

The viscous response of a certified axisymmetric Euler blowup profile is linear at leading order: a measurement by flux observables with theorem-backed nulls

Justin Hill

EPAGOGE LLC, independent researcher

hill.jt@icloud.com

August 13, 2026

Abstract

I measure the viscous response of a certified axisymmetric Euler corner-collapse profile, collapse exponent $\alpha = -0.34240 \pm 3 \times 10^{-5}$ (`boussinesq/ALPHA_RESULT.md`), on pre-registered viscous ladders spanning $\nu = 10^{-9}$ to 10^{-3} across three grid families, and I report three results. First, the physics: at the level of integrated quadratic functionals the viscous response is quasi-static linear response at leading order. In the swirl channel the explicit inviscid quadrature reproduces the measured $\nu \rightarrow 0$ leading coefficient to between 4.1×10^{-5} and 1.9×10^{-4} relative on two grids over $\nu \in [10^{-5}, 10^{-3}]$ at fixed time (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, control C7). In the enstrophy channel the production-subtracted budget ratio gives $\rho_{\text{bulk}} = 1.243/1.095/1.039/1.029$ (Simpson) at $\nu = 10^{-5}$ to 10^{-8} , with $|\rho_{\text{bulk}} - 1|$ falling monotonically to 0.017 on the trapezoid ladder and $\nu \rightarrow 0$ intercepts 1.0035 ± 0.0099 (trapezoid, four rungs) and 0.9999 (Simpson, three rungs) (`ext_inertia_t2/out/flux3/F3_RESULT.md`, section 3); primary reach $[10^{-7}, 10^{-5}]$, the 10^{-8} rung decided by the trapezoid control in class agreement. The joint license is exactly this and no more: linear response at the level of integrated quadratic functionals (C_2 in the r -weighted L^2 pairing, the ω_1 enstrophy flux in the r -weighted H^1 -seminorm pairing), leading coefficient only, sub-leading corrections vanishing no slower than $\nu^{0.4}$; the pointwise mechanism question remains open. Second, there is no finite viscous anomaly in this collapse: the registered deformation $d_\infty \approx 0.19$ was approximately 85% free-slip boundary pedestal, and the de-pedestalled excess is $E = 0$ at $\nu = 10^{-7}$ and 10^{-8} under certified estimators. The apparent anomalies are free-slip wall layers in three independent instances: the pedestal itself; the campaign's circulation-depletion law, $\sup |\Gamma|$ depleting as $\nu^{0.51}$ over 4.5 decades on three grid families with a margin of better than 270:1 of law over grid effects (regate sidecars, `boussinesq/regate_*.json`), its argmax pinned to the wall; and the enstrophy total-flux law $\beta = 0.521$, jackknife CI $[0.511, 0.532]$, with wall share rising $0.704 \rightarrow 0.987$ as ν falls, while the corner-box share never exceeds 0.412 against a 0.5 gate. The wall, not the corner, carries the viscous flux. Third, the method, which I take to be the deepest contribution: in a self-similar collapse any statistic that quotients out the collapse frame can delete its own prediction. The campaign's gauge deflation retained only 0.96 to 28.2% of the predicted response, placed the depth-response mode 88 to 98% inside the gauge span, and inflated variance by factors of 10 to 170 (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 5; `ext_inertia_t2/out/bchan_depth/dmx_depth_gates.json`, ANGLE cells). The repair is ontological, not statistical: measure fluxes of material invariants whose inviscid null is zero by theorem, leaving nothing to deflate (measured gauge share $2.0 \times 10^{-4}\%$ to 0.21% against 71.8 to 99%). All verdicts were issued under frozen pre-registered rules; the campaign logged seven instrument catches; zero of three adversarial reviewers refuted the final ladder, and 63.93 GB of run data re-hashed with zero mismatches (campaign ledger, F3 seal).

Tool and computational resource disclosure

This section follows the Leiden Declaration’s first recommendation [1] and its call to normalize responsible disclosure of AI assistance rather than conceal it. It matches the disclosure format used in the two prior notes of this program.

Tools. Dedalus (spectral PDE solver) for the direct numerical simulations; NumPy/SciPy for every estimator, quadrature and fit; Tectonic for typesetting; no proof assistant was used. An AI assistant (Anthropic Claude, used interactively throughout) wrote and debugged the instrument code for the frame repair, the depth-matched interpolator, and the F2 and F3 flux ladders; derived the exact budgets of Sections 5 and 6 together with their boundary terms; carried out the adversarial verification passes; and drafted this paper’s prose.

