\pi_{\rm nat}$, as a likelihood ratio must be. Two collapses follow, reported separately per protocol. Collapse B, menu completion: adding the rotating-neutron-star hypothesis to the Epoch-1a menu, priced as though it had existed in November 1967, drives $\ln K$ to -0.50, a 8.40-nat drop that satisfies P2. Collapse A, data with the menu held fixed: adding the null orbital-Doppler test and the discovery of a second independent pulsating source moves $\ln K$ to 6.69 and 4.04 respectively, 2.83 with both; the multiplicity channel is the larger mechanism (-3.86 nats against -1.20 for the Doppler null). This is a finding, not a pass: P3 requires $\ln K$ below 1 after Collapse A and the constructed channels do not reach it, so F3 is reported rather than repaired by re-tuning a prior after seeing the number. The achievable range at Epoch 1a, read off the λ sweep bounds, is $[2.63, 8.00]$, 5.37 nats, satisfying P6.

Project Hephaistos. The dossier's open question, whether Suazo et al. (2024) already discussed background-galaxy contamination at publication or whether Ren et al. (2024) introduced it three weeks later, is resolved by reading the discovery paper's own full text (§3.1, §4): the paper estimates ~ 2 contaminants expected among the seven candidates from Hot DOG chance alignment and states plainly that "the possibility of perfect alignments cannot be ruled out." Background contamination is therefore on the Epoch-1 (2024-05-06) menu by construction, not a later addition, which sets every subsequent trajectory. Scored per candidate for the three with a direct-imaging resolution: G reaches $\ln K = +0.92$ (uninformative) at Epoch 1 and collapses to -3.08 once the off-star VLBI radio position resolves it (2024-05-23), never crossing the strong-anomaly band before that datum existed. D and E follow the same Epoch-1 read, rise to $+2.42$ (anomaly, not strong anomaly) once the 2026-07-03 archival diagnostics find no contamination signature for them specifically, in contrast to B and C, and collapse to -3.58 once JWST/MIRI resolves both as background galaxies (2026-07-10). All three satisfy P4: no case reaches the strong-anomaly band ($\ln K \geq 3$) before its resolving datum is public, and the anomaly-band excursion for D and E is exactly what that band exists to register.

Boyajian's star. Scope is the short-timescale dipping only, per the dossier's cut; the secular dimming and the contested plate-fading claim are excluded and named as such. The megastructure hypothesis is not on the 2015-09-11 discovery-paper menu; it enters 2015-10-15 (Wright et al.), so the wavelength-dependence channel is inactive at both early epochs and $\ln K = 0$ there, a null result correctly attributed to the channel not yet existing rather than to evidence. The continuous form required by §B is implemented physically rather than as a hard threshold: an anomalous-diffraction extinction efficiency $Q_{\text{ext}}(x)$ is integrated over a log-uniform dust particle-size prior spanning Boyajian et al. (2018)'s own stated constraint region, and compared against the solid-occulter prediction of zero wavelength dependence at any macroscopic scale. At the 2018-01-02 multiband result the achromaticity statistic is read at the 90th percentile of the dust-prior's own predictive, a conservative rather than most-favorable reading of the paper's "at least some dips" language, giving $\ln K = -17.4$ for the artificial hypothesis, decisively null-favored and consistent with the paper's own "inconsistent with... in-line with" language. P4 is satisfied; the channel was never live before the resolving datum. The exact measured reddening slope and its uncertainty were not extracted from the primary photometry in this pass; the scoring instead uses the qualitative constraint the paper itself commits to, a stated limitation rather than a claimed measurement, in the same sense Appendix C's filter-bandpass approximation is stated rather than sourced.

