

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

v1.0 (last revised 2026-08-16) — 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), and the joint accretion bound (Paper F);
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 $\sim 2\text{--}3 \times 10^5 M_\odot$ 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 dispersion profile, measured to $0.03\text{--}0.05 \text{ km s}^{-1}$, 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; its verdict-dominant bins lie at radii carrying under 10^{-3} of the dark component's enclosed-mass signal, so the leg's formal power is spurious precision, and we retire it from the verdict by a pre-set gate. 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 leg alone, under the assumed tracer cusp, returns a compact optimum ($M_{\text{dark}} \approx 2.0\text{--}2.5 \times 10^4 M_\odot$ 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 (90 per cent coverage 0.88 extended / 0.87 point, 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, under the assumed tracer cusp and at the fiducial bracket only; 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	4
2	Data	5
3	Pre-registered design	5
4	Machinery and validation	6
4.1	Joint likelihood and gate G1	6
4.2	The jerk module, gate G2, and a correction to a published constant	7
4.3	The formation-physics overlay, gate G3	7
5	The joint fit, twice	8
5.1	First contact: the error model is the verdict	8
5.2	Amendment A2 and the gate battery it failed	8
5.3	The fit of record, reported and not promoted	10
6	Diagnosis: where the discriminating power actually is	11
7	The formation-physics overlay	12
8	What would decide it	12
9	Conclusion	13

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 $\lesssim 1.2 \times 10^4 M_\odot$ 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 $2\text{--}3 \times 10^5 M_\odot$ 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 (Colomo 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 (Vital 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_{dark} and Plummer scale a , from the point-mass limit ($a \rightarrow 0$) 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 four 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 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 (A, C, D, E, F), 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 mas yr^{-1} threshold, and the published contamination density, which reproduces the paper’s own expected foreground count (0.0735 against their quoted 0.074). 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) 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 (B, D, K), and those carry what constraining power the leg has. We verified the published reduction chain by recomputing pulsar A’s bound from its P , $\dot{P}$, and the Shklovskii and Galactic terms (1.9796×10^{-9} against the printed $1.98 \times 10^{-9} \text{ m s}^{-2}$). 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, $r = 1.8''\text{--}311''$, statistical errors $0.023\text{--}0.05 \text{ km s}^{-1}$ 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 $8.4 \pm 0.8 \text{ km s}^{-1}$ at $4.7'$) frozen from Häberle et al. (2026); kinematic distance $5,494 \pm 61 \text{ pc}$ from the same survey, with the $5,200 \text{ pc}$ 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 $\omega \text{ Cen}$ 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 ($r > 100''$), where the dark component contributes below one part in 10^3 of the enclosed mass everywhere in the prior grid.

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. Its elements: (i) the model, one dark Plummer component with (M_{dark}, a) continuous over $[10^3, 10^6] M_{\odot} \times [10^{-4}, 3] \text{ pc}$, 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. **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).

The amendments changed the error model. They did not change the priors, the criteria, the data, or the quotability rules, and the second 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:

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

against the published $\xi \simeq 3.04$, a 13.8 per cent difference confirmed by six independent routes (closed form, direct quadrature, a Poisson-field Monte Carlo that never uses the formula, symbolic evaluation, antiderivative differentiation, and 50-digit arithmetic). 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 (18 of 20 checks at 10^{-10} or better; the two “failures” are the ξ comparison itself).

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 $10^{-19} \text{ m s}^{-3}$: the floor, not the precision, is the limit. Reaching a median $\ln K > 3$ requires a census of order 10^2 timed pulsars (indicatively ~ 100 at $2.5 \times 10^5 M_\odot$ and ~ 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 ~ 59 times too high for this cluster, corrected in the same pass as ξ . Consistently, the eight measured TRAPUM $\ddot{v}$ values are inconsistent with cluster jerks at all (Fisher $p = 2.4 \times 10^{-8}$, sign coherence $p = 0.008$), matching the source paper’s own caution that they may be spurious, and they enter 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 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: the error model is the verdict

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 km s^{-1} keeps the extended preference at $\ln K \approx -10$; a 2.0 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 1 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 1 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 ($\approx 1.0 \text{ km s}^{-1}$) is calibrated against the rejected model itself and is therefore circular; the value that flips the verdict (2.0 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 339,145 nats across the parameter plane, lie at radii where the dark component contributes under 10^{-3} of the enclosed mass. The leg’s dominance was spurious precision: its verdict-dominant bins carry no verdict signal.

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 $\text{Normal}(0, 1.0 \text{ km s}^{-1})$ 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 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- $\text{Normal}(0, 0.3 \text{ km s}^{-1})$ 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.

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




Figure 1: 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 does not measure the dark component’s geometry; it measures the analyst’s error model. The pre-registered gates (Section 5.2) formalize that reading, and the leg is retired from the verdict.

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.

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, and covers adequately on the point-mass injection (0.86 at the primary condition), but 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 fires. The fitted discrepancy reaches 2.1–2.4 km s^{−1} in the $\alpha\beta\gamma$ -model configurations, beyond twice its 1.0 km s^{−1} physics budget, and exceeds the budget itself in all twelve configurations. 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.

5.3 The fit of record, reported and not promoted

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 $M_{\text{dark}} = 2.0\text{--}2.5 \times 10^4 M_{\odot}$ at the point-mass limit of the grid, compact posterior mass 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 $2.2\text{--}2.7 \times 10^4 M_{\odot}$; 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.