Division of labour, stated plainly. The instrument code, the budget derivations, and the prose of this paper were produced with AI assistance and would not exist in this form without it. What is the author’s: the epistemic protocol under which every result here was obtained, namely pre-registration of every specification, bar and verdict rule *before* the data that decides it exists, no interim readings, adversarial re-audit instructed to refute, and the standing prohibition on deleting run data; the decisions of what to pursue, what to abandon, what to retract and what to publish; the enforcement of scope against machine-generated drafts, which is why the reach statements in this paper are narrower than the first drafts claimed; and the reframing that produced the paper’s third result. That reframing was a single instruction, that the campaign should measure flux rather than state, issued after three repair cycles on the field statistic had failed. Section 5 is its consequence. Several results here exist *because* the protocol killed earlier machine-produced claims, and those are recorded in Section 8.

What the author can and cannot defend at expert level. The author can defend: the construction of the flux observables and the argument that a material invariant’s drift carries no collapse-frame gauge; why a frame-quotiented statistic can delete its own prediction, which is this paper’s methodological result; the measurement protocol and its failure modes, including the seven instrument artifacts catalogued in Table 5; the reach discipline and why each verdict is bounded where it is; and the retraction record. The author *cannot* defend at expert level the analytical steps of the budget derivations, specifically the production functional and the boundary-term vanishing arguments of Sections 5 and 6, and the Lyapunov-certificate manipulations inherited from the profile campaign. Those were machine-derived under the author’s direction and have not been independently verified by a human expert. Readers should treat them as stated derivations awaiting expert review. The measured ladders do not depend on the derivations being elegant, only on the budget identities holding on the discrete fields, which is gated numerically and reported: closure passes its 10^{-2} bar by two to four orders of magnitude at every rung, and the data and code are public.

Responsibility. Per the Leiden Declaration’s fourth recommendation, responsibility for the correctness and adequacy of everything in this paper, including the machine-assisted portions, rests with the human author.

1 Introduction and results summary

The Luo-Hou axisymmetric corner scenario [14] is the strongest numerical candidate for Euler blowup, and its self-similar profile has since been certified by computer-assisted proof in the nearly self-similar setting [5, 6]. This paper asks the vanishing-viscosity question of that specific certified profile: does viscosity select against the collapse, deform it by a finite amount, or act as a vanishing perturbation, and if it vanishes, is the remainder quasi-static linear response? The question sits in classical territory: Kato’s criterion ties vanishing viscosity to the absence of anomalous boundary dissipation [12], and the Onsager-Duchon-Robert line makes anomalous dissipation the signature of inviscid singularity [8, 9, 16]. For the axisymmetric type-I geometry

specifically, the regularity landscape is set by [10] and [2]. What none of that supplies is a measurement: a certified profile, a viscous ladder against it, and verdict rules frozen before the numbers.

The substrate (Section 2) is my independent, corner-regularized reproduction of the profile, $\alpha = -0.34240 \pm 3 \times 10^{-5}$ against the reference -0.34240009 (`boussinesq/ALPHA_RESULT.md`), a Lyapunov certificate for the linearized corner dynamics, and an archive of inviscid replicates plus viscous ladders on three grid families (D2048/T2048/H1024) spanning $\nu = 10^{-9}$ to 10^{-3} . Every collapse time in this paper uses the fitted $t^* = 3.6185 \times 10^{-3}$, which fits the depth history 2070 times better than the nominal 0.0030 (SSR 5.93×10^{-5} vs 0.1229, `ext_inertia_t2/FLUX_SPEC.md` line 542, the G-F10 disclosure; convention applied throughout `out/flux/FLUX_RESULT.md` G-F10); corrections of record, including that one, are stated in the main text in the numbered ledger of Section 8.

