The Contested Central Dark Mass of Omega Centauri: A Pre-registered Joint Analysis of Fast Stars, Pulsar Timing, and the Limits of the Proper-Motion Dispersion Profile

Tim Swanson

The Omega Centauri Society / Post Oak Labs

tim@postoaklabs.com

August 2026

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

Abstract

The center of Omega Centauri carries the sharpest unresolved mass question in globular-cluster dynamics: seven fast-moving stars inside the central 3″ imply an enclosed dark mass $\gtrsim 8{,}200\,M_\odot$ (Häberle et al., 2024), while joint stellar-kinematics and pulsar-timing modeling favors an extended ~ 2–$3\times10^5\,M_\odot$ remnant component and caps any point mass at $\lesssim 6{,}000\,M_\odot$ (Bañares-Hernández et al., 2025). We report the first joint likelihood analysis of the individual fast stars, the 19-pulsar TRAPUM timing set (Colom i Bernadich et al., 2026) read correctly as one-sided acceleration bounds, and the oMEGACat 40-bin proper-motion dispersion and anisotropy profiles (Häberle et al., 2025), with the dark component parameterized continuously in mass and scale radius from the point-mass limit to the extended-cluster regime. The analysis was pre-registered before real-data contact, with prior grids, decision criteria, and validation gates fixed in advance and five dated amendments disclosed. The pre-registered outcome is the one we report: the data cannot yet decide between a compact and an extended central dark mass, and we can now say why, quantitatively, for each data type. The pre-registered outcome — the data cannot yet decide — is itself a publishable finding that localizes the indecision to specific, named causes per data type. The dispersion profile, measured to 0.03–0.05 km s^{-1}, rejects every smooth spherical non-rotating model ($\chi^2/\nu \approx 473$ in the outer bins) through a coherent, component-differential residual that no radial systematic term can absorb, and its statistical errors sit a factor 10–40 below its own physics-derived systematic budget; on those two grounds, by a pre-set gate, we retire it from the verdict. The pulsar accelerations, all censored, span 3.0 nats across the entire parameter plane. The pulsar spin-frequency second derivatives are inconsistent with cluster jerks and, after correcting a numerical constant in the published nearest-neighbour jerk scale ($\xi = 3.4596$, not 3.04), we show jerk discrimination requires a timed-pulsar census of order 10^2, several times the 19 available, at any timing precision, because the nearest-neighbour floor is a 1-stable process that does not average away. The fast-star and pulsar legs jointly, with the profile leg retired and under the assumed tracer cusp, return a compact optimum ($M_{\rm dark} \approx 2.0$–$2.5\times10^4\,M_\odot$ at the point-mass limit, $\ln K = +10$ to $+12$; range over all prior cells and brackets, Section 5.3). A calibration campaign, now complete, traced the gate battery's one failing check to the retired profile leg itself; re-scored on the verdict configuration under two further dated gate-scope amendments, the battery passes at the fiducial mass-to-light bracket by the point-estimate rule the gate fixes (90 per cent coverage 0.8825 extended, Wilson 0.847–0.910, and 0.865 point, Wilson 0.828–0.895; null and sign-recovery checks passed non-trivially), and the compact preference is quotable as a calibrated 90 per cent statement of consistency between the fast-star and pulsar legs, at the fiducial bracket only and conditional on a Bahcall–Wolf tracer cusp: regenerated at the shallower published stellar-cusp slope the same check returns 0.660 and fails the gate. The same statement is conditional on the TRAPUM pulsar reduction: under the earlier reduction of the same five pulsars, carried here as a labelled sensitivity, the consistency criterion fails in every prior cell of the fit of record and the evidence changes sign. The statement is further conditional on the truncated-polytrope speed-tail family; a hard-truncated isothermal alternative is admissible against the same kinematics and remains untested. It is a statement of consistency, and no resolution of the profile-leg tension, which the gates removed from the verdict, is claimed. A formation-physics overlay sharpens the impasse from the other side: published retention physics builds the compact solution readily and reaches the extended solution only through hybrid configurations that the extended solution's own point-mass cap excludes. We state what would decide the question, and in which order.

Keywords: globular clusters: individual (NGC 5139); intermediate-mass black holes; stellar dynamics; pulsars: timing; methods: statistical; pre-registration

Contents

1 Introduction

ω Cen (NGC 5139) has hosted claims and counter-claims of a central intermediate-mass black hole (IMBH) for nearly two decades: integrated-light kinematics for $\sim 4 \times 10^4\, M_\odot$ (Noyola et al., 2008, 2010), proper-motion modeling capping the mass at $\lesssim 1.2 \times 10^4\, M_\odot$ at $\sim 1\sigma$ ($1.8 \times 10^4\, M_\odot$ at 3σ), with the verdict hinging on the adopted center (van der Marel and Anderson, 2010), radial anisotropy shown to mimic much of the signal (Zocchi et al., 2017), stellar-mass black-hole populations shown to mimic more of it (Zocchi et al., 2019; Baumgardt et al., 2019). The modern form of the question is sharper on both sides. Häberle et al. (2024) found seven fast-moving stars inside the central $3''$, five robust, with velocities above the central escape speed of any IMBH-free model, giving a firm kinematic lower bound of 8,200 $M_\odot$ on an enclosed dark mass. Bañares-Hernández et al. (2025) combined stellar kinematics with millisecond-pulsar accelerations and found the data favor an extended dark component of $2–3 \times 10^5\, M_\odot$ at parsec scale, with a 3σ cap of 6,000 $M_\odot$ on any point mass, below the fast-star bound. The TRAPUM timing release (Colom i Bernadich et al., 2026) expanded the pulsar census to 19 with an IMBH-insensitive upper limit of $10^5\, M_\odot$, sharpening the data without resolving the contradiction.

The methodological history of this exact contest in other clusters is cautionary in both directions. In 47 Tucanae, a pulsar-dynamics IMBH claim (Kızıltan et al., 2017) did not survive multimass remnant modeling (Mann et al., 2019; Smith et al., 2024); in NGC 6397, a claimed IMBH resolved into a diffuse inner subcluster (Vitral and Mamon, 2021); and Aros et al. (2020) showed with mock data that Jeans-type fits fabricate IMBH masses at this scale when anisotropy or mass-to-light structure is mis-modeled. Any new joint analysis of ω Cen enters a literature where the last three comparable verdicts were overturned by systematics, and it should be built accordingly.

This paper reports such an analysis, built with three defenses the prior literature lacked in combination. First, a genuinely joint likelihood: the individual fast stars (velocities, positions, selection function, contamination), the pulsar line-of-sight accelerations read as the one-sided bounds they are, and the dispersion and anisotropy profiles, all constraining one dark component parameterized continuously by mass $M_{\rm dark}$ and Plummer scale a, from the point-mass limit ($a \to 0$) to the extended-cluster regime, so that the compact and extended hypotheses are regions of one parameter plane rather than separate models with separate machinery. Second, pre-registration: the prior grid, the decision criteria, the validation gates, and the amendment protocol were committed to a public repository before any likelihood touched real data, and all five subsequent amendments are dated and disclosed (Section 3). Third, adversarial validation: the pipeline had to reproduce the Aros et al. (2020) failure mode on demand (inject anisotropy, fit isotropic, measure the fabricated mass) and demonstrate calibrated coverage on mock injections before first real-data contact, and a second gate battery governed what became quotable afterward. The pre-contact coverage check was passed at 0.84 and 0.92 before any numerical band existed; once Amendment A2 fixed the band at $[0.85, 0.95]$ that pass was voided and superseded, and Table 2 scores it against the band with everything else.

The outcome is the pre-registered null: the data cannot yet decide. We consider that outcome, reached this way, more useful than another overturnable verdict, because the machinery now localizes the indecision to specific, named causes per data type, and because the same machinery yields the first quantitative statement of what a deciding dataset must contain. Sections 2–4 describe the data, the pre-registered design, and the validated machinery, including a correction to a published constant in the pulsar-jerk formalism that other groups may wish to note independently of anything else in this paper. Sections 5–6 report the joint fits under both pre-registered error models and diagnose, leg by leg, where the discriminating power actually resides. Section 7 adds the formation-physics overlay, Section 8 states what would decide the question, and Section 9 concludes.

2 Data

All inputs were extracted from their primary sources into versioned, provenance-annotated tables before any modeling, with every value carrying its source location and verbatim context; the fit records the content hash of each input it consumed. Four datasets enter.

Fast stars. The seven candidates of Häberle et al. (2024), with the robust/candidate split preserved: the headline analysis uses the five robust stars (stars A, C, D, E, F; star letters and pulsar letters are separate series throughout this paper and are prefixed wherever either could be meant), and the all-seven variant is a reported sensitivity row. Positions, proper motions, and uncertainties come from their Extended Data tables; the selection function is reconstructed from the stated quality-cut pass fraction (157,320 of 241,133 stars, completeness 0.652), the 3″ search radius, the 2.41 $\mathrm{mas\,yr^{-1}}$ threshold, and the published contamination density. The contamination normalization reproduces their expected foreground count (0.0735 against the 0.074 of their main text and the 0.073 of their Extended Data caption), which checks that term and no other part of the selection function. The completeness is the catalogue-wide pass fraction over $r < 90''$ and enters as a single radius-independent factor; Section 6 states the direction and the unmeasured size of the bias that carries into a sample confined to $r < 3''$. The N-body-informed mass range quoted in that paper's methods is model output, not data, and enters nowhere in this analysis.

Pulsar accelerations. The TRAPUM timing solutions (Colom i Bernadich et al., 2026): 8 of 19 pulsars have measured spin-period derivatives; every derived cluster line-of-sight acceleration is a one-sided bound, because the intrinsic spin-down is unknown and non-negative. Seven bounds enter the likelihood; one (pulsar C, which is unrelated to star C above) is excluded because its published table entry is typeset in a form we could not disambiguate, and we do not silently repair source tables. Three of the seven bounds are negative-signed (pulsars B, D, K), and those carry what constraining power the leg has. We verified the published subtraction by recomputing pulsar A's bound from its P, $\dot{P}$, and the tabulated Shklovskii and Galactic terms (1.9796×10^{-9} against the printed $1.98 \times 10^{-9}\,\mathrm{m\,s^{-2}}$); recomputing the Shklovskii terms themselves from the tabulated proper motions is a further check, and Section 6 reports where it does and does not reproduce. The earlier five-pulsar solutions of Dai et al. (2023) are carried as a labelled sensitivity variant; Section 6 reports what that comparison revealed.

Dispersion, anisotropy, rotation. The oMEGACat 40-bin sky-radial and sky-tangential proper-motion dispersion profiles with their anisotropy ratios, $r = 1.8''$–$311''$, statistical errors 0.023–0.05 $\mathrm{km\,s^{-1}}$ in the best-measured bins, retrieved from the survey's machine-readable release and checksum-verified (Häberle et al., 2025); the integrated-light rotation profile (amplitude rising to $8.4 \pm 0.8\,\mathrm{km\,s^{-1}}$ at $4.7'$) frozen from Häberle et al. (2026); kinematic distance $5{,}494 \pm 61$ pc from the same survey, with the 5,200 pc value adopted by Bañares-Hernández et al. (2025) carried as a reported axis rather than harmonized away.

Visible-model structure. No machine-readable radial surface-brightness array for ω Cen exists in the public record; we verified this against the survey chain, the Baumgardt and Hilker (2018) catalogue release, and its archival files before concluding it, and the gap forced Amendment A1 below. The visible model is therefore the reported pair: a Plummer sphere at the catalogue half-light scale, and the published $\alpha\beta\gamma$ profile of Bañares-Hernández et al. (2025) at its released best-fit parameters, each normalized on the outer dispersion bins ($r > 100''$).

The calibration domain, and what it absorbs. Earlier drafts of this paper, and the calibration module itself, justified that $r > 100''$ domain by asserting that the dark component contributes below one part in 10^3 of the enclosed mass there everywhere in the prior grid. That assertion is wrong, and we withdraw it. Recomputed in the model code the fit itself uses, with $M_{\mathrm{vis}} = 3.54 \times 10^6\,M_\odot$ and $b_{\mathrm{vis}} = 7.56$ pc at 5.494 kpc, the enclosed visible mass inside $100''$ (2.66 pc) is $1.30 \times 10^5\,M_\odot$, and the dark fraction $M_{\mathrm{dark}}(< r)/M(< r)$ runs as in Table 1. The claim holds only at the lowest corner of the prior grid and only outside $\sim 250''$; at $100''$ it is exceeded by a factor 7.6 even there, and the extended solution of

Table 1: Dark-component fraction of the enclosed mass, $M_{\rm dark}(<r)/[M_{\rm vis}(<r)+M_{\rm dark}(<r)]$, at four corners of the pre-registered prior grid, over the radii of the visible-model calibration domain. Both components are Plummer spheres; the visible model is the calibrated Plummer member at the kinematic distance. Recomputed twice, in closed form and through the fit's own `M_enc` routine, agreeing to machine precision.

r	r [pc]	$M_{\rm vis}(<r)$ [$M_\odot$]	$10^3\,M_\odot$ $a=10^{-4}$ pc	$2.5\times10^4\,M_\odot$ $a=10^{-4}$ pc	$2.5\times10^5\,M_\odot$ $a=0.7$ pc	$10^6\,M_\odot$ $a=3.0$ pc
100″	2.66	1.30×10^5	0.0076	0.161	0.635	0.6923
150″	4.00	3.62×10^5	0.0028	0.065	0.398	0.5858
250″	6.66	1.02×10^6	0.0010	0.024	0.194	0.4254
311″	8.28	1.43×10^6	0.0007	0.017	0.148	0.3679

Bañares-Hernández et al. (2025) contributes 63 per cent of the enclosed mass in the innermost calibration bin. The consequence is a real bias, and we quantify it rather than assume it away: recalibrating $M_{\rm vis}$ on the same bins with each prior corner's dark component present, under the same A2 marginal profile likelihood the fit of record uses, moves $M_{\rm vis}$ by -0.14, -2.0, -10.3 and -46.2 per cent at the four corners of Table 1. The first two sit inside the 9.7 per cent mass-to-light bracket the fit already carries; the extended corners do not, so the normalization is biased against extended solutions by up to a bracket width at $2.5\times10^5\,M_\odot$ and by more than four bracket widths at $10^6\,M_\odot$. Two consequences run through the rest of the paper. The compact optimum of the fit of record is not affected at the level of its own error budget, because a $2.5\times10^4\,M_\odot$ point mass shifts the calibration by 2 per cent. The retirement of the profile leg is argued in Section 6 on the component-differential residual and the systematics-to-statistics ratio alone, both of which are independent of this claim; the signal-content argument is withdrawn from it. Script: `h/fit/dark_fraction_table.py`, `h/fit/calib_bias.py`. This paper's $M_{\rm vis}=3.54\times10^6\,M_\odot$ agrees with Paper E's $3.55\times10^6\,M_\odot$ Plummer total to <1 per cent; see Paper F App. A for the shared-parameter table across the set.

