

The Component-Differential Residual in the Omega Centauri Proper-Motion Dispersion Profile: A Pre-Registered Campaign That Cannot Identify Its Mechanism

Tim Swanson

The Omega Centauri Society / Post Oak Labs

tim@postoaklabs.com

September 2026

Draft v0.7 (last revised 2026-09-03) — methods note for the AXI campaign;
companion to the mass-tension paper (Paper H); prepared for omegacentauri.me

Abstract

A fit of the oMEGACat proper-motion dispersion profile of ω Cen under a shared smooth discrepancy in the mean leaves a component-differential residual: the model under-predicts the sky-radial dispersion and over-predicts the sky-tangential one, coherently, at every radius and every reading of the grid (Swanson, 2026f). This note records a campaign pre-registered to identify the mechanism behind that residual (Swanson, 2026b), run to its stopping rule, and closed on its own pre-registered outcome that no mechanism can be named. Four hypotheses were specified in advance with separating discriminants: anisotropy structure beyond one Osipkov-Merritt scale (H-A), flattening with azimuthal averaging (H-B), unrelaxed accretion-origin substructure (H-C), and component-dependent measurement systematics (H-D). Stage 0 could run two of the four on the frozen radial table, and both discriminants failed at all twelve residual-evaluation points; H-B and H-C were not tested. Stage 1 fitted M1, a component-resolved discrepancy model that decomposes the shared spline into a shared term δ_S and a differential term δ_D on the same four knots. M1 clears the per-component adequacy test the shared model failed, at all 36 readings, with a minimum p of 0.0301 over 144 tests. Two other gates go against it. Cross-axis stability fails on one of five pre-committed statements, through a discrete switch between the differential and the white systematic term that moves with the δ_D prior width; and the budget check fires on both of its scoring conventions, first at the 201.15'' bin ($\approx 199''$ interpolated), by up to a factor of 3.8, on the as-coded $|\delta_D|$ -versus-band convention, or first at the 142.43'' bin ($\approx 137''$ interpolated), by up to a factor of 7.5 over all 36 readings, on the convention amendment AXI-A3 derives the budget from. The calibrated-region boundary itself, a separate threshold on δ_D alone, is crossed earlier, near 134'' interpolated. Under the pre-registered consequence clauses this makes decision criterion 4, “the campaign cannot identify the mechanism”, the only quotable outcome. The measured object that survives is the δ_D amplitude, 0.16 to 0.45 km s⁻¹ over 36 readings, positive everywhere and inside the calibrated region everywhere, with a median split of 0.29 to 0.33 km s⁻¹ inside the knot span against the 0.27 to 0.31 km s⁻¹ that stage 0 measured independently. Escalation to an axisymmetric model is forbidden by the budget gate’s own clause. Paper H is untouched throughout.

Keywords: globular clusters: individual: NGC 5139 (Omega Centauri); stars: kinematics and dynamics; proper motions; methods: statistical

Epistemic status: This is a methods record of a pre-registered campaign that reached a pre-registered negative outcome. Nothing here is a mechanism claim, and nothing here is a claim about the cluster’s dark component. The one measured quantity offered for quotation is the δ_D amplitude, which is quotable only inside the calibrated region of Section 3.2 and travels with the two gate failures of Section 4.

1 Question and plan

Paper H’s A2 amendment replaced every quadrature floor on the profile leg with a smooth discrepancy in the mean, a four-knot cubic spline $\delta(r)$ in $\log_{10} r$ with knots at $\{5, 20, 80, 300\}''$, shared by both proper-motion components, in the framework of [Kennedy and O’Hagan \(2001\)](#). The FIT-2 record measured what that model leaves behind ([Swanson, 2026f](#)). Runs tests on the standardised profile residuals pass everywhere, at $p = 0.24$ to 0.98 . The sign tests do not: 31 to 34 of 40 sky-radial residuals are positive, at $p = 4.2 \times 10^{-5}$ to 8.4×10^{-6} , while the sky-tangential component is over-predicted in 25 to 29 of 40 bins. The verdict is identical at all three cell readings and in all twelve profile-leg configurations.

A shared $\delta(r)$ cannot represent a residual of that shape, and the marginalisation over a single Osipkov-Merritt anisotropy scale, $r_a \in \{1.5, 3, 6, 15, \infty\}$ pc, does not absorb it either. The object is a coherent split between the two sky components at every radius. The campaign recorded here was pre-registered to ask which named mechanism produces it ([Swanson, 2026b](#)).

1.1 Scope, fixed before any fit

The pre-registration carries a mandatory disclosure into every report the campaign produces, and it is reproduced here once, in its own words:

This campaign improves the visible model and may open the inner bins. The outer profile does not arbitrate compact-versus-extended under any version of it.

The boundary comes from a radius-resolved terminality ruling on Paper H’s profile leg. Beyond about $50''$ the systematics budget runs 10 to 40 times the statistical errors, and those radii carry less than 10^{-3} of the dark component’s enclosed-mass signal, so the bins that dominate the leg carry no verdict signal. Per the terminality ruling, a better error model does not create signal the geometry does not put there.

The boundary permits, narrowly: the inner bins ($\lesssim 10''$), where signal-to-systematics is of order unity, are marginal rather than terminal, and a campaign that closed the component-differential residual there might make them usable. That route runs through decision criterion 5 of the plan and needs its own ratified amendment to Paper H’s pre-registration. It forbids, by pre-commitment: no re-scoring of Paper H’s criteria 2, 3 or 4 from this campaign’s output, no restoration of the profile leg to the configuration that Paper H’s own misspecification gate dropped, and no relaxation of Paper H’s standing embargo list. This note makes no statement of any kind about the cluster’s dark component, and it quotes no dark-component posterior.