Result. P1, P2, P4, P5, and P6 pass; none of F1, F2, or F4 fires, so the framework does not fail validation. P3 fails as constructed: Collapse A's data-driven test leaves LGM-1 at $\ln K = 2.83$ rather than below 1, meaning the discriminating data the discoverers themselves treated as decisive (the Doppler null and the second source) carry real but not overwhelming weight in this channel decomposition, a genuinely reachable outcome the protocol was designed to surface rather than to guarantee against. The framework registered LGM-1's anomaly 2.8 years before the community's own resolution (November 1967 detection against the May 1968 Gold paper), tracked Hephaistos candidates G, D, and E through anomaly and back to null without ever crossing strong-anomaly ahead of their resolving data, and correctly returned zero evidence for Boyajian's star before the wavelength-dependence channel existed. The construction choices flagged throughout, most consequentially the one-hour natural-coherence reference boundary for LGM-1 and the unmeasured Boyajian reddening slope, are stated rather than sourced from a primary catalog or spectrum; a future pass with access to a machine-readable 1967 radio-source catalog and the raw Boyajian et al. (2018) photometry could tighten both without changing the qualitative result.

Acknowledgements and disclosure

The author thanks the maintainers of the NASA Astrophysics Data System and arXiv, on which the citation verification for this work relied. **AI assistance disclosure:** drafting, citation verification, derivation checking, and figure preparation for this manuscript were performed with substantial assistance from a large language model (Claude, Anthropic), under the author's direction; the author reviewed and takes full responsibility for all claims, derivations, and references. Interactive calculators implementing the quantitative material in Sections 2–4 are available as supplementary material at `https://omegacentauri.`

me.

Data availability

No new observational data were generated. All cited measurements are available in the referenced publications. The figure scripts, shared constants, and machine-readable Monte Carlo output described in Appendix D are mirrored publicly at https://omegacentauri.me/papers/source/ and https://omegacentauri.me/papers/figs/, alongside the interactive calculator versions of the same computations.

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.

Amaro-Seoane, P. (2018). Relativistic dynamics and extreme mass ratio inspirals. *Living Reviews in Relativity*, 21:4.

Arca Sedda, M., Amaro-Seoane, P., and Chen, X. (2021). Merging stellar and intermediate-mass black holes in dense clusters: Implications for LISA, LIGO, and the next generation of gravitational wave detectors. *Astronomy & Astrophysics*, 652:A54.

Ashton, G., Burns, E., Dal Canton, T., Dent, T., Eggenstein, H.-B., Nielsen, A. B., Prix, R., Was, M., and Zhu, S. J. (2018). Coincident detection significance in multimessenger astronomy. *The Astrophysical Journal*, 860(1):6.

Bahcall, J. N. and Wolf, R. A. (1976). Star distribution around a massive black hole in a globular cluster. *The Astrophysical Journal*, 209:214–232.

Bahcall, J. N. and Wolf, R. A. (1977). The star distribution around a massive black hole in a globular cluster. ii. unequal star masses. *The Astrophysical Journal*, 216:883–907.

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.

Barausse, E., Cardoso, V., and Pani, P. (2014). Can environmental effects spoil precision gravitational-wave astrophysics? *Physical Review D*, 89:104059.

Bardeen, J. M. (1970). Kerr metric black holes. *Nature*, 226:64–65.

Bardeen, J. M., Press, W. H., and Teukolsky, S. A. (1972). Rotating black holes: locally nonrotating frames, energy extraction, and scalar synchrotron radiation. *The Astrophysical Journal*, 178:347–370.

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.

Binney, J. and Tremaine, S. (2008). *Galactic Dynamics*. Princeton University Press, Princeton, 2nd edition.

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.

Breschi, M., Gamba, R., Carullo, G., Godzieba, D., Bernuzzi, S., Perego, A., and Radice, D. (2024). Bayesian inference of multimessenger astrophysical data: joint and coherent inference of gravitational waves and kilonovae. *Astronomy & Astrophysics*, 689:A51.

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. (2025). The intermediate mass black hole in Omega Centauri: Constraints on accretion from JWST. Submitted to ApJ.

Colom i Bernadich, M., Dai, S., Abbate, F., Kerr, M., Bachetti, M., Buchner, S., Johnston, S., Burgay, M., Possenti, A., Ridolfi, A., Freire, P. C. C., Kramer, M., Stappers, B., et al. (2026). A joint MeerKAT and Parkes view of Omega Centauri: New TRAPUM searches and pulsar timing. *arXiv e-prints*, page arXiv:2603.21845. Under review.