3 Pre-registered design

Epistemic status: This section describes process, not results. The pre-registration file, its git history, its amendments, the referee reports that produced Amendment A2, and every validation and fit report are distributed with the paper source; each claim in this section is checkable against a committed document that predates the result it governs.

The analysis plan was committed before any real-data likelihood evaluation (Nosek et al., 2018). Its elements: (i) the model, one dark Plummer component with $(M_{\rm dark}, a)$ continuous over $[10^3, 10^6]\,M_\odot\times[10^{-4}, 3]$ pc, the point-mass limit included; (ii) prior-sensitivity sweeps over sub-ranges of both axes, with mass-function and retention nuisances bracketed rather than point-chosen; (iii) decision criteria fixed in advance, including the rule that any compact-versus-extended Bayes factor is reported only as a table over all prior cells, that no $|\ln K|<1$ appears in the abstract, and that "the data cannot decide" is a declared, publishable outcome with a mechanical trigger (criterion 3 unmet in at least half the prior cells); (iv) validation gates that had to pass before real-data contact: injection–recovery with calibrated coverage including the Aros et al. (2020) anisotropy reproduction (G1), validation of the jerk-statistics module against the published analytic distributions (G2), and a formation-channel consistency map that never enters the likelihood (G3).

Five amendments were required, all dated, all disclosed here; the first two changed the error model, the next two only the scoring scope of the gates after their own kill switch fired, and the fifth only the reading of how per-bracket coverage is generated and compared (Section 5.3). **A1:** the pre-registered fiducial visible model was defined on a surface-brightness profile that turned out not to exist in machine-readable form; the amendment names the Plummer/$\alpha\beta\gamma$ pair as a reported bracket, neither member promoted. **A2:** the first real-data run exposed that the dispersion profile's assumed error model controlled the compact-versus-extended verdict outright (Section 5); rather than choose a value, we froze the run and took the question to two independent methodological reviews and, on their split, a tiebreak adjudication, whose ruling was adopted verbatim: all quadrature error floors, fixed or fitted, were struck as primary (a coherent one-signed residual is mean-model bias, and inflating variance against a biased mean is equivalent to deleting the leg gradually), replaced by an explicit mean-model discrepancy term in the style of Kennedy and O'Hagan (2001), a four-knot spline in $\log r$ with physics-scaled priors and a small white term, both marginalized, together with four quotability gates: a sign and runs test on the residuals, sign-unanimity of $\ln K$ across every prior cell and bracket, an injection-calibration rerun at real-data precision with deliberate coherent misspecification, and a budget-overrun kill switch that retires the leg from the verdict if the required discrepancy exceeds twice its physics budget. **A3:** when the A2-G-d kill switch fires, gate A2-G-c is evaluated on the verdict configuration, with mandatory disclosures that travel with any quoted region. **A4:** the same re-scoping for A2-G-b, and A2-G-a vacated, the gate having been defined on a leg no longer in the verdict. Both are reproduced verbatim in Appendix B. **A5:** the pre-registration left implicit whether per-bracket coverage is scored on mocks generated at the bracket being fitted; the amendment reads it that way, fixes the bracket-equivalence criterion as a relative one against the fiducial diagonal in the same configuration and sample, and changes no prior, threshold, gate, or verdict (Section 5.3).

A1 and A2 changed the error model. A3 and A4 changed the scope on which the quotability gates are scored, after A2's own kill switch had fired and after those gates had failed as originally scoped; A5 changed a coverage-generation reading. None of the five changed the priors, the model, the prior grid, the decision criteria, or the data. We state the A3/A4 re-scoping in that form because the alternative form, that the amendments left the quotability rules untouched, is not defensible: a gate battery that fails as written and passes once its scope is amended has had its quotability rule changed, whatever the amendment is titled. The re-scoping is defended on its merits in Section 5.3, where the calibration campaign traces the failing check to the retired leg itself. A2 was adopted with its gates fixed before the rerun that they then failed. One further process deviation is disclosed here rather than discovered later: the first run derived its mass-to-light bracket from its own calibration statistics instead of the pre-registered retention-anchored brackets; the deviation was identified in review, folded into the A2 record, and the bracket re-derived under the amended error model (its half-width widening from 4 to 9.7 per cent, the conservative direction).

4 Machinery and validation

4.1 Joint likelihood and gate G1

The three legs share one potential: visible model plus dark Plummer component. The fast-star leg evaluates each star's proper-motion speed against the local high-velocity tail of a truncated-polytrope speed distribution matched to the Jeans solution, with the selection function, completeness, and contamination in the likelihood, and a Bahcall–Wolf tracer cusp (Bahcall and Wolf, 1976) applied identically in generation and fitting inside the dark component's influence radius. The pulsar leg computes the probability that the model line-of-sight acceleration lies below each one-sided bound, with the unknown line-of-sight positions marginalized over the tracer density in the manner of Prager et al. (2017). The profile leg fits both dispersion components with radial anisotropy marginalized over an Osipkov–Merritt scale grid, never fixed isotropic.

Gate G1 required, on mocks: calibrated 90 per cent coverage for a $4 \times 10^4\,M_\odot$ point-mass injection and a $2.5 \times 10^5\,M_\odot$ extended injection (achieved: 0.84 and 0.92, inner quantiles under-covering, declared and flagged as approximate); no manufactured detection on a null injection (0.00); compact/extended false-preference rates below 0.10 (0.02 and 0.00); and the Aros et al. (2020) reproduction, which we regard as the gate's core: mocks generated with radial anisotropy and no dark mass, fitted with isotropy forced, manufacture $8.0 \times 10^4\, M_\odot$ of dark mass and a 76 per cent spurious detection rate; the same data under the pre-registered anisotropy marginalization return a null with a 0 per cent spurious rate. The degeneracy that overturned prior verdicts is reproducible on demand and demonstrably closed by the marginalization.

4.2 The jerk module, gate G2, and a correction to a published constant

Pulsar spin-frequency second derivatives probe a central mass through the jerk field, against a stochastic floor from nearest-neighbour stellar encounters (Prager et al., 2017; Abbate et al., 2019). Implementing that formalism, we found the published value of the dimensionless constant in the nearest-neighbour jerk scale, $\dot{a}_0 = (2\pi\xi/3)\, G\langle m\rangle\sigma n$, to be in error. The defining double integral evaluates in closed form:

$$\xi = \int_0^\infty \mathrm{d}x \int_{-1}^{1} \mathrm{d}\mu \left[1 - e^{-(1+3\mu^2)/(2x^2)}\right] = \sqrt{2\pi}\left[1 + \frac{\ln(2+\sqrt{3})}{2\sqrt{3}}\right] = 3.4596, \tag{1}$$

against the published $\xi \simeq 3.04$, a 13.8 per cent difference verified six ways: one closed-form reduction, four cross-checks of the same integral (direct quadrature, symbolic evaluation, antiderivative differentiation, 50-digit arithmetic), and one Poisson-field Monte Carlo of the physical system that never evaluates the formula and is the only one of the six independent of it. The published constant understates the jerk noise floor by 12.1 per cent, in the direction that inflates the apparent significance of jerk-based central-mass inferences; the derivation in the source papers is otherwise sound, and we have used the corrected value throughout. Gate G2 validated the module against the published distributional identities. Two of its twenty checks are the ξ comparison itself and cannot pass by construction; the remaining 18 agree at 10^{-10} or better. Box 1 gives the derivation and the verification ledger, since this is the one result here that other groups may want to check without reading anything else in the paper.

The module then delivered a negative result we consider as useful as the correction. The line-of-sight nearest-neighbour jerk is a 1-stable (Cauchy) process: it has no finite variance, so it does not average away with more pulsars or better timing. Mapping discrimination power across four decades of timing sensitivity, jerks cannot separate a point mass from an extended component of equal mass anywhere in 10^4–$2.5 \times 10^5\, M_\odot$ at 19 pulsars, and the verdict is nearly unchanged from 10^{-22} to $10^{-19}\,\mathrm{m\,s^{-3}}$: the floor, not the precision, is the limit. Reaching a median $\ln K > 3$ requires a census of order 10^2 timed pulsars (indicatively $\sim$100 at $2.5 \times 10^5\, M_\odot$ and $\sim$200 at $10^5\, M_\odot$; a finalized forecast at the corrected floor is flagged for the released record rather than quoted as settled). The G2-era floor also carried a core-density normalization $\sim$59 times too high for this cluster, corrected in the same pass as ξ. Consistently, the eight measured TRAPUM $\ddot{\nu}$ values are inconsistent with cluster jerks at all (Fisher $p = 2.4 \times 10^{-8}$, sign coherence $p = 0.008$). That exclusion is scored against the stochastic nearest-neighbour floor, which is the larger of the two cluster terms; the deterministic mean-field jerk from the smooth potential, of order $a_{\mathrm{LOS}} v/r \approx 2 \times 10^{-22}\,\mathrm{m\,s^{-3}}$ at the relevant radii, is comparable to the median stochastic floor of $4.5 \times 10^{-22}\,\mathrm{m\,s^{-3}}$ and two to three decades below the measured values, so the exclusion holds against both terms. The measured set matches the source paper's own caution that these values may be spurious, and it enters the analysis only through this goodness statement.

4.3 The formation-physics overlay, gate G3

A buildability map over the (M_{dark}, a) plane was compiled from published formation and retention physics: the merger-driven IMBH growth track of González Prieto et al. (2025), the survival analysis of Martinez

Box 1. The nearest-neighbour jerk constant, derived and checked.

Setup. The nearest-neighbour jerk field of a Poisson stellar background has no finite variance, so its scale is set by a characteristic value rather than an rms. Write the jerk from a single neighbour at separation r moving with relative velocity v as $|\dot{a}| \sim Gmv/r^3$, and ask for the value $\dot{a}_0$ at which the expected number of neighbours producing a larger jerk is unity. Carrying out the angular average over the line of sight and the radial integral over a uniform background of number density n and mean mass $\langle m \rangle$ gives $\dot{a}_0 = (2\pi\xi/3)\, G\langle m \rangle \sigma n$, with the whole geometry collected into one dimensionless double integral, Equation (1).

Reduction. The inner integral over x has an antiderivative in terms of the error function; performing it first leaves a single integral over μ of $\sqrt{\pi/2}\,(1+3\mu^2)^{+1/2}$, which is elementary and gives the inverse hyperbolic sine. Hence $\xi = \sqrt{2\pi}\,[\,1 + \ln(2+\sqrt{3})/(2\sqrt{3})\,] = 3.4595806610$, exactly, with $\ln(2+\sqrt{3}) = \operatorname{arcsinh}\sqrt{3}$.

Verification, six routes, one of them independent of the formula. (i) the closed form above; (ii) adaptive quadrature of the original double integral, agreeing to 2×10^{-13}; (iii) a Poisson-field Monte Carlo that draws 4×10^4 realizations of a random stellar background and measures the scale of the resulting line-of-sight jerk distribution without ever evaluating the integral, returning $\xi = 3.489 \pm 0.028$; (iv) symbolic evaluation; (v) numerical differentiation of the antiderivative used in step (ii); (vi) 50-digit arithmetic on the closed form. Route (iii) also confirms the tail index, with a Hill estimator of 1.025 against the 1-stable value of 1 and a Kolmogorov–Smirnov $p = 0.26$ against the analytic law.

Consequence. The floor scales linearly in ξ, so the floor computed with 3.04 is 12.1 per cent low, and a jerk-based significance computed with it, which scales as $1/\dot{a}_0$, is 13.8 per cent high. Figure 1 shows the two floors against the density they set. The correction does not change any conclusion of this paper, whose jerk result is a non-detection either way, and it does change the noise normalization of any published jerk-based central-mass inference that used the smaller constant.

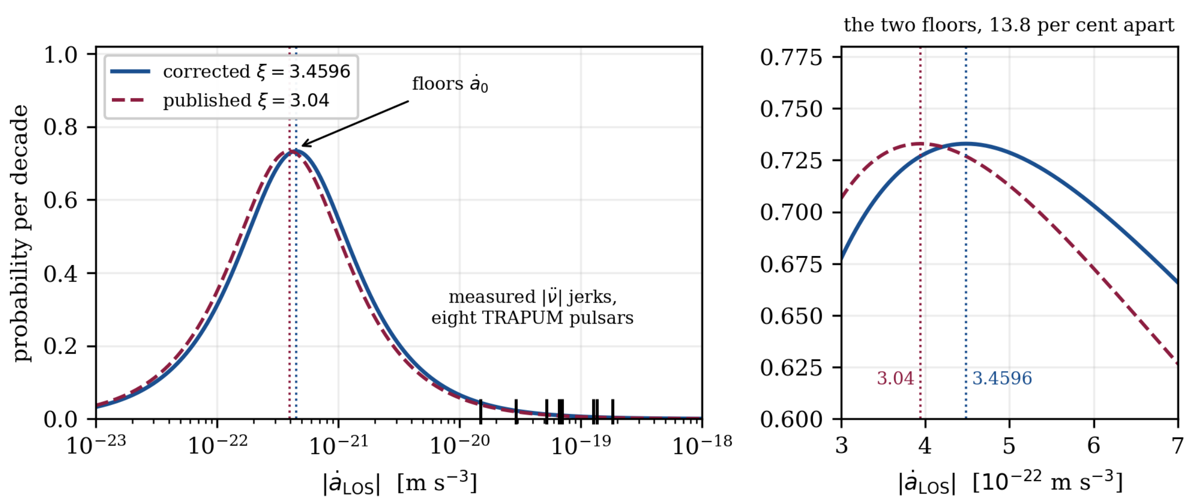

Figure 1: The line-of-sight nearest-neighbour jerk density at the corrected and published values of ξ, evaluated at the median of the eight per-pulsar floors of the released FLAG D record (paper-set dispersion, fiducial mass-segregated millisecond-pulsar density). Left: the full range, with the eight measured TRAPUM $\ddot{\nu}$ jerks as ticks along the axis; per-pulsar margins run 1.5–2.6 decades above the median floor (A 2.2, B 2.5, C 2.6, D 2.1, E 1.8, G 1.5, H 2.2, K 2.5), which is the quantitative form of the statement that they are inconsistent with a cluster origin. The tick set includes pulsar C, which is excluded from the seven-pulsar acceleration likelihood. The law is 1-stable, so the tail falls as $|\dot{a}|^{-2}$ and no amount of averaging over pulsars contracts it. Right: the peak decade, where the 13.8 per cent separation between the two floors is visible. Data and script: `figs/fH_jerk_dist.json`, `figs/fH_expand_v1.py`.

et al. (2026), the retained remnant populations of Dickson et al. (2023, 2024), the hierarchical-merger ceiling of Mai et al. (2026), with the first confirmed stellar-mass black hole in the cluster (Whitaker et al., 2026) as an existence proof and never a rate. The map is a consistency overlay by hard rule: it enters no likelihood and no prior. Its content is taken up in Section 7.