1.2 Four hypotheses, with their separating discriminants

H-A, anisotropy structure beyond one Osipkov-Merritt scale. The true $\beta(r)$ is not in the one-parameter family the visible model carries, and the mismatch shows as a component-differential residual because β enters σ_R and σ_T with opposite sign. Signature: the split tracks the gradient of the measured anisotropy profile rather than the local dispersion amplitude ([Zocchi et al., 2017, 2019](#); [Aros et al., 2020](#)).

H-B, flattening with azimuthal averaging. A spherical model fitted to azimuthally averaged data of a flattened system ($\epsilon = 0.17$, [White and Shawl 1987](#), the largest published value; the intrinsic axial ratio of [van de Ven et al. 2006](#), 0.78 ± 0.03 at inclination 50° , is the sourced flattening cross-check per the plan’s Q2 closure; star-count isophotes give 0.11–0.12 with a rounder core, [Geyer et al. 1983](#); [Pancino et al. 2003](#)) produces a component-differential residual by construction. Signature: this campaign takes residual power at $\cos 2\phi$ in position angle as the diagnostic for this mechanism, with the split amplitude scaling as the flattening of the axisymmetric models in ([van der Marel and Anderson, 2010](#)).

H-C, unrelaxed substructure of accretion origin. ω Cen is the stripped nucleus of an accreted system, and the split is the imprint of a subpopulation that has not phase-mixed ([Ibata et al., 2019](#); [Bekki and Tsujimoto, 2019](#); [Clontz et al., 2024](#)). Signature: a patchy split differing between subsamples at the same radius, the one signature of the four that separates subsamples of the same catalogue.

H-D, component-dependent measurement systematics. The split is in the catalogue rather than in the cluster: residual distortion or a per-epoch orientation imbalance projects unequally onto the two sky directions once they are defined relative to the cluster centre ([Häberle et al., 2024, 2025](#)). Signature: organisation on detector or epoch geometry and on per-bin star counts rather than on cluster-centric radius. Confirming H-D would end the campaign rather than extend the model, so it was tested first.

Decision criterion 2 names a hypothesis as identified only when its own discriminant passes and the other three fail on theirs. Criterion 4 records in advance that “the campaign cannot identify the mechanism” is an acceptable and publishable outcome, to be stated as such if criterion 2 names no hypothesis in at least half the defined cells.

1.3 Data and gates

The data are frozen at listed versions: the 40 adaptive log bins of the oMEGACat VI dispersion and anisotropy table, $r = 1.8209$ to $311.1198''$ over 610,846 stars, undigitised, with its asymmetric sky-radial and sky-tangential errors ([Häberle et al., 2025](#)); the visible-model pair carried as a reported bracket with neither member promoted, a Plummer sphere at the [Baumgardt and Hilker \(2018\)](#) half-light scale and the $\alpha\beta\gamma$ profile of [Bañares-Hernández et al. \(2025\)](#); and both distances, 5494 and 5200 pc. The line-of-sight rotation table ([Häberle et al., 2026](#)) is a line-of-sight observable and enters no likelihood at any stage. No fast-star data and no pulsar data enter the campaign, so both legs of Paper H’s verdict configuration are untouched by construction.

Six gates were fixed before any fit, each with its consequence clause: X-G-a, injection and recovery at the real per-bin error level with component-differential misspecification injections including a sign-flipped member; X-G-b, per-component sign and runs tests at $p > 0.01$ in all four tests at all three readings; X-G-c, cross-axis stability of every reported statement across both visible models, both distances and the halved and doubled δ_D prior rows; X-G-d, a budget check that fires if the required differential exceeds twice its named budget anywhere inside the knot span; X-G-e, a standing scope gate with no statistic, checked by the referee panel; and X-G-f, a pre-real-data check that M1 at $\delta_D \equiv 0$ reproduces the shared-discrepancy fit to the recorded digit. Section 4.3 of the plan opens with the rule that all gates must pass before any statement from the campaign is quotable. Amendment AXI-A4 (Table 3) records the standing exceptions this note relies on: a member’s own consequence clause can make a specific statement quotable under a FAIL or FIRES verdict (X-G-a’s regional licence under AXI-A2; X-G-c’s and X-G-d’s own clauses naming criterion 4 as the quotable outcome), and none of the six gates changed the rule itself.

2 Stage 0: the two discriminants the frozen table can carry

Stage 0 was diagnostic only, with no new likelihood, run against the existing residuals at the three residual-evaluation points crossed with both visible models and both distances (Swanson, 2026a). Its thresholds were committed and dated in the deliverable before any discriminant statistic was computed. The primary statistic is the per-bin split $D_i = e_{R,i} - e_{T,i}$, from which the shared $\delta(r)$ cancels by construction, so it is independent of the shared-term posterior.

Two of the four hypotheses cannot be tested on the frozen radial table at all. H-B needs a position-angle-resolved dispersion field and H-C needs per-star population assignments, and neither is in the table. Stage 0 therefore recorded in advance that it cannot name any hypothesis as identified whatever its statistics return, because two of the three required failures are unavailable. H-B and H-C are not tested and are nowhere called disfavoured.