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 a preference certified by an uncalibrated instrument is not a measurement. 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), 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 gate author’s statement of record:

“Under the fit of record — profile leg dropped per gate A2-G-d — the fast-star sample prefers the compact configuration, with the reporting machinery validated on that configuration by the A2 gate battery (90% coverage 0.88 extended / 0.87 point, null and sign-recovery passed; $n = 400$) at the fiducial M/L bracket. 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 (Amendment A3).”

That statement is dated, and the campaign continued past it. Amendment A5 records what it found: a mechanism rather than a repair. The pessimistic bracket’s under-coverage is mean displacement, not region width. 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.

Table 1 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

Table 1: The fit of record: calibrated 90 per cent regions in (M_{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. All regions about the $a = 10^{-4}$ pc grid edge and are one-sided in a .

Configuration	HPD90 M_{dark} ($10^4 M_{\odot}$)	PL90 M_{dark} ($10^4 M_{\odot}$)	a extent (pc)	M_{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

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. On the fit of record the pre-registered joint-consistency criterion (criterion 3) is met in 31–33 of 36 prior cells per configuration, with both posterior-predictive checks passing in all cells, and criterion 4 does not fire. 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; and the overlap component is weakly discriminating, the pulsar leg contributing approximately 3 nats. The fit of record is, in substance, five stars: five well-measured stars with a modeled selection function, a validated instrument, and an assumed tracer cusp.

On the full sixteen-configuration A2 set, criterion 4 fired in the majority of configurations and the gate battery as originally scoped failed; both facts stand in the record and in Section 5.2. 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.

6 Diagnosis: where the discriminating power actually is

The dispersion profile cannot arbitrate this question at the radii that carry its weight, and there the statement is permanent at this precision. Three independent facts compose the argument. The bins that dominate the leg’s formal weight lie at $50''$ – $250''$, where the dark component’s enclosed-mass contribution is under 10^{-3} : the discriminating signal is not there. 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 precision is not usable. And the residual against every tested model is component-differential: the systematic is not representable by any radial term. 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, and the analysis says so plainly. Read correctly as censored one-sided bounds, seven usable accelerations span 3.0 nats across the entire (M_{dark}, a) plane; alone, they bound the central mass only from above, at a prior-dominated $8.4 \times 10^5 M_{\odot}$. 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 $0.02\text{--}0.04\text{ pc cm}^{-3}$ across all five overlapping pulsars, far beyond either’s formal errors. 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.

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 $\ddot{v}$ 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 everything about their leg behaved: stable optimum across every bracket, selection function reproducing the source paper’s own contamination arithmetic, agreement with an independent analysis chain, and clean single-leg injection recovery. 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 (the all-seven variant fails the count posterior-predictive check in every cell, so the robust-five census is currently the informative one).

7 The formation-physics overlay

The G3 map asks, independently of all kinematics, which regions of the (M_{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 ($8 \times 10^3\text{--}5 \times 10^4 M_{\odot}$, $a < 0.01\text{ pc}$) is buildable under pessimistic retention throughout: it is what the merger-growth models of González Prieto et al. (2025) produce, with recoil from the dominant, capture-driven growth channel four orders of magnitude below the escape velocity. The extended region favored by Bañares-Hernández et al. (2025) is the hard one: as a pure remnant population, $2.5\text{--}3 \times 10^5 M_{\odot}$ 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 $1.5\text{--}6.6 \times 10^4 M_{\odot}$, 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.

Formation physics and the profile-driven kinematics thus point in opposite directions, 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; until then every acceleration-based claim about this cluster inherits it. **(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 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 its error model once under independent review, 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, because its precision lies at radii that carry no discriminating signal 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 (a consistency statement, under an assumed tracer cusp, quoted at the fiducial bracket), 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 both 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., Buchner, S., Johnston, S., Burgay, M., Possenti, A., Ridolfi, A., Freire, P. C. C., Kramer, M., Stappers, B., et al. (2026). A joint MeerKAT and Parkes view of Omega Centauri: New TRAPUM searches and pulsar timing. *arXiv e-prints*, page arXiv:2603.21845. Under review.
- 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. *Monthly Notices of the Royal Astronomical Society*.
- 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*.
- 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. Bibcode 1983A&A...125..359G.
- González Prieto, E., Rodríguez, 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. arXiv:2607.28809; unpublished as of 2026-08-16.
- Häberle, M., Neumayer, N., Seth, A., Bellini, A., Libralato, M., Anderson, J., Whitaker, M., van de Ven, G., Watkins, L. L., Clontz, C., et al. (2024). Fast-moving stars around an intermediate-mass black hole in ω Centauri. *Nature*, 631(8020):285–289.
- 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*.
- 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 Kiroğlu, 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.
- 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 the population of dark stellar remnants in the globular clusters 47 Tuc and Terzan 5 using pulsar timing. *The Astrophysical Journal*, 975.
- 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.
- Vital, 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, K. et al. (2026). A long period stellar-mass black hole binary in ω Centauri. *The Astrophysical Journal Letters*, 1006:L1.
- White, R. E. and Shawl, S. J. (1987). Axial ratios and orientations of 100 galactic globular clusters. *The Astronomical Journal*, 94:30–34. Bibcode 1987AJ.....94...30W; 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.