5 The joint fit, twice

5.1 First contact: error-model dependence

The first pre-registered run returned a result about the machinery rather than the cluster, and we report it as such. No smooth, spherical, non-rotating, single-mass Jeans model survives the oMEGACat profile at its measured precision: $\chi^2/\nu = 473$ in the outer bins (179 over all 40) for the better visible model, with one-signed pulls reaching 20σ, model below data across $50''$–$250''$, robust to swapping the visible model and the distance. With no systematic term, the first-contact configuration is dominated by this misfit and returns an extended component at $\ln K \approx -220$; adding a quadrature floor of $1.0\,\mathrm{km\,s^{-1}}$ keeps the extended preference at $\ln K \approx -10$; a $2.0\,\mathrm{km\,s^{-1}}$ floor flips the verdict to compact at +4; deleting the profile leg gives compact at +11. The compact-versus-extended answer was controlled entirely by an assumed error parameter, and the analysis was frozen at that point under the amendment protocol rather than tuned. (The floor values quoted in this paragraph are the first run's; Figure 2 plots the amended rerun's sensitivity-appendix re-evaluation of the same ladder, under the re-derived visible-model normalization and with the amendment's marginalized white term present, which compresses the no-floor extreme from −220 to −11 without changing any sign along the ladder.)

Figure 2 shows this ladder, together with what became of it under Amendment A2. Three details of the first run bear on everything after. The floor value that best repairs χ^2 ($\approx 1.0\,\mathrm{km\,s^{-1}}$) is calibrated against the rejected model itself and is therefore circular; the value that flips the verdict ($2.0\,\mathrm{km\,s^{-1}}$) is rejected by the data from the other side ($\chi^2/\nu = 0.27$, a significant under-dispersion); and the bins that dominate the leg's formal weight, contributing a dynamic range of 3.4×10^5 nats across the parameter plane (the digits beyond the leading two are an artifact of summing 40 bins, not a meaningful precision), carry statistical errors a factor 10–40 below the leg's own physics-derived systematic budget. The leg's dominance was spurious precision, and the reason is the error budget rather than the signal content: Table 1 shows those bins do carry dark-component signal, and for the extended hypothesis they carry most of it.

5.2 Amendment A2 and the gate battery it failed

The A2 rerun replaced every floor with the reviewed discrepancy model: a four-knot spline mean term $\delta(r)$ with Normal(0, $1.0\,\mathrm{km\,s^{-1}}$) knot priors, the scale fixed in advance as the quadrature budget of three named physics terms (energy-equipartition and multimass effects, 0.5–0.8 $\mathrm{km\,s^{-1}}$ from the survey's own equipartition profile; flattening and azimuthal averaging at the cluster's projected ellipticity $\epsilon = 0.17$, per White and Shawl 1987 as compiled in the Harris catalogue, the largest of the published values, which range from 0.08 to 0.17 with method and radius (Geyer et al., 1983; Pancino et al., 2003), so the budget term is conservative; the intrinsic axial ratio is 0.78 ± 0.03 at inclination 50° (van de Ven et al., 2006); second-order rotation leakage), plus a half-Normal(0, $0.3\,\mathrm{km\,s^{-1}}$) white term, both marginalized everywhere, across sixteen configurations (both visible models, both distances, the primary and three mandatory companion rows). Rotation does not enter the mean model, on a mechanism ruling worth recording: the misfit lives in proper-motion dispersions computed about per-bin means, which remove ordered rotation to first order, and the measured rotation is line-of-sight in any case; the data agree, since residual rotation would inflate the tangential component and the tangential component is the low one.

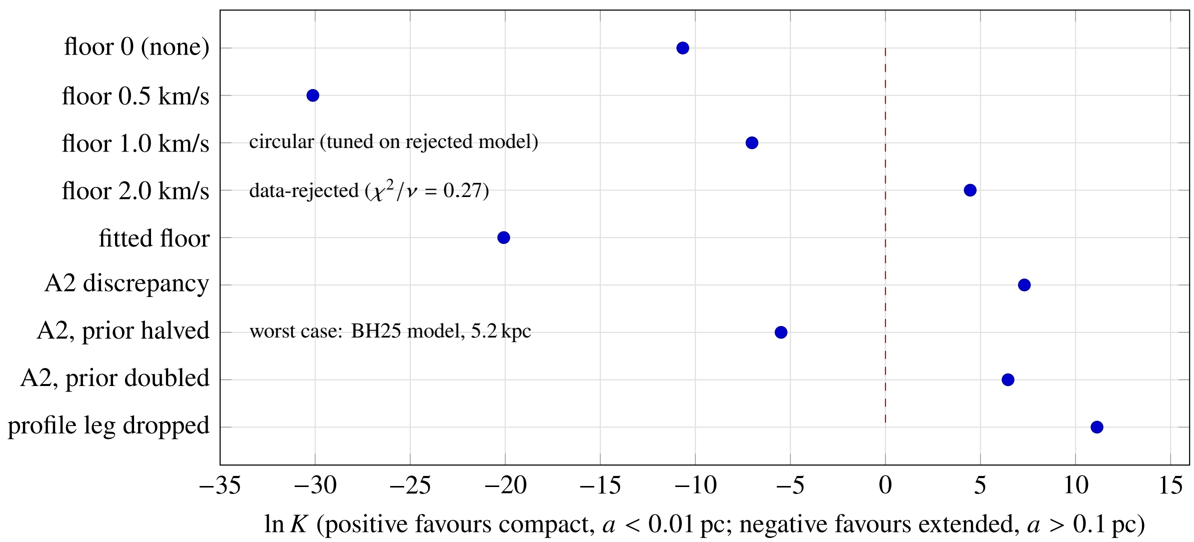

Figure 2: The compact-versus-extended Bayes factor as a function of the profile leg's error model: each row is one labelled configuration from the amended rerun's sensitivity appendix (fiducial bracket, full prior range, Plummer visible model at 5,494 pc except where labelled; the leg-dropped row spans +10.0 to +12.2 across the full bracket set, of which one configuration is plotted). The first-contact run's more extreme no-floor value (−220; Section 5.1) differs by normalization, not sign. Every row is an assumption choice, and the answer changes sign along the ladder. This figure is the paper's diagnosis: the profile leg's output tracks the analyst's error model rather than the dark component's geometry. The pre-registered gates (Section 5.2) formalize that reading, and the leg is retired from the verdict. The ladder is not monotone in the assumed floor, and should not be read as a smooth deflation: a 0.5 km s^{-1} floor sits 19.5 nats further into the extended region than no floor at all, and the fitted floor at −20.08 lies outside the envelope of its fixed-floor neighbours. The bins carry statistical errors of 0.023–0.05 km s^{-1}, so a floor of a few tenths of a km s^{-1} does not deflate the leg gradually; it flattens the bin-weight hierarchy in one step, and which bins dominate changes with it. We report the non-monotonicity rather than model it, because the leg is retired either way and every row below a 2.0 km s^{-1} floor carries the same sign. The four A2 rows are the pre-FLAG-D record; the reconciled rerun that Section 5.3 quotes moves them by up to 0.3 nats and no sign. Data and script: `figs/fH_ladder.json`, `h/fit2/export_ladder_figure.py`.

All four quotability gates failed, each informatively.

The residual is component-differential. The sign test fails in every configuration ($p = 8 \times 10^{-6}$ to 4×10^{-5}), and it fails by component: the model under-predicts the sky-radial dispersion in 33 of 40 bins and over-predicts the sky-tangential in 26 of 40, at every cell of the parameter plane, including the cell where the profile leg fits best. A shared radial discrepancy $\delta(r)$ cannot represent a component-differential residual by construction. Whatever the profile is expressing (anisotropy structure beyond one Osipkov–Merritt scale, flattening, unrelaxed substructure from the cluster's accretion origin), it is not absorbable by any radial error term, fixed or fitted.

The verdict-controlling assumption reproduces one level up. $\ln K$ fails sign unanimity (52 of 288 defined cells carry the minority sign), and the minority concentrates in the halved-prior companion row: at half the knot-prior scale, the worst configuration returns $\ln K = -5.5$ with 99 per cent extended posterior mass, against +7.3 and 97 per cent compact for the primary. Striking the quadrature floor moved the controlling assumption from an error-bar width to a prior width. We take the reproduction of the regress as a result, and we do not pursue a third amendment: the leg's answer is not stable under any error model we or three independent reviews could defend, because the leg does not contain the answer. Figure 3 shows the sixteen configurations at the fiducial bracket, which is the level of the record at which the reproduction is visible in one panel.

The calibration gate localizes the risk. The injection rerun at real-data precision passes its null test with a factor-4 margin (median $|\ln K| = 0.26$ on null injections against a threshold of 1: no manufactured preference, even when the injections carry the same coherent misspecification that broke the profile leg, so the criterion-4 verdict is not an artifact of the miscalibration), recovers the correct sign in 100 per cent of signal injections; the original gate-G1 point-mass injection passed at 0.840 and the extended injection at 0.920 (both $n = 50$; Table 2, rows "G1 point" and "G1 extended"), but that pass is voided and superseded by the later A2-G-c band and rescored by the uniform point rule below, where 0.840 is a FAIL. Under the A2-G-c protocol at real-data precision the machinery under-covers on the extended injection: 90 per cent regions contain the truth 79.5 per cent of the time under coherent misspecification (Wilson interval 0.734–0.845, established at $n = 200$, entirely below the 0.85 requirement). The machinery is uncalibrated on the extended alternative, the one it would need most.

The kill switch condition is met. The fitted discrepancy reaches 2.1–2.4 $\mathrm{km\,s^{-1}}$ in the $\alpha\beta\gamma$-model configurations, beyond twice the flat 1.0 $\mathrm{km\,s^{-1}}$ budget that Amendment A2 fixed as the operative trigger (set conservatively from the interior of the 0.5–1.6 $\mathrm{km\,s^{-1}}$ physics-budget range quoted in Section 6), and exceeds that budget itself in all twelve configurations. Appendix A tabulates the firing status under the flat trigger and under both ends of the range; the recorded verdict used the flat trigger. Per the pre-set fallback, the profile leg is retired from the verdict, surviving only as the visible-model calibration and as the diagnosis of this section.

Three disclosures required verbatim by Amendment A2: the discrepancy term absorbs smooth signal by construction; the profile leg's outer bins ($r > 50''$) carry no verdict-relevant signal at any systematics level compatible with the named physics; and a criterion-4 outcome under these gates is the pre-registered and publishable result. The second of those is reproduced as A2 fixed it and is a statement about recoverable signal at the leg's systematics level, not about the signal present in those bins; Table 1 gives the latter, which is not small.

5.3 The fit of record

With the profile leg retired, the fit of record is the fast-star and pulsar joint analysis. It is stable in a way nothing else in this paper is: across both visible models, both distances, and all mass-to-light brackets, it returns $M_{\mathrm{dark}} = 2.0\text{–}2.5 \times 10^4\, M_\odot$ at the point-mass limit of the grid, compact posterior mass

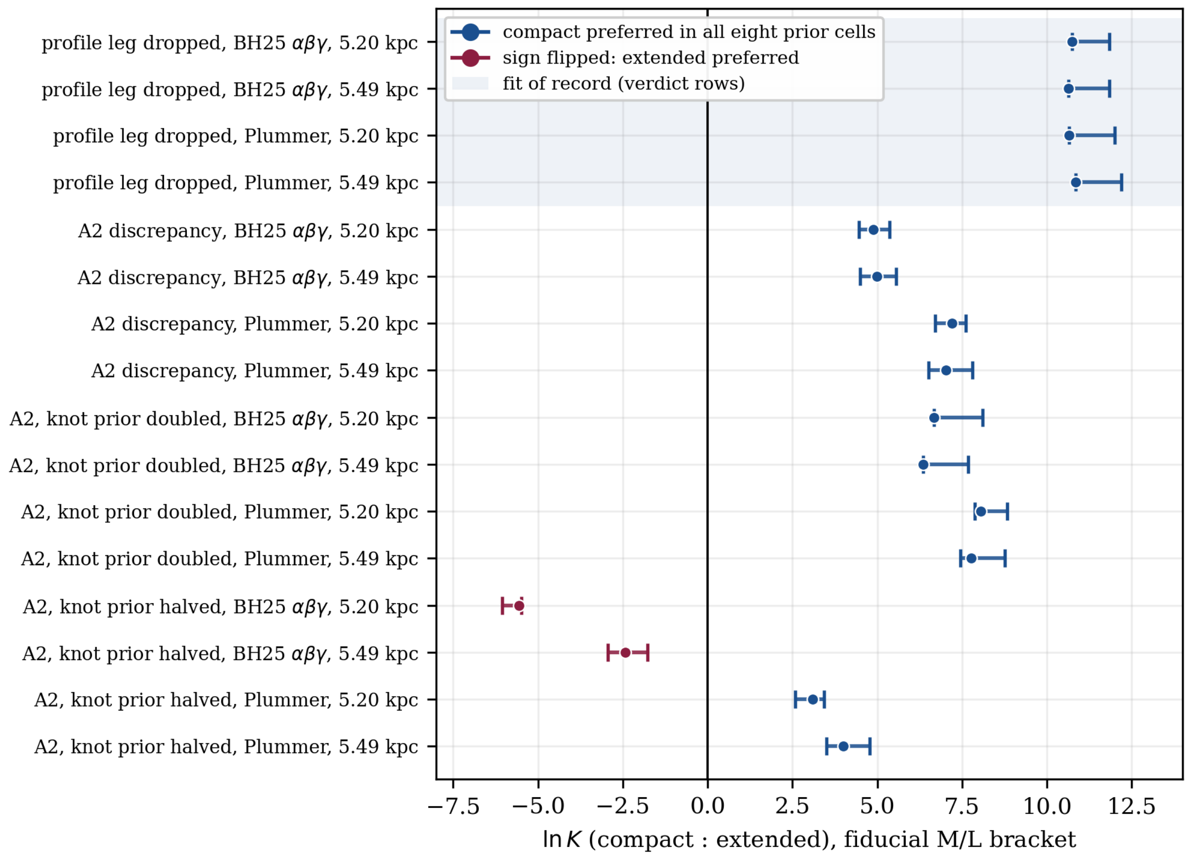

Figure 3: The sixteen pre-registered configurations of the amended rerun at the fiducial mass-to-light bracket: four error-model schemes crossed with both visible models and both distances. The dot is the full-prior-range cell and the bracket spans the eight defined prior cells of that configuration. Where Figure 2 varies one assumption at a time within one configuration, this figure is the pre-registered grid the gates were scored on. Two rows carry the minority sign, both at half the knot-prior scale on the $\alpha\beta\gamma$ visible model, and they are what gate A2-G-b failed on. The four verdict rows, with the profile leg dropped, span +10.6 to +12.2 here and are one-signed in every cell; the wider +9.96 to +12.22 quoted in Section 5.3 is the range across all three brackets. Data and script: `figs/fH_lnk_ladder.json`, `figs/fH_expand_v1.py`.