H-A: FAIL. The correlation of the split with the anisotropy gradient runs $|\rho| = 0.081$ to 0.185 at $p = 0.248$ to 0.614 , against a committed threshold of $|\rho| \geq 0.40$ at $p < 0.01$ at all twelve readings. The signature condition, that the split track the anisotropy gradient more strongly than the dispersion amplitude, holds at none of the twelve. The verdict is stable at all three smoothing bandwidths.

H-D: FAIL. The correlation with per-bin star counts runs $|\rho| = 0.144$ to 0.233 at $p = 0.147$ to 0.370 , and its radius-controlled partial runs 0.082 to 0.145 against a threshold of 0.30 . The surface-density form of the predictor behaves the same way: its raw correlation runs $|\rho| = 0.239$ to 0.323 and its radius-controlled partial runs 0.273 to 0.283 , both under threshold at every reading (Swanson, 2026a). The committed stop-or-continue rule stops the campaign only on a passing H-D discriminant, so it did not fire.

What the table cannot separate. The three predictors and radius are close to the same variable in rank space over the frozen 40 bins: the anisotropy gradient correlates with radius at $\rho = -0.986$ at the primary bandwidth and at -1.000 at the widest, star counts with radius at $+0.986$, and dispersion with radius at -0.970 . An uncontrolled rank correlation with any of them is close to a rank correlation with radius, which is why the partial-correlation controls were committed in advance and why both hypotheses fail those controls on the secondary standardised statistic while passing the raw correlation. This is a measurement-design fact about the azimuthally averaged table rather than a finding about the cluster.

The scale check. The one component-differential effect measured in the same catalogue is the energy-equipartition asymmetry, $\eta = 0.088 \pm 0.017$ at the centre falling to 0.049 ± 0.009 at the half-light radius over the mass range 0.288 to $0.690 M_\odot$, with the radial component of equipartition declining much faster with radius than the tangential (Häberle et al., 2025). Against the predicted differential scale $\Delta_{\text{pred}} = \sigma \eta \frac{1}{2} \ln(m_{\text{max}}/m_{\text{min}})$, the observed median $|D|$ of 0.52 to 0.53 km s^{-1} gives a ratio of 0.724 to 0.735 at every reading, inside the committed consistency band of $[0.5, 2.0]$. The median signed split is $+0.27$ to $+0.31 \text{ km s}^{-1}$, with the sky-radial component as the under-predicted one, reproducing the direction of the FIT-2 record. The comparison carries no verdict for any hypothesis and was committed as a check on the prior derivation of Section 3.

3 Stage 1: the M1 model and the gate battery

M1 is the minimal extension of the shared discrepancy model, in the same framework. The shared mean discrepancy is decomposed on the same four knots,

$$\sigma_{R,\text{model}}(r) \rightarrow \sigma_{R,\text{model}}(r) + \delta_S(r) + \delta_D(r), \quad (1)$$

$$\sigma_{T,\text{model}}(r) \rightarrow \sigma_{T,\text{model}}(r) + \delta_S(r) - \delta_D(r), \quad (2)$$

with δ_S keeping the existing $\text{Normal}(0, 1.0 \text{ km s}^{-1})$ independent knot priors, the half- $\text{Normal}(0, 0.3 \text{ km s}^{-1})$ white systematic s retained unchanged, and both terms marginalised in every reported quantity. M1 reduces to the shared model exactly at $\delta_D \equiv 0$. It adds four nuisance dimensions and no new physics machinery, and by construction it measures the differential rather than explaining it.

The δ_D knot prior is $\text{Normal}(0, 0.5 \text{ km s}^{-1})$, derived before any fit as the upper edge of the equipartition-differential term from the frozen η profile rather than as half of the shared budget. Flattening is excluded from that scale. The halved and doubled prior rows, $\tau_D \in \{0.25, 0.5, 1.0\} \text{ km s}^{-1}$, are pre-committed and mandatory and were noted non-trivial in advance, on the precedent that halving the shared model's knot prior had flipped evidence signs across whole configurations.

3.1 X-G-f, the pre-real-data reproduction check: PASS

Against the artifact the gate originally named, the check failed: log-evidence moved by 0.065 to 0.291 nats in all 16 configurations. The traced cause is not M1. The profile leg is the only object M1 changes, and the M1 profile-leg likelihood cube reproduces the reference exactly in all 16 configurations, with the whole discrepancy sitting in the acceleration leg, which the campaign does not touch, and accounted for by a reconciliation of the two pulsar spatial distributions committed one day after the reference was recorded (Swanson, 2026d). No code in the repository could reproduce the superseded kernel, so the gate as written tested a property of its reference. Amendment AXI-A1 repointed it to the same driver re-executed under the reconciled kernel, against which the same run reproduces every scalar, every posterior array, every log-likelihood cube, all 288 prior cells and the gate record. No code and no number changed between the failure and the pass.

3.2 X-G-a, injection and recovery: FAIL, with a calibrated region

Nine signed members were run at $n = 400$ mock realisations each: a point-mass injection, an extended injection, and a null, each carrying a coherent component-differential misspecification at $+0.5$, $+1.0$ and -0.5 km s^{-1} . Five of the nine fail a committed threshold, in two modes that separate cleanly (Swanson, 2026d).