Curtis, O., Adams, V. H., Angerhausen, D., Bates, J., Berea, A., Dick, S. J., Elvis, M., Khatri, S. P., Linares, R., Muco, M., Seager, S., and Wright, J. T. (2026). The Dyson Minds 2025 workshop: SETI around black holes. *Publications of the Astronomical Society of the Pacific*, 138:046001.

de Vita, R., Trenti, M., and MacLeod, M. (2018). Wandering off the centre: a characterization of the random motion of intermediate-mass black holes in star clusters. *Monthly Notices of the Royal Astronomical Society*, 475(2):1574–1586.

Dicken, D. et al. (2024). JWST MIRI flight performance: Imaging. *Astronomy & Astrophysics*, 689:A5.

Dvali, G. and Osmanov, Z. N. (2023). Black holes as tools for quantum computing by advanced extraterrestrial civilizations. *International Journal of Astrobiology*, 22(6):617–640.

Fishbach, M., Holz, D. E., and Farr, B. (2017). Are LIGO's black holes made from smaller black holes? *The Astrophysical Journal Letters*, 840:L24.

Forgan, D. H., Wright, J. T., Tarter, J. C., Korpela, E. J., Siemion, A. P. V., Almar, I., and Piotelat, E. (2019). Rio 2.0: revising the Rio scale for SETI detections. *International Journal of Astrobiology*, 18(4):336–344.

Fragione, G., Leigh, N. W. C., Ginsburg, I., and Kocsis, B. (2018). Tidal disruption events and gravitational waves from intermediate-mass black holes in evolving globular clusters across space and time. *The Astrophysical Journal*, 867:119.

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.

Gerosa, D. and Berti, E. (2017). Are merging black holes born from stellar collapse or previous mergers? *Physical Review D*, 95:124046.

Gnedin, O. Y. and Ostriker, J. P. (1999). On the self-consistent response of stellar systems to gravitational shocks. *The Astrophysical Journal*, 513:626–637.

González Prieto, E., Rodriguez, C. L., and Cabrera, T. (2025). Growing the intermediate-mass black hole in Omega Centauri. *The Astrophysical Journal Letters*, 990:L69.

Grimaldi, C. and Marcy, G. W. (2018). Bayesian approach to SETI. *Proceedings of the National Academy of Sciences*, 115(42):E9755–E9764.

Häberle, M., Neumayer, N., Bellini, A., Libralato, M., Clontz, C., Seth, A. C., et al. (2024a). oMEGACat. II. Photometry and proper motions for 1.4 million stars in Omega Centauri and its rotation in the plane of the sky. *The Astrophysical Journal*, 970:192.

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., Anderson, J., Whitaker, M., van de Ven, G., Watkins, L. L., Clontz, C., et al. (2024b). Fast-moving stars around an intermediate-mass black hole in ω Centauri. *Nature*, 631(8020):285–289.

Hamilton, C. and Rafikov, R. R. (2019). Secular dynamics of binaries in stellar clusters – I. General formulation and dependence on cluster potential. *Monthly Notices of the Royal Astronomical Society*, 488:5489–5511.

Harris, W. E. (1996). A catalog of parameters for globular clusters in the Milky Way. *The Astronomical Journal*, 112:1487. 2010 edition: arXiv:1012.3224.

Heggie, D. and Hut, P. (2003). *The Gravitational Million-Body Problem: A Multidisciplinary Approach to Star Cluster Dynamics*. Cambridge University Press, Cambridge.

Hughes, S. A. and Blandford, R. D. (2003). Black hole mass and spin coevolution by mergers. *The Astrophysical Journal Letters*, 585:L101–L104.

Inoue, M. and Yokoo, H. (2011). Type III dyson sphere of highly advanced civilisations around a super massive black hole. *Journal of the British Interplanetary Society*, 64:58–62.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.

Lacki, B. C. (2025). Ground to dust: Collisional cascades and the fate of Kardashev II megaswarms. *The Astrophysical Journal*, 985:191.

Li, J.-K., Tao, Z.-Z., and Zhang, T.-J. (2026). Narrowband radio technosignature search toward 3I/ATLAS with FAST.

Mahida, A. D. et al. (2026). No evidence for accretion around the intermediate-mass black hole in Omega Centauri. *The Astrophysical Journal*, 996:122.