96.3–97.6 per cent, $\ln K = +10.0$ to $+12.2$ (the range over the sixteen-configuration bracket set of the amended rerun, under the reconciled mass-segregated pulsar tracer density of the released record), and a nominal, uncalibrated (per the gate of Section 5.2) 90 per cent upper limit of 2.2–$2.7\times10^4\ M_\odot$; its own decision-criterion tally passes in 31–33 of 36 prior cells. The optimum sits above the Häberle et al. (2024) firm lower bound and a factor of a few below their model-informed range, from an analysis that shares no machinery with theirs. It also inherits the assumed Bahcall–Wolf tracer cusp inside the influence radius, a profile no dataset in this analysis measures: the mass range quoted here is conditional on that cusp, and the sensitivity to its form is a reported bracket rather than a resolved question. Shifting the tail exponent by -0.25 and $+0.5$ moves $\ln K$ by $+0.15$ and -0.26 nats and the 90 per cent upper limit by $-1{,}000$ and $+1{,}500\ M_\odot$ about the values quoted here, with the sign of the evidence positive in every prior cell. The bracket is carried within the truncated-polytrope family; of the two Maxwellian alternatives the King-form lowered Maxwellian cannot be matched to the cusp kinematics at any scale while the hard-truncated isothermal can, and the latter is a named untested axis rather than an excluded one (Section 5.4).

At the first writing of this section those numbers were reported and not promoted: the injection-calibration gate under-covered on the extended alternative, and an uncalibrated instrument cannot certify the preference. A dedicated calibration campaign then traced that failure to its source: the under-coverage is caused by the profile leg itself (dropping it moves extended-injection coverage from 0.820 to 0.882 on identical realizations), the same leg the kill switch had already removed from the verdict. Scoring the gate battery on the verdict configuration, formalized by gate-scope Amendments A3 and A4 (each ruled by the gate's author, each carrying the scoring change and nothing else, each with mandatory disclosures including the named-configuration failures the re-scoping does not hide; both reproduced verbatim in Appendix B), the battery passes: 90 per cent coverage 0.8825 extended and 0.865 point at $n = 400$, null and sign-recovery checks passed non-trivially, at the fiducial mass-to-light bracket.

The content of that ruling, stated in our own voice rather than quoted: under the fit of record, with the profile leg dropped per gate A2-G-d, the fast-star and pulsar legs jointly prefer the compact configuration, and the reporting machinery is validated on that configuration by the A2 gate battery at the fiducial mass-to-light bracket; the single-leg decomposition of that preference is not quoted here, though the pulsar leg alone contributes at most the 3.0 nats of dynamic range it spans (Section 6); the pessimistic bracket is uncalibrated and carries no validated region; and the coverage check's misspecification axis is satisfied trivially by construction for the leg-dropped configuration, per Amendment A3. An earlier draft printed this as a verbatim block quote attributed to the gate's author. The analysis record (FIT3_REPORT, FLAG F-4) states that no such text exists in the repository, only a session transcript, and that a paraphrase was deliberately not reconstructed as if it were the ruling. We have therefore removed the quotation marks and the attribution and kept the content, which the amendments themselves carry in the appendix.

One rule, one table. The A2-G-c band is $[0.85, 0.95]$ on 90 per cent coverage, and the rule is that the point estimate must lie inside it; the Wilson 95 per cent interval is reported beside every entry and is not itself the decision rule. Table 2 applies that rule to every coverage check this programme produced, which an earlier draft did not: 0.795 was failed on the interval, 0.820 on the point estimate, and the result of record was passed on the point estimate with no interval printed. Under an interval rule no row in the table would pass, the result of record included, so the point rule is the one stated and it is applied uniformly. Two rows in that table matter beyond bookkeeping. The result of record, 0.8825 at $n = 400$, sits 1.82 standard errors above the floor (one-sided $p = 0.034$) with a Wilson interval whose lower end, 0.847, lies below it: the pass is real by the stated rule and it is not comfortable. And the cusp-family row at $\gamma = 1.30$, the slope the González Prieto et al. (2025) formation models measure, returns 0.660 with a Wilson interval of 0.522–0.776 entirely below the floor, 3.8 standard errors below it: scored against the gate, which no earlier draft did, the calibration fails under that published alternative cusp. The quotable statement of

Table 2: Every 90 per cent coverage check of the Paper H programme, scored against the A2-G-c band $[0.85, 0.95]$ by one rule: the point estimate must lie inside the band. Wilson 95 per cent intervals and the one-sided z against the 0.85 floor are reported for every row. The gate-G1 rows predate the band and are scored here for completeness; their original pass is voided and superseded.

Check	Configuration	Cusp	Bracket	n	Coverage	Wilson 95%	z	Verdict
G1 point	all legs	1.75	fiducial	50	0.840	0.715–0.917	−0.20	FAIL
G1 extended	all legs	1.75	fiducial	50	0.920	0.812–0.968	+1.39	PASS
A2-G-c ext.	all legs	1.75	fiducial	200	0.795	0.734–0.845	−2.18	FAIL
A2-G-c ext.	all legs	1.75	fiducial	400	0.820	0.779–0.855	−1.68	FAIL
A2-G-c ext.	all legs	1.75	pessimistic	400	0.680	0.633–0.724	−9.52	FAIL
A2-G-c ext.	leg-dropped	1.75	fiducial	400	0.8825	0.847–0.910	+1.82	PASS
A2-G-c point	leg-dropped	1.75	fiducial	400	0.865	0.828–0.895	+0.84	PASS
cusp family	all legs	1.75	fiducial	50	0.840	0.715–0.917	−0.20	FAIL
cusp family	all legs	1.30	fiducial	50	0.660	0.522–0.776	−3.76	FAIL

this section is therefore conditional on $\gamma \approx 1.75$ specifically, not on the tracer cusp generically. Data and script: `figs/fH_calibration_table.json`, `h/cal2/calibration_table.py`.

The ruling reproduced above is dated, and the campaign continued past it. Amendment A5 records what it found: a mechanism, without a repair. The pessimistic bracket's under-coverage is mean displacement of the recovered mass; the region width is unchanged. The recovered mass compensates for the visible mass the bracket removes, at an effective matching radius of 2.2 pc, while the 90 per cent region stays two grid cells wide at every bracket. Generated at the normalization being fitted, the three leg-dropped diagonals agree in coverage to 0.004. What costs coverage is the offset between generating and fitted normalization: on the extended injection a one-step offset costs up to 7 points and a two-step offset up to 23, monotone in the offset; the point injection, whose compact dark component cannot compensate, is close to flat. Absolute leg-dropped coverage sits near 0.87 at every bracket, a property of the configuration, and the joint cube's under-coverage is bracket-independent and stays embargoed.

Grid refinement as a repair for the all-legs under-coverage is untested. The referee prediction that it would not help is on record and is stated here as a prediction rather than a result.

The coverage statement carries the same tracer-cusp condition as the mass, and the size of that condition is now measured. Rerunning gate G1's point injection ($4 \times 10^4\, M_\odot$ at $a = 10^{-4}$ pc, 50 realizations per family, the full three-leg likelihood at the fiducial bracket, r_a marginalized) with the generation-side tracer cusp varied and the fit always assuming the shipped Bahcall–Wolf form, 90 per cent coverage is 0.840 ± 0.052 under the shipped $\gamma = 1.75$ cusp, consistent at 0.5σ with the 0.86–0.88 quoted above. Generating instead at the shallower $\gamma = 1.30$ stellar-cusp slope the González Prieto et al. (2025) models measure, coverage degrades to 0.660 ± 0.067 and the median recovered mass carries a -12.3 per cent bias against -5.1 per cent under the shipped cusp. Scored against A2-G-c by the rule of Table 2, that is a FAIL: 0.660 lies 3.8 standard errors below the 0.85 floor and its Wilson interval, 0.522–0.776, lies entirely below the band. Note also that the shipped-cusp coverage number cannot test this, because generation and fitting apply the same cusp; the $\gamma = 1.30$ run is the only real test of the assumption and it does not pass. The displacement is one-sided: under a shallower true cusp the recovered mass moves down, so the mass quoted here is an underestimate under this misspecification, while the sign of the preference and the point-mass-edge location of the mode are unchanged. Data and script: `figs/fH_calc7_cuspfam.json`, `figs/fH_calc7_cuspfam.py`.

Table 3 carries the calibrated regions. $\ln K(\text{compact} : \text{extended})$ is positive in every one of the 96 defined prior cells of the fit of record, ranging from +9.96 to +12.22 across both visible models and both distances; the prior-scale rows of gate A2-G-b are vacated by the leg drop (Amendment A4), and on the all-configuration set including profile-leg fits the same gate had failed (52 of 288 minority-sign cells), the verdict-control finding the leg drop remedies. Forty-eight of those fifty-two cells lie in the halved knot-prior rows, sixteen on the $\alpha\beta\gamma$ visible model at each distance (all eight pessimistic and all eight fiducial cells) and eight on the Plummer model at each distance (pessimistic cells only); the remaining four lie in the primary $\alpha\beta\gamma$ row at 5.20 kpc. The doubled knot-prior rows carry none, and no cell of the leg-dropped fit of record carries the minority sign. The full configuration-by-cell matrix ships with the analysis chain (`flagd/results/fit2_msp`; machine-readable mirror `figs/fH_calc7_records.json`). On the fit of record the pre-registered joint-consistency criterion (criterion 3, the overlap test) is met in 31–33 of 36 prior cells per configuration, and criterion 4 does not fire on that test. The two posterior-predictive checks are not both clean at this scope: the acceleration-extrema check passes in all cells, while the count check passes in 24 of 36 cells under the robust-five census and in 0 of 36 under the all-seven census (Section 6 details both). This is a statement of consistency between the fast-star and pulsar legs, not a resolution of the profile-leg tension, which gate A2-G-d removed from the verdict.

The overlap component is weakly discriminating and its power should be read off rather than assumed. Across the four verdict configurations the fast-star 90 per cent region occupies 5–52 of the 1,891 grid cells and the pulsar region 370–1,420, so the test asks whether a small region falls inside a large one and will usually say yes. It is not, however, automatic, and the cells where it fails are informative: every one of the 3–5 misses per configuration lies in the full 10^3–$10^6\,M_\odot$ mass sweep, and all but two lie in the pessimistic mass-to-light bracket. The pulsar leg's mild preference is for the top of the mass range and the fast-star optimum is at $2\times10^4\,M_\odot$; when the sweep spans all three decades the two regions concentrate in different ones and separate, while every sub-range capped at 10^5 or $3\times10^5\,M_\odot$ overlaps. So criterion 3 is a real but low-power test, and what it reports is the absence of gross tension rather than agreement on a mass. The whole statement is additionally conditional on the TRAPUM reduction of the pulsar bounds: Section 6 shows criterion 3 failing in 36 of 36 cells under the Dai reduction of the same pulsars. The fit of record rests on five stars: five well-measured stars with a modeled selection function, a validated instrument, and an assumed tracer cusp.

Figure 4 shows the same four regions in the plane they live in, which the table can only report as intervals. The regions are one-sided: every one of them abuts the lower a grid edge, the point-mass limit of the parameterization, so the data bound the scale radius from above and not from below (the exact point-mass rows of Table 3 price that limit: restricting a to zero raises the evidence by 0.83–0.87 nats relative to the continuous model), and a reader should not read the quoted a extents as intervals with two ends. The profile-likelihood region also contains the HPD region in all four configurations, so the disclosed secondary estimator is the wider of the two everywhere here.

The scope on which each thing is scored. Two scopes are in use in this paper and the rule that assigns them needs stating, because Amendments A3 and A4 moved the gates from one to the other and did not move the criteria. The quotability gates A2-G-a, -b and -c are scored on the verdict configuration, the leg-dropped fit of record: they exist to certify the instrument that produces a quoted number, and the instrument that produces the quoted number is the leg-dropped one. The pre-registered decision criteria 1–4 are scored on the full sixteen-configuration A2 set, because they answer the question the pre-registration asked, which is whether the data as a whole decide between a compact and an extended dark component, and that question is not made easier by deleting the dataset that disagrees.