Mode 1, the amplitude interval. Empirical two-sided 90 per cent interval coverage of the δ_D amplitude is 0.8500 to 0.8775 at an injected $\pm 0.5 \text{ km s}^{-1}$, inside the required $[0.85, 0.95]$, and 0.7125 to 0.7200 at $+1.0 \text{ km s}^{-1}$ in every injection, with Wilson intervals reaching only 0.755 to 0.762. The measured cause is a 7 to 8 per cent multiplicative shrinkage of the amplitude from the knot prior, against an interval width that does not grow to match it.

Mode 2, inherited coverage. Coverage of the injected dark component under the extended injection sits at 0.8425 to 0.8475 at $n = 400$, below the band with Wilson intervals containing it. Amendment AXI-A2 authorised one pre-committed extension of the three extended members to $n = 1600$ on a continued seed

stream, committed and dated before it was run, with the verdict scored on the pooled sample whatever it returned and no second extension permitted. It resolves below the band: 0.8237 to 0.8325 pooled, with Wilson upper limits of 0.8416, 0.8452 and 0.849997 for the -0.5 , $+0.5$ and $+1.0$ members respectively — the last excludes 0.85 by three parts in a million and is read as touching the threshold rather than clearing it (Swanson, 2026d). The coverage curve reproduces to about 0.01 at every nominal level the compression measured for the shared model before any differential block existed (Swanson, 2026e), and it is constant along the differential axis, so this is an inherited property of the profile leg rather than an effect of δ_D . The campaign quotes no dark-component posterior, so the consequence attaches to criterion 5, which stays closed.

What passes. The null member manufactures no evidence at any differential, with median $|\ln K| = 0.16$ to 0.18 against a threshold of 1 and a breach fraction of at most 0.0025. Correct-sign recovery is 400 of 400 in both signal injections, and the sign-flipped member added at referee request recovers at least as well as its mirror, symmetric to 0.025.

The regional ruling. Under AXI-A2 the amplitude axis is scored regionally. δ_D amplitudes are quotable only inside the calibrated region $|\delta_D| \leq 0.5 \text{ km s}^{-1}$, where all six ± 0.5 members pass; the $+1.0$ members remain failed and mark the boundary. Any amplitude outside the region is reported with the measured calibration attached, a bias of about -8 per cent of amplitude and an interval about 15 per cent narrow at 1.0 km s^{-1} , and is not quotable as a measurement. Every δ_D figure carries the region boundary, and the doubled-prior row is reported alongside the primary as the shrinkage sensitivity check.

The A3 non-triviality disclosure, measured rather than argued. X-G-a's binding condition, that the differential-injection axis be non-trivial for the configuration being scored or the gate reports as vacated, is checked directly. Holding the mock survey identical by the seed pairing of note X-G-a-N5 and varying only the injected differential amplitude moves the δ_D amplitude coverage from 0.850 to 0.713 in the null member and moves the recovered amplitude from 0.465 to 0.920 (Swanson, 2026d). The axis is therefore non-trivial, a coverage number along it is a measurement, and X-G-a reports as failed rather than as vacated. This disclosure travels with the gate whether or not it is quoted elsewhere.

3.3 X-G-b, the adequacy gate the shared model failed: PASS

Twelve fits were run, the four primary configurations crossed with the three δ_D prior rows, each read at the three residual-evaluation points, for 36 readings and 144 tests. The fit driver and the gate scorer were committed unexecuted before the first fit ran, so neither gate is scored on a statement list chosen after the numbers were seen (Swanson, 2026c).

X-G-b: PASS, at all 36 readings, with a minimum p of 0.0301 over the 144 tests (Table 1). The tightest value anywhere in the battery is now a runs test rather than a sign test. The profile-leg best cell is where a sceptic should look, because the joint MAP is chosen by a different leg, and M1 passes there as well.

One qualification travels with the pass. Clearing the adequacy gate says the residual is absorbed rather than explained, because a discrepancy term absorbs component-differential signal by construction, as the plan's own description of M1 states in advance. The pass is evidence that the differential axis was the missing degree of freedom, and not evidence about which mechanism supplies it.

Table 1: The per-component adequacy tests, M1 against the shared discrepancy model on the same 40 bins. Both models pass every runs test; the sign test is where they differ.

statistic	shared model (A2)	M1
sign test, sky-radial	31 to 34 of 40 positive	20 to 25 of 40 positive
sign test, sky-radial, p	4.2×10^{-5} to 8.4×10^{-6}	0.154 to 1.000
sign test, sky-tangential	25 to 29 of 40 over-predicted	19 to 26 of 40 positive
sign test, sky-tangential, p	failing at every reading	0.0807 to 1.000
runs tests, both components	0.24 to 0.98	0.0301 to 0.932
minimum p at any reading	sign test, 8.4×10^{-6}	runs test, 0.0301
readings passing	0 of 12 configurations	36 of 36 readings

3.4 X-G-c, cross-axis stability: FAIL, with the mechanism measured

Five statements were fixed in the committed scorer: the X-G-b verdict; whether any bin inside the knot span has a δ_D interval excluding zero, and with which sign; the X-G-d verdict; whether the maximum $|\delta_D|$ inside the span lies inside the calibrated region; and the sign of the amplitude with whether its interval excludes zero. Four hold across all twelve rows. The fourth does not, so **X-G-c: FAIL**.