Mandel, I., Brown, D. A., Gair, J. R., and Miller, M. C. (2008). Rates and characteristics of intermediate mass ratio inspirals detectable by Advanced LIGO. *The Astrophysical Journal*, 681:1431–1447.

McInnes, C. R. (2026). Stellar engines and Dyson bubbles can be stable. *Monthly Notices of the Royal Astronomical Society*, 546:1–18.

Merritt, D. (2005). Brownian motion of a massive binary. *The Astrophysical Journal*, 628:673–678.

Narayan, R., Chael, A., Chatterjee, K., Ricarte, A., and Curd, B. (2022). Jets in magnetically arrested hot accretion flows: geometry, power and black hole spin-down. *Monthly Notices of the Royal Astronomical Society*, 511(3):3795–3813.

Nitschai, M. S., Neumayer, N., Clontz, C., Häberle, M., Seth, A. C., Husser, T.-O., et al. (2023). oMEGACat. I. MUSE spectroscopy of 300,000 stars within the half-light radius of ω Centauri. *The Astrophysical Journal*, 958:8.

Opatrný, T., Richterek, L., and Bakala, P. (2017). Life under a black sun. *American Journal of Physics*, 85(1):14–22.

Raval, T. and Srikanth, D. (2024). Exploring the Dyson ring: Parameters, stability and helical orbit. Published in Proc. IEEE SPACE 2024.

Ripperda, B., Liska, M., Chatterjee, K., Musoke, G., Philippov, A. A., Markoff, S. B., Tchekhovskoy, A., and Younsi, Z. (2022). Black hole flares: Ejection of accreted magnetic flux through 3D plasmoid-mediated reconnection. *The Astrophysical Journal Letters*, 924:L32.

Sheikh, S. Z., Smith, S., Price, D. C., DeBoer, D., Lacki, B. C., Czech, D. J., et al. (2021). Analysis of the Breakthrough Listen signal of interest blc1 with a technosignature verification framework. *Nature Astronomy*, 5:1153–1162.

Spitzer, Lyman, J. (1987). *Dynamical Evolution of Globular Clusters*. Princeton University Press, Princeton, NJ.

Stone, N. C. and Metzger, B. D. (2016). Rates of stellar tidal disruption as probes of the supermassive black hole mass function. *Monthly Notices of the Royal Astronomical Society*, 455:859–883.

STScI (2026). JWST user documentation: MIRI sensitivity. https://jwst-docs.stsci.edu/jwst-mid-infrared-instrument/miri-performance/miri-sensitivity. Table 1, ETC 6.0; accessed 2026-07-30.

Swanson, T. (2026a). The economics of inward migration: Relocation versus densification for computation-maximizing civilizations. Preprint; omegacentauri.me.

Swanson, T. (2026b). Inward resolutions of the Fermi paradox: A critical review of migration down thermodynamic gradients. Preprint; omegacentauri.me/paper.html.

Swanson, T. (2026c). The macro transcension hypothesis: Spinning black holes in dense stellar clusters as thermodynamic attractors for advanced civilizations, with Omega Centauri as an observational test bed. Preprint; omegacentauri.me/paper.html.

Swanson, T. (2026d). A multi-messenger technosignature and anomaly-detection campaign for Omega Centauri. Preprint; omegacentauri.me/campaign.html.

Tchekhovskoy, A., Narayan, R., and McKinney, J. C. (2011). Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole. *Monthly Notices of the Royal Astronomical Society: Letters*, 418:L79–L83.

Thorne, K. S. (1974). Disk-accretion onto a black hole. II. Evolution of the hole. *The Astrophysical Journal*, 191:507–519.

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.

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.

Wright, J. T. (2020). Dyson spheres. *Serbian Astronomical Journal*, 200:1–18.

Yang, H. et al. (2026). GRMHD simulations of magnetized accretion disk/jet: variabilities of black holes and spectral energy distributions in magnetic states. *Universe*, 12(5):142.

Yuan, F. and Narayan, R. (2014). Hot accretion flows around black holes. *Annual Review of Astronomy and Astrophysics*, 52:529–588.