That asymmetry has a consequence the analysis record raises and an earlier draft of this paper passed over. On the verdict configuration alone, criterion 3 (overlap) is met in 31–33 of 36 prior cells, with the acceleration-extrema check passing in all cells and the count check passing in 24 of 36 (robust-five census) or 0 of 36 (all-seven census); criterion 4 does not fire on the overlap test there; FIT3_REPORT records this and files it as FLAG F-3, "not decided here". We adjudicate it here, in favour of the all-configuration scope,

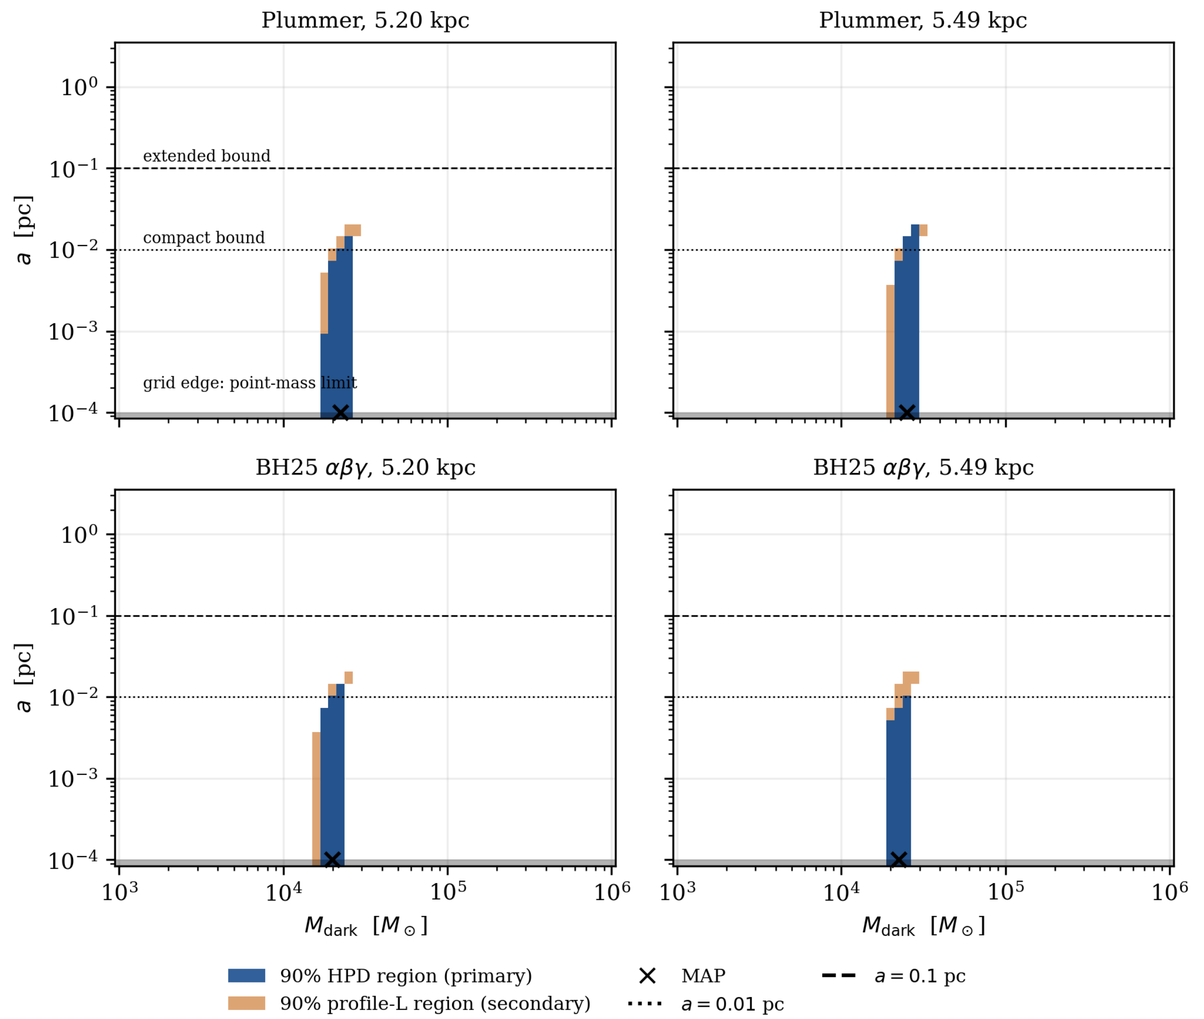

Figure 4: The fit of record in the $(M_{\rm dark}, a)$ plane, one panel per verdict configuration, at the fiducial mass-to-light bracket over the full pre-registered prior grid. Filled cells are the 90 per cent HPD region, the primary estimator; the lighter cells are the 90 per cent profile-likelihood region, the disclosed secondary. The HPD region occupies 39 to 49 of the 1,891 grid cells and the profile-likelihood region 45 to 58, and no cell of either lies above $a = 0.1$ pc in any configuration. The shaded strip along the bottom is the $a = 10^{-4}$ pc grid edge, which every region touches. Regions were recomputed from the stored cubes and checked against the released table before plotting. Data and script: `figs/fH_posterior_plane.json`, `figs/fH_expand_v1.py`.

Table 3: The fit of record: calibrated 90 per cent regions in ($M_{\rm dark}, a$), fiducial M/L bracket. Under self-consistent generation, with the mocks drawn at the normalization being fitted, the pessimistic bracket is calibrated equivalently to the fiducial one in the verdict configuration (pooled diagonals 699 of 800 each; Fisher $p = 0.76$ on the point and 0.74 on the extended injection; Amendment A5). Its rows remain labelled sensitivity rows outside the quotable set, and the joint cube stays uncalibrated at every bracket. HPD is the primary estimator; the profile-likelihood (PL) region is the disclosed secondary. The two interval estimators and the upper limit answer different questions and can order either way: HPD90 is the shortest region holding 90 per cent of the marginal posterior, PL90 is the region above a fixed drop in profile likelihood, and the 90 per cent upper limit is the one-sided quantile with 90 per cent of the posterior below it. A one-sided quantile printed below an HPD upper edge is the ordinary behaviour of a right-skewed marginal, not an inconsistency between the columns. The restricted rows evaluate the hypothesis the data prefer, $a = 0$ exactly, under identical priors; their $\ln K$ relative to the continuous optimum is +0.83 to +0.87 in all four configurations, so the continuous parameterization spends about a nat of Occam factor on a scale radius the data do not use. All regions abut the $a = 10^{-4}$ pc grid edge and are one-sided in a.

Configuration	HPD90 $M_{\rm dark}$ ($10^4\,M_\odot$)	PL90 $M_{\rm dark}$ ($10^4\,M_\odot$)	a extent (pc)	$M_{\rm dark}$ 90% UL
Plummer, 5.20 kpc	1.78–2.51	1.78–2.82	≤ 0.012	2.42×10^4
Plummer, 5.49 kpc	2.24–2.82	2.00–3.16	≤ 0.017	2.71×10^4
BH25 $\alpha\beta\gamma$, 5.20 kpc	1.78–2.24	1.59–2.51	≤ 0.012	2.16×10^4
BH25 $\alpha\beta\gamma$, 5.49 kpc	2.00–2.51	2.00–2.82	≤ 0.009	2.42×10^4
Exact point-mass restricted fits ($a \equiv 0$), same priors:				
Plummer, 5.20 kpc	1.78–2.51	—	$a \equiv 0$	2.38×10^4
Plummer, 5.49 kpc	2.24–2.82	—	$a \equiv 0$	2.68×10^4
BH25 $\alpha\beta\gamma$, 5.20 kpc	1.78–2.24	—	$a \equiv 0$	2.13×10^4
BH25 $\alpha\beta\gamma$, 5.49 kpc	2.00–2.51	—	$a \equiv 0$	2.40×10^4

as a deliberate conservatism rather than a mechanical trigger: on the full A2 set criterion 4 fired in the majority of configurations and the gate battery as originally scoped failed, and a criterion-4 verdict scored on the configuration that survives the gates would be a verdict about the surviving probe, not about the data. Two further facts support retaining it. The verdict configuration's own criterion-3 tally is not clean, missing in 3–5 cells of 36. And the pulsar reduction of Section 6 flips the verdict configuration's own answer: under the Dai reduction of the same five pulsars, criterion 3 fails in 36 of 36 cells and criterion 4 fires on the fit of record as well. The multi-dataset question of the title therefore remains undecided; what the calibration campaign changed is that the surviving probe's preference is now a calibrated statement rather than a reported curiosity.

5.4 Speed-tail and tracer-cusp sensitivity of the fit of record

The fast-star likelihood evaluates each star against the high-velocity tail of the local speed distribution, so the verdict could in principle ride on that family rather than on the stars. Two tests bound the dependence. Holding the Bahcall–Wolf tracer density, the escape-speed truncation, and the selection function fixed, we shift the polytrope exponent by −0.25 and +0.5, a heavier and a lighter tail than the shipped law under the same Jeans second-moment matching: $\ln K$ moves from +10.85 to +11.00 and +10.59, the posterior mass in the compact region stays within 0.09 points of 96.3 per cent, the mass mode does not leave the point-mass edge, and $\ln K$ is positive in every one of the 24 defined prior cells under all three laws. The same shifts move the 90 per cent upper limit by −1,000 and +1,500 $M_\odot$ about the quoted value. Repeating the same exponent shifts on the $\alpha\beta\gamma$ model at 5.49 kpc reproduces that result on the second visible base: $\ln K$ moves by +0.14 and −0.26 nats about +10.64, the compact posterior mass stays within 0.13 points of 97.33 per cent, the sign of the evidence is positive in every cell of every family, and the mass mode stays at the point-mass edge throughout. Data: `figs/fH_calc7_tail_families_abg5494.json`.

One of the two Maxwellian-based families can be matched to this model's kinematics and one cannot, and an earlier draft of this paper excluded both on a cap that is wrong for each. The Bahcall–Wolf consistency relation $\sigma_{1d}^2 = \Psi/(1+\gamma)$ used throughout demands $\langle v^2\rangle = 3\sigma_{1d}^2 = \frac{3}{2.75}\Psi = \frac{3}{5.5}\,v_{\rm esc}^2 = 0.5\overline{45}\,v_{\rm esc}^2$ (0.545, rounded here and below to 0.55; the rounding is immaterial to both verdicts) inside the influence radius, where $\gamma + 1 = 2.75$. The truncated second moment of each family was stated as capping at $v_{\rm esc}^2/2$. Recomputed by numerical sweep in the truncation-scale ratio, with the limits confirmed analytically, the supremum is $\frac{3}{5}v_{\rm esc}^2 = 0.600\,v_{\rm esc}^2$ for the hard-truncated isothermal and $\frac{3}{7}v_{\rm esc}^2 = 0.4286\,v_{\rm esc}^2$ for the King-form lowered Maxwellian, each approached as the scale runs to infinity. The King family is therefore genuinely excluded, $0.4286 < 0.55$. The hard-truncated isothermal is not: it meets the requirement at $\sigma = 0.839\,v_{\rm esc}$, and the claim that the matching problem has no solution is withdrawn. We have not carried it as a third tail family. That would require rerunning the fast-star leg, and the Round-8 revision was authorized for exactly one rerun, spent on the pulsar reduction of Section 6; the hard-truncated isothermal is named here as an admissible, untested robustness axis rather than a closed one, and the disclosed tracer-cusp bracket is the pair of exponent-shift numbers within the truncated-polytrope family only. Data and scripts: `figs/fH_calc7_tail_families.json`, `figs/fH_calc7_tail_families.py`, `figs/fH_calc7_tail_caps.json`, `h/calc7/fH_calc7_tail_caps.py`.

6 Diagnosis: locating the discriminating power

The dispersion profile cannot arbitrate this question at the radii that carry its weight, and there the statement is permanent at this precision. Two facts compose the argument, and a third that earlier drafts used has been withdrawn. The statistical errors (0.03–0.05 km s^{-1}) sit a factor 10–40 below the leg's own physics-derived systematic budget (0.5–1.6 km s^{-1}, the range from which A2's flat 1.0 km s^{-1} trigger was conservatively set); and the residual against every tested model is component-differential, so no radial error term can represent it. The withdrawn third fact was the claim that the verdict-dominant bins at $50''$–$250''$ carry a dark-component contribution under 10^{-3} of the enclosed mass. Table 1 recomputes it: the true fraction at $100''$ runs from 0.0076 at the lowest prior corner to 0.635 at the extended solution, so those bins carry substantial signal for the hypothesis the leg would be needed to test. What disqualifies the leg is that the signal is buried under a systematic budget it cannot separate from, not that the signal is absent. The statement is radius-resolved: it is terminal for the verdict-dominant bins beyond $\sim 50''$, while the innermost bins ($\lesssim 10''$) sit at signal-to-systematics near unity and are marginal rather than foreclosed. Deeper catalogs leave the first fact untouched and make the second worse at the verdict-dominant radii; only the inner bins can benefit, and they are a different, smaller measurement. The profile's proper role in this problem is visible-model calibration, which it performs well: the calibrated models reproduce the catalogue central escape velocity to 2–3 per cent without it being an input, though the catalogue value derives from models fit to overlapping kinematic data, so this is a consistency check rather than an independent one.

The pulsar accelerations are nearly uninformative. Read correctly as censored one-sided bounds, seven usable accelerations span 3.0 nats across the entire $(M_{\rm dark}, a)$ plane; alone, they yield no informative bound within the pre-registered $[10^3, 10^6]\,M_\odot$ grid, with the 90 per cent quantile running 8.3–$8.6\times10^5\,M_\odot$ across the three kernels tested below and the posterior rising toward the prior's upper edge rather than falling.

That limit disagrees with the source paper's own, from the same data, and the disagreement needs stating. Colom i Bernadich et al. (2026) report that their likelihood peaks at $M_{\rm BH} = 0$ and place an upper limit of $10^5\,M_\odot$ at 90 per cent confidence. Our leg, from seven of the same pulsars, returns $8.4\times10^5\,M_\odot$:

8.4 times weaker, and rising rather than falling in mass. FIT_REPORT attributed the difference to the line-of-sight marginalisation weighting, which concentrates weight near $z = 0$ where $a_{\rm LOS}$ vanishes. We tested that attribution in the authorized rerun and it does not survive: re-evaluating the leg under the fiducial mass-segregated millisecond-pulsar kernel, under the light-star tracer density carrying the Bahcall–Wolf cusp, and under a flat kernel with no central concentration at all returns 8.35, 8.26 and $8.60 \times 10^5\, M_\odot$, all rising in mass. The attribution is withdrawn. What the difference does reflect is that the two analyses ask different questions of the same table. Colom i Bernadich et al. (2026) fit a King model with a point mass against a fixed cluster potential; we read every entry as the one-sided bound it is, with the potential varying over the prior grid, and four of the seven bounds are positive-signed and satisfiable by any mass in the grid. The three negative-signed bounds (pulsars B, D, K) are the only ones that require the model to *reach* a value, and a larger central mass reaches them more easily, which is the direction of the mild rise. Pulsar C, dropped here for an ambiguous table entry, is retained by them. We take their limit as the better-constrained statement about a point mass and ours as a statement about what the bounds alone exclude across the full $(M_{\rm dark}, a)$ plane, and we note the consequence for the impasse this paper is about: their $10^5\, M_\odot$ point-mass limit sits below the 2–3 $\times\, 10^5\, M_\odot$ extended component Bañares-Hernández et al. (2025) favour, so the two are not in conflict, while it sits above Bañares-Hernández et al.'s own 6,000 $M_\odot$ point-mass cap and is therefore the weaker of the two constraints on that quantity.

The reduction chain, checked further than before. An earlier draft claimed to have "verified the published reduction chain" by recomputing pulsar A's bound from its P, $\dot{P}$ and the tabulated Shklovskii and Galactic terms. That verifies one subtraction while re-using the two terms, and we downgrade the claim accordingly: it verifies Equation (2)'s subtraction. Recomputing the Shklovskii term itself, $\dot{P}_{\rm Sh} = \mu^2 dP/c$ from each pulsar's tabulated proper motion at 5.494 kpc, reproduces the published value for four of the seven (B, E, G, K, all to three digits) and does not for three: pulsar A at a factor 2.03, pulsar D at 0.84 and pulsar H at 0.80. Pulsar H carries a second inconsistency: the Galactic term implied by its tabulated $\dot{P}_{\rm G}$ is $-1.04 \times 10^{-9}\,{\rm m\,s^{-2}}$, ten times the -1.036 to $-1.041 \times 10^{-10}\,{\rm m\,s^{-2}}$ that the other six imply for a quantity that must be common to all cluster members, and its published bound of $4.58 \times 10^{-9}\,{\rm m\,s^{-2}}$ is reproduced by our terms (4.573×10^{-9}) and not by its own tabulated ones (5.44×10^{-9}). The bounds themselves are what enters our likelihood, and six of the seven reproduce from the terms as printed; we flag the three Shklovskii discrepancies and the pulsar-H Galactic term as items for the timing-level reconciliation of Section 8(ii), not as corrections we are entitled to make to another group's table. Data and script: `figs/fH_shklovskii_check.json`, `h/fit/shklovskii_check.py`.