In the four rows at $\tau_D = 0.5$ and the four at $\tau_D = 1.0 \text{ km s}^{-1}$ (eight rows, all three readings each) the maximum $|\delta_D|$ inside the span is 1.140 to 1.151 km s^{-1} , outside the region. In the four rows at $\tau_D = 0.25 \text{ km s}^{-1}$ (twelve readings) seven are low, inside the region: the posterior-mean and joint-MAP cells of three configurations (Plummer at 5494 and 5200 pc, $\alpha\beta\gamma$ at 5494 pc), plus the posterior-mean cell alone of the fourth ($\alpha\beta\gamma$ at 5200 pc) — bin-grid maximum 0.335 to 0.412 km s^{-1} , outer-knot posterior mean 0.345 to 0.427 km s^{-1} . The remaining five readings at $\tau_D = 0.25$ are high, outside the region: the profile-leg-best cell in all four configurations (1.140 to 1.144 km s^{-1}), plus $\alpha\beta\gamma$ at 5200 pc’s joint-MAP cell, which splits from its own posterior-mean cell (0.403 low against 1.140 high) rather than following it.

Read per cell, this is the two-mode signature directly: at the halved prior the posterior-mean cell samples the low mode (inside the calibrated region) in all four configurations, the joint-MAP cell joins it in three of the four ($\alpha\beta\gamma$ at 5200 pc instead samples the high mode at its joint MAP, splitting from its own posterior-mean cell), and the profile-leg best cell samples the high mode (outside the region) in every configuration; at the primary and doubled priors every reading samples the high mode. The per-row, per-reading classification is in this note’s released record companion (`calc7/axi_calc7_switch.json`); no mechanism statement follows from it.

The mechanism is a degeneracy between δ_D and the white systematic s , and it is measured rather than inferred (Table 2). The outer bins carry a component split that the fit must assign to one term or the other. At $\tau_D = 0.5$ or 1.0 the fit assigns the split to a large differential and a small white systematic; at $\tau_D = 0.25$ it assigns the split instead to a white systematic three times larger, with inflated errors. The switch is bracketed by the two sampled prior widths 0.25 and 0.5 km s^{-1} ; no intermediate τ_D was tested, so the transition is bracketed rather than located.

The pre-registration recorded before any fit that the shared model had replaced one assumption controlling the verdict with another, that a component-resolved discrepancy adds a third, and that X-G-c exists to make prior-width control visible if it occurs. It occurred, X-G-c detected it, and it is reported as the anticipated mechanism rather than as an anomaly. The failure is confined to one statement, so what it removes is the ability to say where the differential sits at the outer radii, which is the region the scope boundary already places outside arbitration.

Table 2: The switch behind the X-G-c failure, at the fiducial-bracket MAP cell, Plummer at 5494 pc. Each row is unimodal in the white systematic s ; the preferred solution switches between rows. Prior shrinkage cannot produce the move, because the outermost knot carries a posterior standard deviation of 0.0217 to 0.0509 km s^{-1} across the full 36-reading battery (0.0217 to 0.0424 for the nine Plummer-5494 readings; 0.0424, 0.0261 and 0.0260 at the three τ_D rows in this table’s own fiducial-bracket joint-MAP cell). The two right-hand columns are different objects: the bin-grid maximum is evaluated over the forty catalogue bins inside $[5, 300]''$ (the gate convention; the 300'' knot is not a bin), while the outer-knot column is the spline’s posterior mean at the 300'' knot itself. Both are correctly computed; an earlier version of this table carried only the knot value under a single “outer knot δ_D ” heading, where §3.4’s text separately quotes the bin-grid maximum. The “mass at $s < 0.06$ ” column is computed directly from the posterior samples; the `s_edge_mass` field released alongside this note is a separate grid-edge tail-mass check (10^{-52} to 10^{-57} in every reading) and is not this divider.

τ_D [km s^{-1}]	s mode	s median	mass at $s < 0.06$	bin-grid max, in span [km s^{-1}]	outer knot, posterior mean [km s^{-1}]
0.25	0.09	0.090	0.034	0.344	0.354 ± 0.042
0.50	0.03	0.030	0.968	1.145	1.207 ± 0.026
1.00	0.03	0.030	0.969	1.147	1.209 ± 0.026

3.5 X-G-d, the budget check: FIRES

The gate was scored two ways, both fixed before the run: against the constant 0.5 km s^{-1} scale named for the δ_D prior, and against the radius-dependent equipartition band it was derived from. It fires if either is exceeded by more than a factor of two anywhere inside the knot span $[5, 300]''$; the two bins outside the span, where the spline is a linear extrapolation, are excluded from the verdict. Amendment AXI-A3 (Table 3) records that the equipartition derivation is a statement about the split $\sigma_R - \sigma_T = 2\delta_D$, not about δ_D alone; the scorer as coded compares $|\delta_D|$ to the band directly, a factor of two tighter than the quantity the band was derived to bound. Both conventions are reported.

On the as-coded $|\delta_D|$ -versus-band convention, the maximum $|\delta_D|$ inside the span runs 0.335 to 1.151 km s^{-1} over the 36 readings, at $280.9''$ in every one of them. Twenty-nine of the 36 readings exceed twice the constant scale, and the same 29 exceed twice the equipartition band, with a maximum ratio to the band of 1.10 to 3.77. On the AXI-A3-consistent $|2\delta_D|$ -versus-band convention, all 36 of the 36 readings exceed twice the band, including the seven low-mode readings of Section 3.4 (ratio 2.20 to 2.59), with a maximum ratio of 7.54. **X-G-d: the clause FIRES on both conventions.** The seven readings that stay inside the band on the as-coded convention are the large- s solutions of Section 3.4, and they exceed the band on the AXI-A3-consistent convention as well, so the two gate outcomes have one mechanism between them under either scoring.