1.1 The measurement problem: the frame is the mechanism

The campaign’s central difficulty is not resolution, it is observability. The sealed Lane G record established that the reduced viscous forcing drives the gauge functionals of the collapse frame 37 times harder than any in-class shape deformation: viscosity’s leading act on this profile is collapse-frame drift (the $37\times$ figure is sealed in `ext_inertia_t2/GAUGE_DRIVE_SPEC.md` line 7 and the campaign ledger, Lane G; verdict artifact `ext_inertia_t2/out/laneg_verdict.json`). The depth-matching instrument confirmed the same fact from the other side: after removing the snapshot-pairing label residual by 99.88%, the signal’s fitted depth-equivalent shift fell by only 14.7 to 44.1%, so the depth content is intrinsic to the signal, not an artifact of pairing (`ext_inertia_t2/out/bchan_depth/`). If the frame is the mechanism, then any statistic that fits and quotients out the frame is structurally blind to the response it was built to detect. That is exactly what happened: the field-statistic channels spent their prediction into the gauge quotient (Section 4), and every downstream symptom, the frame-fit hypersensitivity, the non-null null, the deflation-created degeneracy, follows from that single cause. This is why the campaign’s answer required new observables: fluxes of material invariants whose inviscid null is zero by theorem, so there is no frame to fit and nothing to deflate (Sections 5 and 6).

1.2 The three claims

Claim I (physics). On the certified corner-collapse profile, the viscous response of integrated quadratic functionals is quasi-static linear response at leading order. In the swirl channel (F2), the explicit inviscid quadrature a_2^0 , which contains no operator inverse, no fitted frame, and no gauge basis, reproduces the measured $\nu \rightarrow 0$ leading coefficient to between 4.1×10^{-5} and 1.9×10^{-4} relative, on two grids, over the whole fixed-time window (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, control C7). In the enstrophy channel (F3), the production-subtracted budget ratio on the wall-excised bulk gives, at the deepest common depth window $[1200, \Omega^*)$, $\rho_{\text{bulk}} = 1.243/1.095/1.039/1.029$ (Simpson) and $1.242/1.094/1.036/1.017$ (trapezoid) at $\nu = 10^{-5}$ to 10^{-8} , so $|\rho_{\text{bulk}} - 1|$ falls monotonically from 0.242 to 0.017 on the trapezoid ladder, with $\nu \rightarrow 0$ intercepts 1.0035 ± 0.0099 (trapezoid four-rung fit, jackknife spread) and 0.9999 (Simpson three-rung fit) (`ext_inertia_t2/out/flux3/F3_RESULT.md`, section 3). The joint license, verbatim and no more: linear response at the level of integrated quadratic functionals (C_2 in the r -weighted L^2 pairing, the ω_1 enstrophy flux in the r -weighted H^1 -seminorm pairing), leading coefficient only, sub-leading corrections vanishing no slower than $\nu^{0.4}$. The pointwise mechanism question remains open. Reach: the F2 statement holds at fixed time over $\nu \in [10^{-5}, 10^{-3}]$ and is undecidable from this archive at $\nu \leq 10^{-6}$, where the inviscid floor 1.8×10^{-10} binds; the F3 statement holds at matched depth with primary reach $[10^{-7}, 10^{-5}]$, the 10^{-8} rung decided by the trapezoid control in class agreement (Table 1).

Claim II (physics). There is no finite viscous anomaly in this collapse, and the apparent anomalies are free-slip wall layers, in three independent instances. (1) The registered deformation $d_\infty \approx 0.19$ (`boussinesq/HANDOFF_DEFORMATION.md`) was approximately 85% free-slip boundary pedestal, bounded at or above 83% at the verdict cell; scored de-pedestalled on identical node sets, the excess is $E = -0.0019 \pm 0.0037$ at $\nu = 10^{-8}$ and -0.0073 ± 0.0059 at 10^{-7} , that is $E = 0$ under the certified estimators, with the full ladder $E = 0.2248 (10^{-5}) \rightarrow 0.0588 (10^{-6}, \text{an upper bound}) \rightarrow 0 \rightarrow 0$, monotone (sealed viscous-anomaly record, campaign ledger). (2) The campaign’s circulation-depletion law: $\sup |\Gamma|$ depletes as $\nu^{0.51}$ across 4.5 decades of ν , three grid families, with a margin of better than 270:1 of law over grid effects (at the $\nu = 10^{-6}$ cell, four grid configurations agree to 0.84% while one decade of ν moves the depletion by the factor 3.29; regate sidecars, Section 7); its argmax is pinned at $(r, z) = (1.000000, 0.041667)$, on the wall, at every write of every snapshot series, so it is a wall number and never a corner result (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 6.1). (3) F3’s total-flux law: the unsplit budget ratio ρ_{tot} diverges as ν falls with $\beta = 0.521$, jackknife CI $[0.511, 0.532]$, while the wall share of the dissipation rises $s_{\text{wall}} = 0.704/0.882/0.959/0.987$; the corner-box share never exceeds 0.412 against the frozen 0.5 gate, so corner language is not licensed anywhere in this paper (`ext_inertia_t2/out/flux3/F3_RESULT.md`, sections 3 and 5). The wall carries the viscous flux. Reach: instance (1) over $\nu \in [10^{-8}, 10^{-5}]$ at matched depth with three certified estimators never mixed; instances (2) and (3) as tabulated in Table 1.