The leg's real contribution is a data-quality discovery: the pre-registered posterior-predictive check on the anomalous negative accelerations discriminates between the Dai et al. (2023) and TRAPUM reductions of the same pulsars (under the Dai values, pulsar D's bound is unreachable by any model in the plane), rather than between dark-component geometries. The two reductions differ systematically in dispersion measure by 0.02–0.04 pc cm^{-3} across all five overlapping pulsars, far beyond either's formal errors. Table 4 ships the comparison. The offsets run 0.024–0.037 pc cm^{-3}, four positive and one negative, and range from 4.7σ to 73σ of the combined formal error; pulsar B's spin-down itself agrees between the reductions ($\dot{P} = -5.433(4)$ against $-5.4336(6) \times 10^{-20}$), which places the discrepancy in the dispersion measures and in what the acceleration chain derives from them. Until that discrepancy is resolved at the timing level, no dynamical analysis should treat either reduction's accelerations as settled input; we froze on TRAPUM as the newer, longer-baseline solution and carry Dai as the labelled sensitivity.

How much that choice is worth, on the fit of record. FIT-1 recorded that the Dai substitution inflates the acceleration leg's dynamic range from 3.0 to 1,379 nats and drops the criterion-3 overlap count from 36 cells to 12, and dismissed it on the grounds that "the joint posterior is unmoved because the profile leg dominates both". That leg has since been retired, so the dismissal was stale, and we spent the one rerun this revision was authorized for on repairing it. Substituting the Dai values on the fit of record (fast stars plus pulsars, profile leg dropped, all four verdict configurations, the same deterministic pipeline) reverses it. The pulsar leg spans 1,379 nats over the four pulsars Dai and TRAPUM share, and it controls the answer: $\ln K$ moves from +10.6 to +10.8 under TRAPUM to −21.1, −25.6, −34.8 and −39.4 under Dai; the maximum-posterior cell moves from (2.0–$2.5 \times 10^4\, M_\odot,\ 10^{-4}$ pc) to (7.1–$8.9 \times 10^5\, M_\odot,\ 2.1$ pc); the 90 per cent upper limit moves from 2.2–2.7×10^4 to 7.4–$8.6 \times 10^5\, M_\odot$; and criterion 3 is met in 0 of 36 prior cells in every configuration, so criterion 4 fires on the verdict configuration as well. The compact preference of Section 5.3 is therefore conditional on the TRAPUM reduction in the strongest sense available: the two published reductions of the same five pulsars put the fit of record on opposite sides of the question. We do not read the Dai row as a competing result. Its own posterior-predictive check fails (under the Dai values pulsar D's bound is unreachable by any model in the plane), which is why the pre-registration's check flagged that reduction in the first place, and a leg whose 1,379 nats come from an unreachable bound is registering a data problem rather than a mass. We read it as the size of the assumption, and it is larger than every other assumption priced in this paper. Data and script: `h/fit4/results/r8_rerun.json`, `h/fit4/run_r8_rerun.py`. Figure 5 puts the leg's weakness and its residual content in one panel: at every pulsar's projected radius, both the compact and the extended model can reach an acceleration below the bound, so no bound is violated by either hypothesis, and the two envelopes separate by less than the spread of the bounds themselves outside the innermost half-parsec, where no timed pulsar sits.

Table 4: Dispersion measures for the five overlapping pulsars, Dai et al. (2023) against TRAPUM (Colom i Bernadich et al., 2026), as published. Δ = TRAPUM − Dai; the significance column divides Δ by the combined formal error. Data: `figs/fH_calc7_records.json`.

Pulsar	Dai DM	TRAPUM DM	Δ	combined σ	significance
A	100.2899(7)	100.32670(40)	+0.0368	8×10^{-4}	46σ
B	100.2500(10)	100.28060(80)	+0.0306	1.3×10^{-3}	24σ
C	100.6400(50)	100.66430(130)	+0.0243	5.2×10^{-3}	4.7σ
D	96.5180(20)	96.54620(90)	+0.0282	2.2×10^{-3}	13σ
E	94.3690(4)	94.33967(5)	−0.0293	4×10^{-4}	73σ

The jerks are floor-limited, not precision-limited (Section 4.2): a 1-stable nearest-neighbour floor that averaging cannot beat, discrimination requiring 5–10 times the current pulsar census, and the currently measured $\ddot{\nu}$ set inconsistent with cluster origin altogether. With the corrected ξ, published jerk-based significance estimates should be revisited by their authors; the correction is arithmetically small and directionally unfavorable.

The fast stars are the live probe. They are few, but their leg passed every check run against it: stable optimum across every bracket, contamination normalization reproducing the source paper's own expected foreground count, agreement with an independent analysis chain, and clean single-leg injection recovery. The second of those is a check on the contamination term alone, and we no longer list it as a check the selection function passed: the number is a density times an area, $0.0026\,\mathrm{arcsec}^{-2} \times \pi(3'')^2 = 0.0735$, and it exercises neither the completeness nor the proper-motion threshold. Of the two named obstacles, calibration on the extended alternative has since been solved for the verdict configuration at the fiducial bracket (the campaign of Section 5.3); what remains is sample size.

What the leg conditions on. Four assumptions carry it, and they are collected here rather than left in passing clauses. *The tracer cusp.* A Bahcall–Wolf $\gamma = 1.75$ profile inside the influence radius, measured by no dataset in this analysis, applied identically in generation and fitting, so the shipped coverage number

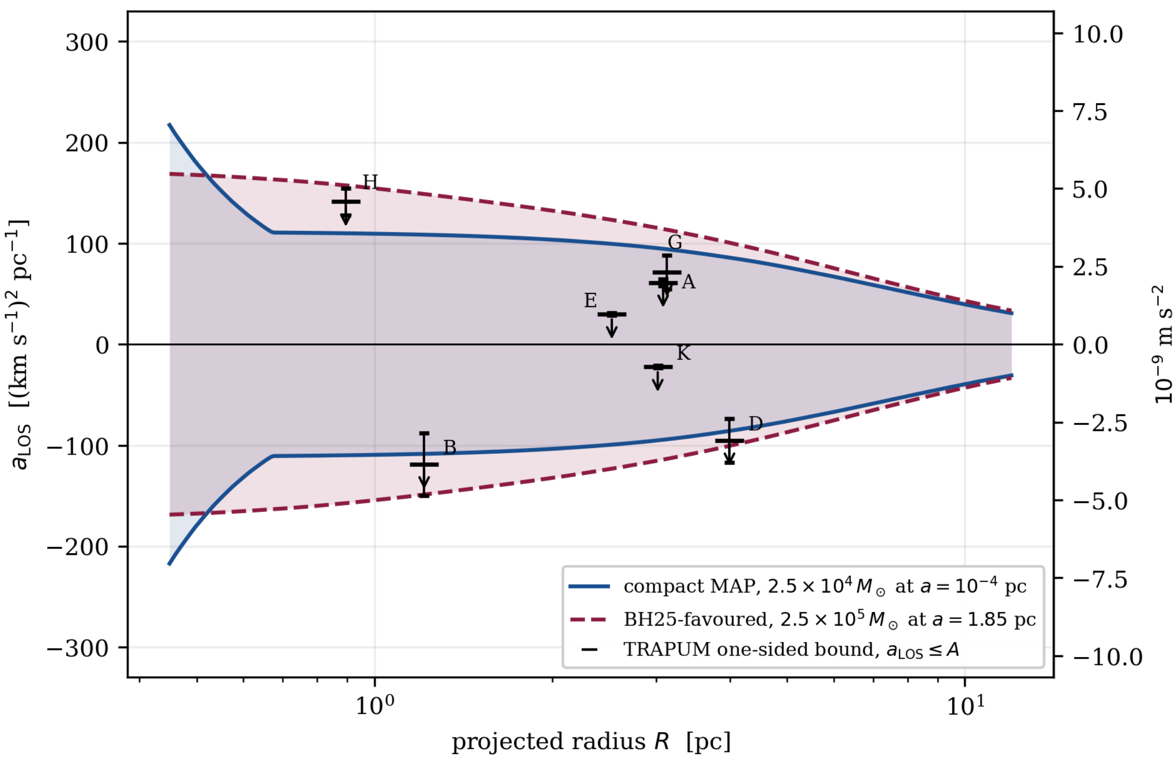

Figure 5: The seven usable TRAPUM line-of-sight accelerations against projected radius, each a one-sided bound with the arrow marking the permitted direction, together with the largest $|a_{\mathrm{LOS}}|$ the model can produce along the line of sight at that radius for a compact and for an extended dark component at the fiducial bracket. Three bounds are negative-signed (B, D, K) and those are the ones a model has to reach. The envelopes are computed by the same line-of-sight quadrature the likelihood uses, and reproduce the released posterior-predictive extrema of the record cell to a relative agreement better than 10^{-8}. The compact and extended curves cross at 0.52 pc, inside the innermost timed pulsar at 0.89 pc, which is the geometric statement behind the leg's 3 nats of dynamic range. Data and script: `figs/fH_pulsar_bounds.json`, `figs/fH_expand_v1.py`.

cannot test it; at the shallower published $\gamma = 1.30$ the calibration gate fails and the recovered mass is biased low by 12.3 per cent (Table 2). That FAIL was established on the all-legs G1 injection machinery ($n = 50$, Section 5.2), not on the leg-dropped fit of record; no leg-dropped $\gamma = 1.30$ injection run exists, and the shipped-cusp run applies the same cusp in generation and fitting, so it cannot test its own assumption. The quotable mass and coverage of Section 5.3 are therefore conditional on $\gamma \approx 1.75$ as assumed, not as measured; leg-dropped coverage under cusp misspecification is an open, named axis rather than a scored one. *The completeness.* A single radius-independent factor 0.652, measured over $r < 90''$ and applied to a sample confined to $r < 3''$ in the most crowded field in the Galaxy; the extraction record already flags those as two different radii for two different purposes, and the fit uses one constant outside the radial integral anyway. HST completeness in this core falls inward, so the true inner value is a bound from above; a lower inner completeness requires more intrinsic fast stars to yield five detections, hence more mass, so the quoted $M_{\rm dark}$ is biased low by an amount this analysis does not measure and the count check is correspondingly mis-normalized. Bracketing it with a radial ramp is a rerun, and is named in Section 8 rather than done here. *The membership of two candidates.* The result assumes Häberle et al.'s two excluded candidates are non-members. Under the alternative the optimum barely moves ($2.51 \times 10^4\ M_\odot$) and $\ln K$ shifts by 4.6 nats, but the count posterior-predictive check then fails in all 36 cells rather than 12. That is a failure of the preferred model to produce the observed count, not a reason the data are wrong, and it is reported here as the conditionality it is; the robust-five census already fails that check in 12 of 36 cells, so it is not clean as used either. *The tail family.* The truncated polytrope, bracketed by exponent shifts, with the hard-truncated isothermal admissible and untested (Section 5.4). Figure 6 shows the leg. Its geometry is worth stating explicitly, because the figure invites a comparison it does not licence: the curves are escape velocities at three-dimensional radius plotted at each star's projected radius, which overstates $v_{\rm esc}$ at the star's true radius, while the plotted speeds are two-dimensional and understate the space velocity. The likelihood does not compare the two the way the eye does; it evaluates each star against the local high-velocity tail of the speed distribution with the selection function folded in.

Figure 7 re-plots the same seven stars and curves with the selection threshold marked as a line, so the margin between each star's measured speed and the three escape-velocity models can be read directly.

7 The formation-physics overlay

The G3 map asks, independently of all kinematics, which regions of the $(M_{\rm dark}, a)$ plane published formation and retention physics can populate. Its two headline statements frame the impasse from the opposite side. The compact region implied by the fast stars (8×10^3–$5 \times 10^4\ M_\odot$, $a < 0.01$ pc) is buildable under pessimistic retention throughout: it is what the merger-growth models of González Prieto et al. (2025) produce, with recoil we estimate from their capture-channel mass ratios to lie four orders of magnitude below the escape velocity. The recoil figure is our estimate under standard kick scaling, not a number González Prieto et al. (2025) print. The extended region favored by Bañares-Hernández et al. (2025) is the hard one: as a pure remnant population, 2.5–$3 \times 10^5\ M_\odot$ exceeds the retained black-hole mass that Dickson et al. (2024) infer for ω Cen even at the optimistic edge of their credible range, the remainder of the region admits an all-remnant reading only at the upper edge of the published retention credible range, and every hybrid configuration that reaches the upper masses requires an embedded IMBH of 1.5–$6.6 \times 10^4\ M_\odot$, which the extended solution's own 3σ point-mass cap excludes. The upper half of the extended region is reachable, on published physics, only through configurations the extended fit itself rules out.