The consequence is applied as written: the component-differential model is declared misspecified beyond this campaign’s scope, and the campaign does not escalate to an axisymmetric model on the strength of it. The excess is confined to large radius, at a boundary later than the calibrated-region crossing of Section 3.2: the region boundary itself, $|\delta_D| = 0.5 \text{ km s}^{-1}$, is crossed near $134''$ interpolated, while the budget-overrun crossing follows at the $201.15''$ bin ($\approx 199''$ interpolated; 5.36 pc at 5494 pc, 5.07 pc at 5200 pc) on the as-coded convention, or at the $142.43''$ bin ($\approx 137''$ interpolated) on the AXI-A3-consistent convention. Inside $50''$, in binned medians over the groups $[5, 10)$, $[10, 25)$ and $[25, 50)''$, the differential sits at 0.17, 0.09 and 0.04 km s^{-1} against a budget of 0.80, 0.76 and 0.72 km s^{-1} . A four-knot spline that has to reach 1.2 km s^{-1} at $300''$ is absorbing something other than the equipartition asymmetry it was scaled against, at radii where the systematics budget is an order of magnitude above the statistical errors. That is a misspecification statement about M1 at large radius, and under Section 1.1 it is not a statement about the dark component.

4 Outcome: criterion 4

Three of the six gates stand against escalation and none stands for it. X-G-c’s consequence clause makes criterion 4 the only quotable outcome, X-G-d’s clause forbids escalation to an axisymmetric model, and X-G-a stood failed with its regional licence before the real-data fit began. The campaign’s outcome is therefore **criterion 4, the pre-registered “the campaign cannot identify the mechanism”**, reported with the mechanism of the X-G-c failure named, as the clause requires. Criterion 2 is not scored: two of its four discriminants cannot be run on the frozen radial table, and nothing in stage 1 adds either. H-B and H-C remain not tested.

The outcome carries two parts, both already recorded earlier in this note. The first is the plain failure-to-identify: on the frozen radial table, H-A and H-D fail their committed discriminants at all twelve readings (§2), and H-B and H-C cannot be run on that table at all. The second is the structural reason a positive result was never available from this table alone: the anisotropy gradient and the star-count predictor are rank-collinear with cluster-centric radius, $|\rho| = 0.86$ to 1.00 across the bandwidths tested (§2, “What the table cannot separate”; (Swanson, 2026a)), so an azimuthally averaged radial table cannot separate either testable discriminant from a radius trend by construction. Neither statement adds a mechanism claim or a compact-versus-extended statement; both restate §2.

4.1 The measured object




Figure 1: The primary output, criterion 1: the δ_D posterior against radius at the primary prior row, Plummer at 5494 pc, fiducial bracket, read at the joint MAP as a residual-evaluation point. The other eleven rows and the other two brackets differ in the third decimal except where Section 3.4 says otherwise. The curve crosses the calibrated-region boundary at about 130'', and **the pointwise curve beyond that crossing is outside the calibrated region and is not quotable as a measurement** (Section 3.2). Values from `fit_results/m1_fit_results.json`.

The δ_D amplitude, defined as the injection gate defined it, runs 0.16 to 0.45 km s^{-1} over the 36 readings, positive everywhere, with the 90 per cent interval excluding zero everywhere. Every value sits inside the calibrated region, so no amplitude reported here needs the out-of-region calibration attached. The pointwise curve is a different matter (Figure 1): $\delta_D(r)$ crosses the region boundary at about 130'' and reaches 1.15 km s^{-1} at 280.9'', so the outer half of the curve is outside the calibrated region, is drawn with the boundary marked, and is not quotable as a measurement. The regional calibration was measured on the amplitude rather than on the pointwise curve, which is a limitation of the calibration and is stated as one.

Two regimes separate in the posterior width. Inside about 50'' the binned-median posterior standard

deviation, over the groups $[5, 10)$, $[10, 25)$ and $[25, 50)''$, is 0.187 , 0.085 and 0.054 km s^{-1} (per-bin range 0.039 to 0.460), of order the prior scale, so the inner differential is prior-dominated and the data barely constrain it. Beyond about $100''$ the binned-median falls to 0.019 and 0.018 km s^{-1} over $[100, 200)$ and $[200, 300)''$ (per-bin range 0.013 to 0.028), with the posterior at 0.5 to 1.0 km s^{-1} . The inner-to-outer ratio of the binned medians is about 3 to 10, not the 20 to 50 an unweighted read of the two headline numbers would suggest; against the 0.5 km s^{-1} δ_D prior scale the outer regime is about 26 to 28 times smaller. The differential the profile leg requires is concentrated at the radii the terminality ruling of Section 1.1 already places behind a systematics floor 10 to 40 times the statistical errors. The inner-region statement here is about the width of a prior-dominated posterior and is not an inner-bin usability claim, which routes through criterion 5 alone.

The untuned external check. Stage 0 measured the split $D = e_R - e_T$ at a median of $+0.27$ to $+0.31 \text{ km s}^{-1}$ before M1 existed, and M1 parameterises that split as $2\delta_D$. Over the twelve primary-row readings the fitted median $2\delta_D$ inside the knot span is $+0.291$ to $+0.334 \text{ km s}^{-1}$, recovering the stage-0 split to about 8 per cent with no part of the pipeline tuned to make it do so. The ratio to the budget does not carry across as cleanly: on the absolute statistic stage 0 used, the median $|2\delta_D|$ over the median budget is 0.41 to 0.47 here against 0.72 to 0.74 there. Stage 0 took the absolute value of a noisy per-bin residual, which inflates its median, while M1's δ_D is a spline that has already smoothed that noise away. The signed medians, which are not subject to that inflation, agree.