Claim III (methods). In a self-similar collapse, any statistic that quotients out the collapse frame can delete its own prediction. The campaign’s gauge deflation retained only 0.96 to 1.51% of the predicted response in the θ channel, 13.3 to 28.2% in the sealed A channel, and about 5% combined; it placed the depth-response mode 88 to 98% inside the gauge span and inflated variance by factors of 10 to 170. That single defect manufactured every downstream symptom: a verdict decided by a frame fit 4.5×10^3 times less sensitive than its statistic, a zero-viscosity null matching the prediction at -0.9534 against the real pair’s -0.2813 , and a deflation-created degeneracy $\rho_D = 0.947$ to 0.969 after deflation against 0.0096 before (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 5; `ext_inertia_t2/out/bchan_depth/`). The repair is ontological, not statistical: measure fluxes of material invariants, whose inviscid null is zero by theorem, leaving nothing to deflate; the measured gauge share fell from 71.8 to 99% down to $2.0 \times 10^{-4}\%$ to 0.21%. Two further exportable tools ship with the paper: the instrument-catch taxonomy, seven catches in one campaign, each an instrument reporting itself and each caught by asking what data enters a gate and whether it can fail (Section 4); and pre-registered frozen specs with VOID-beats-shaky verdict rules and adversarial post-run refutation (Section 10). Zero of three adversarial reviewers refuted the final ladder; every headline number was independently reproduced, and 63.93 GB of run data re-hashed with zero mismatches (campaign ledger, F3 seal).

1.3 Reach and verdicts

Table 1 is the paper’s summary. Every verdict carries its viscosity range and its matching convention; nothing extrapolates beyond the measured range. Depth is quoted as $s = -\log(1 - t/t^*)$ or as the depth label Ω , always on the fitted $t^* = 3.6185 \times 10^{-3}$.

1.4 The honest frontier

What this paper does not claim. The pointwise question, whether $\omega_{1,\nu} - \omega_{1,\text{inv}} = \nu \delta + o(\nu)$ holds anywhere, least of all at the corner ring, remains open at every ν ; both confirmations are integrated statements. The sub-leading exponent is not resolved: F2 measures its swirl-channel correction at $O(\nu^{0.71})$ at the deepest fixed times reached (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 6.3), and F3’s bulk correction fits $q = 0.425$ to 0.438 , at the edge of the frozen 0.4 band (`ext_inertia_t2/out/flux3/F3_RESULT.md`, section 9). Nothing here

Channel / statistic	Verdict (headline number)	ν reach	Convention	Artifact
Profile deformation d_q ; de-pedestalled excess E	No finite anomaly: $E = 0$ at $10^{-7}, 10^{-8}$; pedestal $\approx 85\%$ of regis- tered $d_\infty \approx 0.19$	$[10^{-8}, 10^{-5}]$	matched depth; three certified estimators, never mixed	sealed anomaly record; HANDOFF_ DEFORMATION.md
ω_1 field statistic, θ channel (T2c), de- flated overlap	VOID: instrument uncalibrated in 11 of 12 rung-channel cells; verdict de- cided by a frame fit $4.5 \times 10^3 \times$ less sensitive than the statistic	$[10^{-8}, 10^{-5}]$	fitted collapse frame	out/bchan/ out/bchan_ repair/
ω_1 field statistic, sealed A channel (T2b), deflated overlap	WITHDRAWN: deflation retains 13.3 to 28.2% of the prediction; zero-viscosity null -0.9534 vs. real pair -0.2813	$[10^{-8}, 10^{-5}]$	fitted collapse frame	lane2_ final.json; out/bchan_ depth/
Swirl Casimir C_2 (F2), $a_2(\nu)$	Leading coefficient CONFIRMED to $4.1 \times 10^{-5}..1.9 \times 10^{-4}$ rel., two grids; correction exponent 0.7130 (jackknife 0.029) at the deepest fixed time ($t =$ 2.9×10^{-3} , $s = 1.62$); the F2d crossing-band claim is retracted (Sec- tion 8, correction 5)	decidable $[10^{-5}, 10^{-3}]$; undecidable $\nu \leq 10^{-6}$ (floor 1.8×10^{-10})	fixed time; fit- ted t^*	out/flux/FLUX_ RESULT.md
ω_1 enstrophy flux, bulk (F3, ρ_{bulk})	F3a: linear response confirmed; in- tercept 1.0035 ± 0.0099 (trapezoid) / 0.9999 (Simpson)	primary $[10^{-7}, 10^{-5}]$; 10^{-8} rung by trapezoid con- trol in class agreement	matched depth (Ω -matched); fitted t^*	out/flux3/F3_ RESULT.md
ω_1 enstrophy flux, total (F3, ρ_{tot})	F3b: wall-layer result; $\beta = 0.521$, CI $[0.511, 0.532]$; $s_{\text{wall}} = 0.704 \rightarrow 0.987$; corner share $\leq 0.412 < 0.5$	$[10^{-8}, 10^{-5}]$	matched depth; fitted t^*	out/flux3/F3_ RESULT.md