Figure 8 draws the map with the fit of record's regions on top of it. The four HPD regions fall wholly inside the region the published growth models populate under pessimistic retention, while the region the

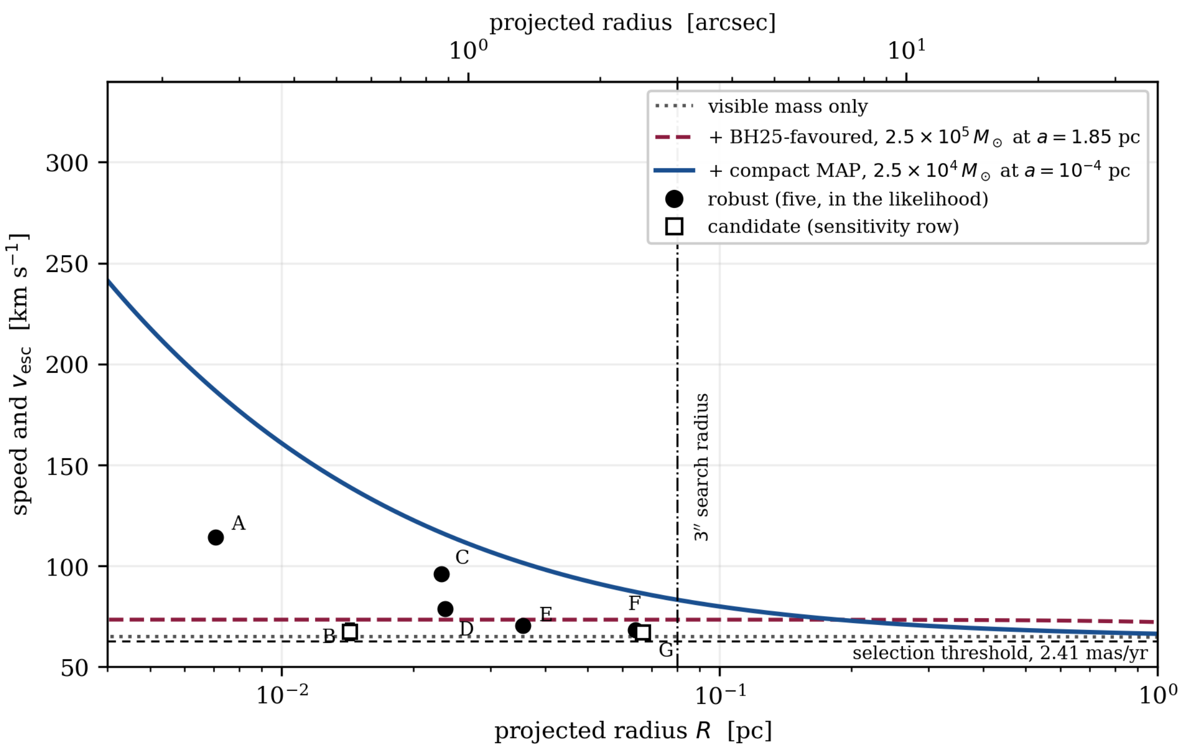

Figure 6: The seven Häberle et al. (2024) fast stars at 5.49 kpc against central escape-velocity curves for the visible mass alone, the visible mass plus the compact optimum of the fit of record, and the visible mass plus a representative point of the extended region favoured by Bañares-Hernández et al. (2025). Filled circles are the five robust stars that enter the likelihood; open squares are the two candidates carried as a sensitivity row. Only the compact component lifts $v_{\rm esc}$ appreciably inside 0.05 pc, which is where five of the seven stars lie, and this radial concentration of the signal is what makes the leg a compact-mass probe rather than a total-mass probe. The visible-only curve reproduces the Baumgardt and Hilker (2018) catalogue central escape velocity to 4.6 per cent without being fitted to it. Data and script: `figs/fH_fast_stars.json`, `figs/fH_expand_v1.py`.

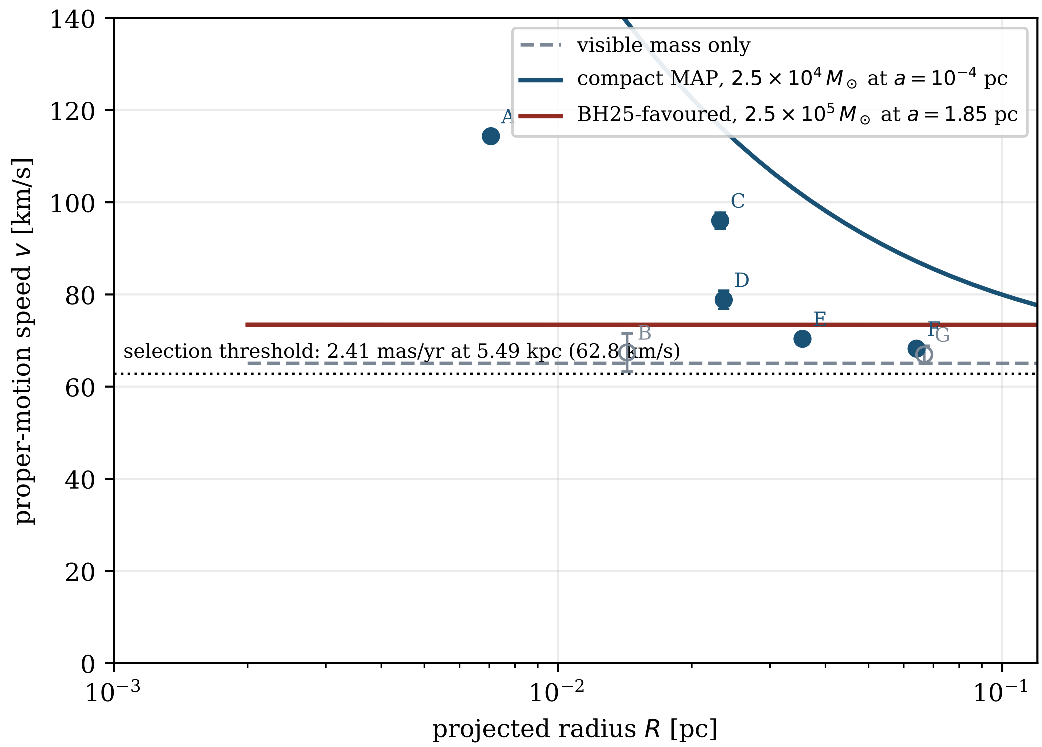

Figure 7: The fast-star overlay: the seven Häberle et al. (2024) stars (filled blue, the five robust members of the likelihood; open grey, the two candidates) against the three escape-velocity curves of Figure 6 — visible mass only (dashed grey), the visible mass plus the compact MAP ($2.5 \times 10^4\ M_\odot$ at $a \to 0$; blue), and the visible mass plus the Bañares-Hernández et al. (2025)-favoured extended component ($2.5 \times 10^5\ M_\odot$ at $a = 1.85$ pc; red) — with the selection threshold (2.41 mas/yr at 5.49 kpc, 62.8 km s^{-1}, from Baumgardt and Hilker 2018's catalogued central velocity dispersion, distance-scaled; Paper G's X-ray population forecast instead uses the Yin et al. (2024) scaling-relation escape velocity, 62.2 km s^{-1}, a different convention agreeing to ~1 per cent) marked as the dotted line. Curves, stars, and threshold are taken from the released record `figs/fH_fast_stars.json` without recomputation; the same projected-radius caveat as Figure 6 applies.

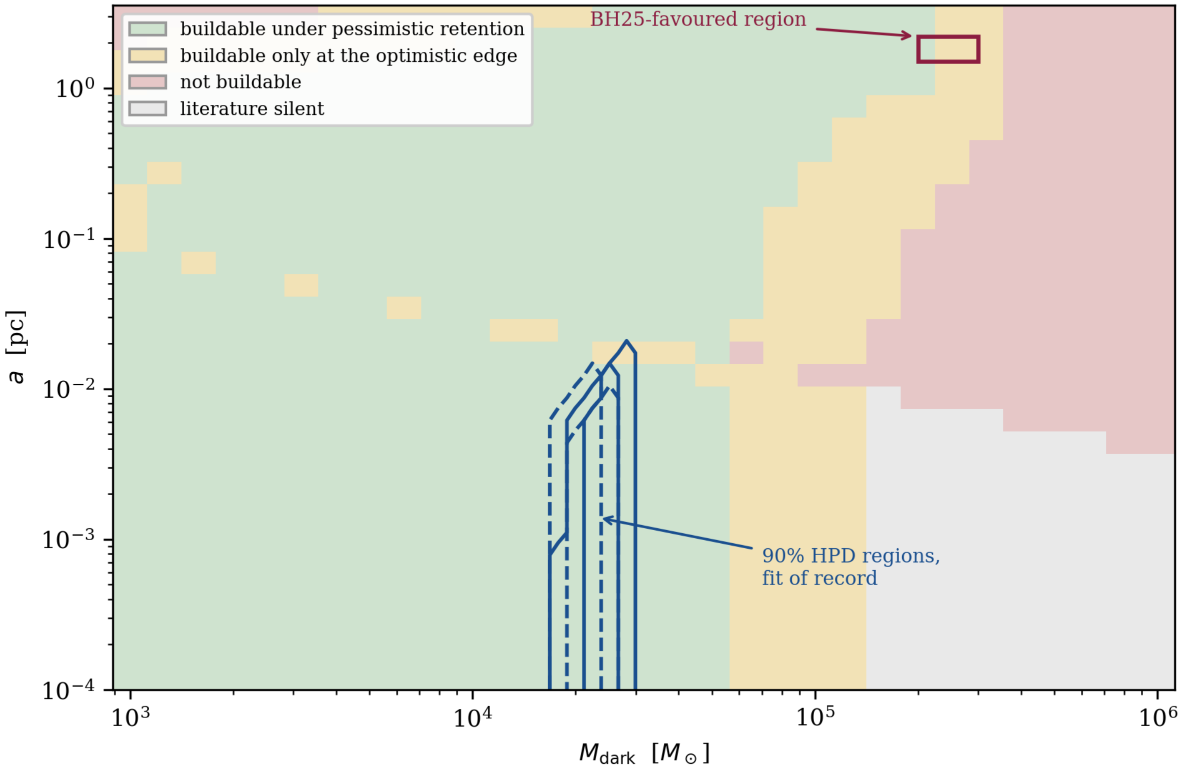

Figure 8: The G3 buildability map with the four 90 per cent HPD regions of the fit of record overlaid (solid outlines, Plummer visible model; dashed, $\alpha\beta\gamma$). Each of the 961 cells carries the verdict of the published formation and retention physics compiled in Section 4.3, with 557 buildable under pessimistic retention, 153 only at the optimistic edge, 140 not buildable and 111 unaddressed by the literature. The map enters no likelihood and no prior, by the gate's own hard rule, and the overlay is a visual consistency check rather than a combination. Data and script: `figs/fH_g3_overlay.json`, `figs/fH_expand_v1.py`.

profile-driven analysis favours sits in the band that only the optimistic edge of the retention literature reaches, adjacent to cells no published channel populates at all.

Formation physics and the profile-driven kinematics thus disagree, and the overlay is carried as exactly that: a live disagreement, reported alongside the posterior, tilting nothing (the map never enters the likelihood). A dynamical caution cuts the same way: an extended $3.5 \times 10^4\ M_\odot$ remnant configuration at 0.2 pc scale, the kind of solution the unrepaired profile leg preferred, has a subsystem relaxation time of order 10^5 yr and does not persist; where the misfitting fit parked its posterior was not a physical configuration at all.

8 What would decide it

In order of leverage per unit effort. **(i) Calibrate the fast-star machinery on the extended alternative: done.** The campaign ran and closed (Section 5.3): the failure traced to the retired profile leg, the verdict configuration is calibrated at the fiducial bracket, the fit of record's preference is quotable at the strength that survived, and the mass-to-light brackets are calibrated equivalently to each other under self-consistent generation (Amendment A5), with the joint cube still uncalibrated and embargoed at every bracket. **(ii) Resolve the Dai/TRAPUM timing discrepancy.** A timing-level reconciliation of the DM offsets and of pulsar B's derivative sign, by the groups that own the data; and, in the same pass, of the three Shklovskii terms and the one Galactic term that Section 6 could not reproduce from the published proper motions. This item is now the highest-leverage one on the list rather than a housekeeping note: the authorized rerun shows the reduction choice flips the verdict of the fit of record outright. Until it is settled every acceleration-based claim about this cluster inherits it, including ours. **(iii) Enlarge the fast-star census.** The robust five carry the result; the existing astrometric archive plus one more epoch of the same quality would either populate the tail or bound it. **(iv) An anisotropy-structured profile analysis — run, and closed without a mechanism.** The component-differential residual is a named, measured target (radial under-predicted, tangential over-predicted, at every radius); the separately pre-registered campaign that this item called for has since been run, and the paragraph below reports what it settled. It would improve the visible model and could open the marginal inner bins; on the terminality argument above, the verdict-dominant outer profile should not be expected to arbitrate under any version of it. **(v) More timed pulsars.** At a census of order 10^2, jerks activate as a genuinely independent channel (Abbate et al., 2019; Chen et al., 2025); the corrected floor makes the requirement concrete, with the finalized forecast flagged for the released record.

That campaign was pre-registered on its own plan, run to its stopping rule, and closed on its criterion 4, the outcome that no mechanism can be identified (Swanson, 2026). Its component-resolved discrepancy model clears the per-component sign test this paper's shared radial term failed, at all 36 readings: the differential axis was the missing degree of freedom. The differential is real and one-signed, with a median split of 0.29 to 0.33 $\mathrm{km\,s^{-1}}$ inside the knot span, and it concentrates beyond $100''$, where the terminality argument above already places the systematics floor an order of magnitude above the statistical errors; the width of its prior controls where in the outer profile it sits. The budget gate fires beyond about $130''$, forbidding escalation to an axisymmetric model under the gate's own clause, and pointwise amplitudes past that radius lie outside the calibrated region. Nothing from that campaign enters this paper's verdict configuration.

9 Conclusion

We pre-registered a joint analysis of the contested central dark mass of ω Cen, validated it adversarially, ran it, amended it five times under independent review, twice in its error model and twice in the scope on which its gates are scored, and report the pre-registered null: the data cannot yet decide between a compact and an extended central dark mass. The analysis localizes the indecision. The dispersion profile, the formally strongest dataset, cannot arbitrate the question at any error model, not because the dark component is absent from its bins (Table 1 finds it substantial, 0.008–0.69 at $100''$ across the prior grid) but because that signal is buried under a physics-derived systematic budget (0.5–1.6 $\mathrm{km\,s^{-1}}$) an order of magnitude above the statistical errors (0.03–0.05 $\mathrm{km\,s^{-1}}$), and its residual structure is not radial; the pulsar accelerations are censored into near-silence; the jerks sit under a 1-stable floor that no timing precision escapes, with a corrected constant that other groups can verify in one line of algebra. What remains live is small and specific: five robust fast stars whose stable compact preference now carries a demonstrated calibration at the fiducial bracket, quotable as a consistency statement conditional on a $\gamma \approx 1.75$ tracer cusp, on the fiducial bracket, on the TRAPUM pulsar reduction, and on the truncated-polytrope speed-tail family (a hard-truncated isothermal alternative is admissible and untested), each of which we price; a formation-physics overlay that builds the compact solution easily and the extended one only through self-excluded hybrids, and a short, ordered list of what would settle the rest. The mass tension of ω Cen is not a disagreement between datasets; it is a disagreement between assumptions that the datasets are not yet strong enough to overrule, and this analysis measures how strong they would need to be.

Data availability

The complete analysis chain is distributed with the paper source at omegacentauri.me: the pre-registration with its full git history and all five amendments; the provenance-annotated data extractions with content hashes; the validation reports for gates G1–G3 with fixed seeds; both fit reports with every configuration, gate verdict, and sensitivity row; the three independent methodological reviews behind Amendment A2; and a self-audit script that re-derives every number quoted in the fit reports from the released result files (140 checks); superseded result blocks are marked in situ in the released files. The provenance chain includes its own error record: a factor-of-ten transcription slip in an early digitized anisotropy anchor was caught by the subsequent checksum-verified file retrieval and is documented, with both values, in the released extraction, which is the discipline this paper's conclusions depend on working in practice.