Figure 2: The mandatory shrinkage sensitivity check, all three δ_D prior rows at the same configuration and reading as Figure 1. The doubled row lies within the line width of the primary. The halved row departs from both beyond about $60''$ and stays inside the calibrated region at the outer bins, which is the disagreement gate X-G-c failed on and the switch measured in Table 2. Beyond the boundary crossing the pointwise curves are not quotable as measurements.

4.2 Embargoed statements

Stated explicitly because negative outcomes invite overreading. No mechanism is identified, and “cannot identify the mechanism” is the finding rather than a stand-in for one. No pointwise δ_D value beyond about $130''$ is a measurement, because it sits outside the calibrated region. No dark-component number appears anywhere in this note, and criterion 5 stays closed. No compact-versus-extended statement is made in any register, sourced to the profile leg or otherwise. H-B and H-C are not tested and are not disfavoured; the frozen radial table carries neither their data nor their discriminants.

5 Disclosed deviations

Four amendments were ruled and ratified during the campaign, all after adoption and all disclosed here as deviations rather than folded in silently (Table 3). None changes a threshold, a prior or a model. AXI-A1 did change one member’s verdict, X-G-f from FAIL to PASS, by a reference repointing rather than by any change to M1 or the data; the other three amendments license what is quotable under a standing FAIL or FIRES verdict (AXI-A2, AXI-A4) or document a scoring convention without rescoreing anything (AXI-A3), and change no verdict.

A fifth process note belongs with them. The standing note added with AXI-A1 requires any future gate or amendment in this campaign that names a numerical reference to name a post-reconciliation artifact, or to state explicitly which kernel that artifact was produced under.

6 Downstream use

Feeds. The characterisation of the profile leg’s misfit that a shared discrepancy could not provide, for Paper H’s successor analysis. The visible-model bracket for any future fit, since a component-resolved discrepancy with calibrated uncertainty is a better statement of what the visible model does not know than a shared radial term. The methods record itself: a measured, gated answer to what the component-differential residual is, publishable whether or not a mechanism is named.

A successor instruction. Any future fit that draws on this campaign’s δ_D values must respect the calibrated-region rule of Section 3.2: only $|\delta_D| \leq 0.5 \text{ km s}^{-1}$ is a measurement, and any pointwise value beyond the boundary of Figure 1 (about $130''$) is not, whatever the fit that produced it. No downstream consumer should take an uncalibrated value from this record as an input without applying the measured calibration of Section 3.2 first.

Residual systematics. Stage 0 could test H-D only against per-bin star counts and their surface-density form, the two proxies the frozen radial table carries, so the orientation-dependent reading of the hypothesis is untested here rather than excluded. A consumer wanting to bound it should run the split against per-epoch and detector-geometry predictors in the parent catalogue (Häberle et al., 2024, 2025); nothing in this record licenses treating H-D as closed.

Cross-calibration with Paper H. This campaign’s mocks carry the profile leg and their coverage of the injected dark component sits below the required band (Section 3.2), while Paper H’s calibration passes only after the profile leg is dropped (Swanson, 2026e). The two results score different configurations, and each holds for the configuration it was run on. The compression measured here reproduces the shared-model compression Paper H measured before any differential block existed, which is the property Paper H removed by dropping the leg, so the two records agree once the configuration is named.

Does not feed. Paper H’s verdict machinery. Not criterion 2, not criterion 3, not criterion 4, not the fast-star leg’s quotable set, not the standing embargo list, and not the terminality ruling, which this campaign is scoped inside rather than against. The fit of record stands where its own amendments left it.

Disposition. Stage 1 is closed. Escalation to the axisymmetric model M2 is forbidden by X-G-d’s clause. The position-angle data work unit is neither triggered nor cleared by this fit, and if the question is ever

Table 3: Amendments to the plan of record, all ratified by the project owner. AXI-A1 through AXI-A3 are dated 2026-08-15; AXI-A4 is dated 2026-09-03, ruled during this note’s referee-round repair rather than during the campaign itself. Each was ruled by the author of the gate whose scope was in question, on the precedent the plan’s amendment protocol names.

	kind	what changed	why
AXI-A1	gate-scope clarification	X-G-f’s digit-level reference re-pointed from the original fit record to the same driver re-executed under the reconciled pulsar kernel	the original reference was produced under a kernel superseded one day later, so no code could reproduce it and the gate tested a property of its reference rather than of M1
AXI-A2	gate-scope clarification	X-G-a’s amplitude axis scored regionally, with the calibrated region $ \delta_D \leq 0.5 \text{ km s}^{-1}$; one pre-committed extension of the three extended members to $n = 1600$, with no second extension under any outcome	the failing members were the doubled-scale probes and mapped where the estimator breaks, while the campaign’s own target lies inside the region. The extension resolves whether the coverage edge was a sampling excursion, under an optional-stopping condition set in advance
AXI-A3	documentation clarification	records that the equipartition derivation is a statement about the split $\sigma_R - \sigma_T$, which under M1’s convention is $2\delta_D$, so the δ_D knot prior is a factor of two looser than the derivation’s literal edge	the campaign’s own fit surfaced the ambiguity. The looser reading is the conservative direction, with less shrinkage and wider intervals, and every gate was scored on the prior as run, so no rerun follows
AXI-A4	quotability override, dated 2026-09-03	states explicitly, against Section 4.3’s “all gates must pass” rule, what this note may quote while X-G-a stands FAIL and X-G-c and X-G-d stand FAIL/FIRES: the δ_D amplitude inside the calibrated region (X-G-a’s own regional consequence clause, AXI-A2), and the criterion-4 outcome with the X-G-c mechanism named (X-G-c’s and X-G-d’s own consequence clauses). No dark-component number, no mechanism identification and no compact-versus-extended statement is licensed by this override	the note as drafted offered these two statements as quotable without a rule change from Section 4.3 authorizing the offer explicitly; each was already licensed by its own gate’s consequence clause, so this amendment states that licence rather than creating one