Acknowledgements

This analysis rests on the public data products of the oMEGACat survey, the TRAPUM collaboration, and the SDSS-V Local Volume Mapper, and on the published formalisms of Prager et al. and Abbate et al., whose jerk framework we correct in one constant and otherwise reproduce with admiration.

References

Abbate, F., Spera, M., and Colpi, M. (2019). Intermediate mass black holes in globular clusters: effects on jerks and jounces of millisecond pulsars. *Monthly Notices of the Royal Astronomical Society*, 487:769–781.

Aros, F. I., Sippel, A. C., Mastrobuono-Battisti, A., et al. (2020). Dynamical modelling of globular clusters: challenges for the robust determination of IMBH candidates. *Monthly Notices of the Royal Astronomical Society*, 499:4646–4665.

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.

Bañares-Hernández, A., Calore, F., Martín Camalich, J., and Read, J. I. (2025). New constraints on the central mass contents of Omega Centauri from combined stellar kinematics and pulsar timing. *Astronomy & Astrophysics*, 693:A104.

Baumgardt, H., He, C., Sweet, S. M., Drinkwater, M., Sollima, A., Hurley, J., Usher, C., Kamann, S., Dalgleish, H., Dreizler, S., and Husser, T.-O. (2019). No evidence for intermediate-mass black holes in the globular clusters ω Cen and NGC 6624. *Monthly Notices of the Royal Astronomical Society*, 488:5340–5351.

Baumgardt, H. and Hilker, M. (2018). A catalogue of masses, structural parameters, and velocity dispersion profiles of 112 Milky Way globular clusters. *Monthly Notices of the Royal Astronomical Society*, 478:1520–1557.

Chen, X., Vázquez-Aceves, V., et al. (2025). Detecting intermediate-mass black holes in globular clusters with miniature pulsar timing arrays. arXiv:2507.08201; unpublished as of 2026-08-16.

Colom i Bernadich, M., Dai, S., Abbate, F., Kerr, M., Bachetti, M., Bhargava, Y., Buchner, S., Johnston, S., Burgay, M., Possenti, A., Nag, R., Ridolfi, A., Carleo, A., Corongiu, A., Freire, P. C. C., Camilo, F., Chen, W., Cadelano, M., Risbud, D., Padmanabh, P. V., Champion, D. J., Kramer, M., Stappers, B., Serylak, M., Balakrishnan, V., Bailes, M., Dutta, A., Vleeschower Calas, L., Venkatraman Krishnan, V., and Men, Y. (2026). A joint MeerKAT and Parkes view of Omega Centauri: New TRAPUM searches and pulsar timing. *arXiv e-prints*, page arXiv:2603.21845.

Dai, S., Johnston, S., Kerr, M., Berteaud, J., Bhattacharyya, B., Camilo, F., and Keane, E. (2023). Timing of pulsars in the globular cluster ω Centauri. *Monthly Notices of the Royal Astronomical Society*, 521:2616–2626.

Dickson, N., Hénault-Brunet, V., Baumgardt, H., Gieles, M., and Smith, P. J. (2023). Multimass modelling of Milky Way globular clusters – I. implications on their stellar initial mass function above 1 $M_{\odot}$. *Monthly Notices of the Royal Astronomical Society*, 522(4):5320–5339.

Dickson, N., Smith, P. J., Hénault-Brunet, V., Gieles, M., and Baumgardt, H. (2024). Multimass modelling of Milky Way globular clusters – II. present-day black hole populations. *Monthly Notices of the Royal Astronomical Society*, 529(1):331–347.

Geyer, E. H., Nelles, B., and Hopp, U. (1983). The shape of NGC 5139 (omega Centauri): the ellipticity as a function of radius. *Astronomy & Astrophysics*, 125:359–364.

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.

Häberle, M. et al. (2026). SDSS-V local volume mapper (LVM): The integrated light and internal rotation of omega Centauri. *Astronomy & Astrophysics*. Accepted for publication; arXiv:2607.28809; status checked 2026-08-26.

Häberle, M., Neumayer, N., Seth, A., Bellini, A., Libralato, M., Baumgardt, H., Whitaker, M., Dumont, A., Alfaro-Cuello, M., Anderson, J., Clontz, C., et al. (2024). Fast-moving stars around an intermediate-mass black hole in ω Centauri. *Nature*, 631(8020):285–288.

Häberle, M., Neumayer, N., Seth, A. C., et al. (2025). oMEGACat. VI. analysis of the overall kinematics of ω Centauri in 3D: velocity dispersion, kinematic distance, anisotropy, and energy equipartition. *The Astrophysical Journal*, 983:95.

Kennedy, M. C. and O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*, 63:425–464.

Kızıltan, B., Baumgardt, H., and Loeb, A. (2017). An intermediate-mass black hole in the centre of the globular cluster 47 Tucanae. *Nature*, 542:203–205.

Mai, A., Kremer, K., and Kıroğlu, F. (2026). Shadows of the colossus: Hierarchical black hole mergers in a 10-million-body globular cluster simulation. *The Astrophysical Journal*, 998:138.

Mann, C. R., Richer, H., Heyl, J., Anderson, J., Kalirai, J., Caiazzo, I., Möhle, S., Knee, A., and Baumgardt, H. (2019). A multi-mass velocity dispersion model of 47 Tucanae indicates no evidence for an intermediate mass black hole. *The Astrophysical Journal*, 875:1.

Martinez, M. A. S., González Prieto, E., and Rasio, F. A. (2026). Survival analysis of intermediate-mass black holes in dense star clusters. *arXiv e-prints*, page arXiv:2602.23431. Submitted to ApJ.

Nosek, B. A., Ebersole, C. R., DeHaven, A. C., and Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11):2600–2606.

Noyola, E., Gebhardt, K., and Bergmann, M. (2008). Gemini and Hubble Space Telescope evidence for an intermediate-mass black hole in ω Centauri. *The Astrophysical Journal*, 676:1008–1015.

Noyola, E., Gebhardt, K., Kissler-Patig, M., Lützgendorf, N., Jalali, B., de Zeeuw, P. T., and Baumgardt, H. (2010). Very large telescope kinematics for omega Centauri: Further support for a central black hole. *The Astrophysical Journal Letters*, 719:L60–L64.

Pancino, E., Seleznev, A., Ferraro, F. R., Bellazzini, M., and Piotto, G. (2003). The multiple stellar population in ω Centauri: spatial distribution and structural properties. *Monthly Notices of the Royal Astronomical Society*, 345:683–690.

Prager, B. J., Ransom, S. M., Freire, P. C. C., Hessels, J. W. T., Stairs, I. H., Arras, P., and Cadelano, M. (2017). Using long-term millisecond pulsar timing to obtain physical characteristics of the bulge globular cluster Terzan 5. *The Astrophysical Journal*, 845:148.

Smith, P. J., Hénault-Brunet, V., Dickson, N., Gieles, M., and Baumgardt, H. (2024). Probing populations of dark stellar remnants in the globular clusters 47 Tuc and Terzan 5 using pulsar timing. *The Astrophysical Journal*, 975(2):268.

Swanson, T. (2026). The component-differential residual in the omega centauri proper-motion dispersion profile: A pre-registered campaign that cannot identify its mechanism. https://omegacentauri.me/papers/omega-centauri-axi-note.pdf. Methods note, The Omega Centauri Society.

van de Ven, G., van den Bosch, R. C. E., Verolme, E. K., and de Zeeuw, P. T. (2006). The dynamical distance and intrinsic structure of the globular cluster ω Centauri. *Astronomy & Astrophysics*, 445:513–543.

van der Marel, R. P. and Anderson, J. (2010). New limits on an intermediate-mass black hole in Omega Centauri. II. Dynamical models. *The Astrophysical Journal*, 710:1063–1088.

Vitral, E. and Mamon, G. A. (2021). Does NGC 6397 contain an intermediate-mass black hole or a more diffuse inner subcluster? *Astronomy & Astrophysics*, 646:A63.

Whitaker, M., Kerr, E., Seth, A., Häberle, M., Strader, J., Anderson, J., Bellini, A., Clontz, C., Freeman, Z., Griggio, M., Kamann, S., Libralato, M., Neumayer, N., González Prieto, E., Rodriguez, C. L., Saracino, S., Smith, P., van de Ven, G., and Wang, Z. (2026). A long period stellar-mass black hole binary in ω Centauri. *The Astrophysical Journal Letters*.

White, R. E. and Shawl, S. J. (1987). Axial ratios and orientations of 100 galactic globular clusters. *The Astronomical Journal*, 94:30–34. compiled in the Harris Milky Way GC catalogue.

Zocchi, A., Gieles, M., and Hénault-Brunet, V. (2017). Radial anisotropy in ω Cen limiting the room for an intermediate-mass black hole. *Monthly Notices of the Royal Astronomical Society*, 468:4429–4440.

Zocchi, A., Gieles, M., and Hénault-Brunet, V. (2019). The effect of stellar-mass black holes on the central kinematics of ω Cen: a cautionary tale for IMBH interpretations. *Monthly Notices of the Royal Astronomical Society*, 482:4713–4725.

A Kill-rule firing status under alternative budget readings

Amendment A2 fixed a flat 1.0 km s^{-1} budget with a 2× kill threshold; the ratified Round-7 panel (P-12) kept that reading as the operative historical trigger and asked for the alternatives to be disclosed. The table reports, per configuration, the maximum fitted $|\delta(r)|$ inside the knot span and whether the kill rule fires under the flat reading ($2 \times 1.0 = 2.0$ km s^{-1}) and under the two ends of the physics-budget range ($2\times0.5 = 1.0$ km s^{-1}; $2\times1.6 = 3.2$ km s^{-1}), from the shipped record (`fH_calc7_records.json`). Under the low end every configuration fires; under the flat trigger four fire, which with the budget exceedance in all twelve was the recorded basis for retiring the leg; under the high end none fires, a reading the panel rejected because it would reopen a settled amendment chain.

Configuration	max $|\delta|$ (km s^{-1})	fires at 2.0	fires at 1.0	fires at 3.2
A2 $\alpha\beta\gamma$, 5.20 kpc	2.28	yes	yes	no
A2 $\alpha\beta\gamma$, 5.49 kpc	2.11	yes	yes	no
A2 Plummer, 5.20 kpc	1.67	no	yes	no
A2 Plummer, 5.49 kpc	1.46	no	yes	no
τ=0.5 $\alpha\beta\gamma$, 5.20 kpc	1.42	no	yes	no
τ=0.5 $\alpha\beta\gamma$, 5.49 kpc	1.32	no	yes	no
τ=0.5 Plummer, 5.20 kpc	1.36	no	yes	no
τ=0.5 Plummer, 5.49 kpc	1.20	no	yes	no
τ=2.0 $\alpha\beta\gamma$, 5.20 kpc	2.36	yes	yes	no
τ=2.0 $\alpha\beta\gamma$, 5.49 kpc	2.17	yes	yes	no
τ=2.0 Plummer, 5.20 kpc	1.78	no	yes	no
τ=2.0 Plummer, 5.49 kpc	1.56	no	yes	no

Table 5: Firing status of the A2-G-d kill rule under the three budget readings of the ratified P-12 disclosure. The recorded verdict used the flat 1.0 km s^{-1} trigger.

B Amendments A3 and A4, operative text

Section 5.3 scores the A2 gate battery on the verdict configuration under two dated gate-scope amendments. Their operative clauses are reproduced here from the pre-registration file's amendment log, verbatim except where a bracketed ellipsis marks a result-of-record line whose numbers are stated in Section 5.3 and Table 3. Both are ruled by the gate's author and ratified on the same date; neither changes the model, the likelihood, the priors, the grid, or the data.

> **A3 — 2026-08-14: A2-G-c scope under a fired A2-G-d (gate-scope clarification only; no change to model, likelihood, priors, grid, or data).** *Reason:* A2-G-c as written names the all-legs configuration. A2-G-d fired, and its consequence clause drops the profile leg from the verdict, so every quoted number comes from the leg-dropped configuration, which A2-G-c as written never scores. OCS-H-CAL-1 scored both (CAL_REPORT §6). *Change:* when A2-G-d fires, A2-G-c is evaluated on the verdict configuration (profile leg dropped), which must pass all G-c checks at $n \geq 200$. [Result of record: PASS.] *Mandatory disclosures traveling with any quoted region:* (1) the leg-dropped coverage check is exactly invariant to the profile-only mock misspecification, so along that axis it is satisfied trivially; the null and sign-recovery checks are non-trivial and passed on their own terms; (2) the all-legs configuration the gate names FAILs (extended 0.820, Wilson 0.779–0.855, $n = 400$), mechanism: profile-leg over-confidence (CAL_REPORT §2); (3) calibration holds at the fiducial M/L bracket only; the pessimistic bracket is uncalibrated in every configuration (0.680 extended, all legs). The profile-likelihood 90% region is reported as a disclosed secondary estimator alongside the HPD region

wherever a 90% region is quoted; no inner-quantile statement is made from it; decision criteria 1 and 3 remain in posterior terms.

A4 — 2026-08-14: A2-G-a and A2-G-b scope under a fired A2-G-d (gate-scope clarification only; no change to model, likelihood, priors, grid, or data). *Reason:* A2-G-a and A2-G-b, as written, name axes that exist only while the profile leg is in the verdict. A2-G-d fired; per A3's principle, gates are scored on the verdict configuration. *Change:* (a) When A2-G-d fires, A2-G-b is evaluated on the verdict configuration: $\ln K$ sign identical in every defined prior cell, both visible models (per A1), both distances. The halved/doubled knot-prior rows are vacated, not passed: τ enters the profile leg only, and invariance to it is exact by construction — the verdict-control channel those rows were built to expose is removed by the leg drop itself. [Result of record: PASS, all 96 defined cells positive.] The consequence clause reads over the verdict configuration's defined cells and does not fire. (b) *Mandatory disclosures traveling with any criterion-2 statement:* (1) on the all-configuration set including profile-leg-inclusive fits, A2-G-b FAILed (52 of 288 minority-sign cells, FLAGD §5); mechanism: discrepancy-prior control of profile-leg-inclusive verdicts, the finding that A2-G-d's leg drop remedies, not an anomaly the re-scoping hides; (2) the prior-scale rows are satisfied trivially on the verdict configuration, in the same sense as A3 disclosure (1); (3) the unanimity of record spans a strictly smaller axis set than the gate as originally written. (c) A2-G-a is vacated when A2-G-d fires (no profile residuals in the verdict likelihood). It failed 0 of 12 on the all-legs set. No consequence clause; disclosed in methods in one sentence.