taken up again that extraction goes forward as an independent work unit judged on its own merits. The orbit-superposition model M3 was rejected on the record at pre-registration and stays rejected. What the campaign leaves behind for a successor is the degeneracy of Section 2: on an azimuthally averaged radial table the candidate mechanisms are close to the same variable in rank space, and the axes that break the degeneracy are position angle and subpopulation, neither of which is a function of radius by construction.

Compute. The twelve fits took 1.2 minutes of wall clock on CPU, dominated by four model-table builds at 5 to 9 seconds each, because the base tables are shared across the three prior rows and M1’s marginalisation is closed-form; the figure pass added a further 1.5 minutes. The gate battery is where the cost sits: the main $n = 400$ battery (nine members, 3600 rows) ran 51.2 minutes, and the pre-committed extension ran a further 47.7 minutes at eight workers, 1200 new (injection, k) groups producing 3600 new rows (three members per group), at 1.38 to 2.47 s per row.

References

- 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.
- Bañares-Hernández, A., Calore, F., Martín Camalich, J., and Read, J. I. (2025). New constraints on the central mass contents of Omega Centauri from combined stellar kinematics and pulsar timing. *Astronomy & Astrophysics*, 693:A104.
- Baumgardt, H. and 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.
- Bekki, K. and Tsujimoto, T. (2019). A new formation model for ω Centauri: A complex interplay of astrophysical processes. *The Astrophysical Journal*, 886:121.
- Clontz, C., Seth, A. C., Dotter, A., Häberle, M., Nitschai, M. S., Neumayer, N., Feldmeier-Krause, A., Latour, M., Wang, Z., Souza, S. O., Kacharov, N., Bellini, A., Libralato, M., Pechetti, R., van de Ven, G., and Alfaro-Cuello, M. (2024). oMEGACat. IV. Constraining the ages of Omega Centauri subgiant branch stars with HST and MUSE. *The Astrophysical Journal*, 977:14.
- 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.
- 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., Bellini, A., Libralato, M., Clontz, C., Seth, A. C., et al. (2024). oMEGACat. II. Photometry and proper motions for 1.4 million stars in Omega Centauri and its rotation in the plane of the sky. *The Astrophysical Journal*, 970:192.
- Häberle, M., Neumayer, N., Clontz, C., Seth, A. C., Bellini, A., et al. (2025). oMEGACat. VI. Analysis of the overall kinematics of Omega Centauri in 3D: Velocity dispersion, kinematic distance, anisotropy, and energy equipartition. *The Astrophysical Journal*, 983:95.

- Ibata, R. A., Bellazzini, M., Malhan, K., Martin, N., and Bianchini, P. (2019). Identification of the long stellar stream of the prototypical massive globular cluster ω Centauri. *Nature Astronomy*, 3:667–672.
- Kennedy, M. C. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*, 63:425–464.
- 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.
- Swanson, T. (2026a). AXI campaign, stage 0: H-A and H-D diagnostics on the FIT-2 residuals. The Omega Centauri Society, internal record. Work unit OCS-AXI-STAGE0-1, 2026-08-14. Internal record, `paper/axi/stage0/STAGE0_REPORT.md`.
- Swanson, T. (2026b). Axisymmetric / anisotropy profile campaign: pre-registration. The Omega Centauri Society, internal record. Adopted 2026-08-14; amendments AXI-A1, AXI-A2 and AXI-A3 dated 2026-08-15. Internal record, `paper/AXI-PREREGISTRATION.md`.
- Swanson, T. (2026c). OCS-AXI-M1-FIT-1: the M1 fit on the frozen real table, and gates X-G-b, X-G-c, X-G-d. The Omega Centauri Society, internal record. Work unit OCS-AXI-M1-FIT-1, 2026-08-16. Internal record, `paper/axi/stage1/M1_FIT_REPORT.md`.
- Swanson, T. (2026d). OCS-AXI-M1-GATES-1: M1 build and the pre-real-data gate battery. The Omega Centauri Society, internal record. Work units OCS-AXI-M1-GATES-1 to GATES-3, 2026-08-15. Internal record, `paper/axi/stage1/GATES_REPORT.md`.
- Swanson, T. (2026e). Paper H CAL-1: calibration of the A2 fit at real precision. The Omega Centauri Society, internal record. 2026-08-14. Internal record, `paper/h/cal/CAL_REPORT.md`.
- Swanson, T. (2026f). Paper H FIT-2: the A2 discrepancy model on the frozen profile leg. The Omega Centauri Society, internal record. 2026-08-13. Internal record, `paper/h/fit2/FIT2_REPORT.md`.
- 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.
- 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.