Table 1: Reach and verdicts for every channel of the campaign. Paths are relative to the campaign root; out/ paths live under ext_inertia_t2/. Depth conventions use the fitted $t^* = 3.6185 \times 10^{-3}$ throughout; the nominal 0.0030 is not the singularity time (Section 8, correction 2).

concerns generic Navier-Stokes blowup; nothing concerns no-slip walls beyond the wall-swap probe of Section 7; growth of perturbations is never used as a falsifier, since blowup solutions are hydrodynamically unstable by theorem [17]. The azimuthal stability program (R2) is parked with its certified Tier-0 route sketched in Section 9. Where the archive cannot decide, I say so: F2 below $\nu = 10^{-6}$, and everything below $\nu = 10^{-8}$ or deeper than the stored collapse window.

Corrections are part of the result, not an appendix: the numbered ledger in Section 8 states, in main text, the in-run circulation gate that never bit on viscous data, the nominal collapse time that is not the singularity time, the duplicate-write defects, the retracted F2d crossing and the retracted reach label, the withdrawn sealed T2b, and the latent F3 code defects, none outcome-changing, all disclosed.

Section 2 fixes the substrate and its certificates; Section 3 the anomaly verdict; Section 4 the instrument arc and the root-cause and catch tables; Sections 5 and 6 the flux observables and the joint license; Section 7 the wall layers; Section 8 the corrections ledger; Section 9 scope and open problems; Section 10 methods.

2 The certified substrate

Every measurement in this paper is made on one flow: the Luo-Hou corner collapse of 3D axisymmetric Euler with boundary [14, 15], and its viscous continuation. This section fixes the system, the certified self-similar profile, the linear-stability certificate, the run archive, and

the time and depth conventions every later section inherits. All solver line references are to `boussinesq/dedalus_axisym.py`.

2.1 The system in Luo-Hou variables

I solve the transformed axisymmetric Euler system in the rescaled variables $u_1 = u^\theta/r$, $\omega_1 = \omega^\theta/r$, $\psi_1 = \psi^\theta/r$, exactly as constructed in the solver (docstring and equation assembly, `dedalus_axisym.py` lines 10 to 24 and 234 to 250; the inviscid form is Luo-Hou PNAS eqs. 2a to 2d):

$$\partial_t u_1 + u^r \partial_r u_1 + u^z \partial_z u_1 = 2 u_1 \partial_z \psi_1 + \nu \Delta_5 u_1, \quad (1)$$

$$\partial_t \omega_1 + u^r \partial_r \omega_1 + u^z \partial_z \omega_1 = \partial_z(u_1^2) + \nu \Delta_5 \omega_1, \quad (2)$$

$$-(\partial_{rr} + \frac{3}{r} \partial_r + \partial_{zz}) \psi_1 = \omega_1, \quad (3)$$

$$u^r = -r \partial_z \psi_1, \quad u^z = 2 \psi_1 + r \partial_r \psi_1. \quad (4)$$

Here $\Delta_5 = \partial_{rr} + (3/r) \partial_r + \partial_{zz}$: for an axisymmetric field f/r the exact viscous operator $(\Delta - 1/r^2)$ on f becomes Δ_5 on f/r , so both u_1 and ω_1 diffuse by $\nu \Delta_5$ with no extra terms (the Hou-Li 5D reduction; solver docstring, lines 217 to 220). Setting $\nu = 0$ recovers the validated inviscid path untouched (lines 209 and 792 to 795).

Geometry. The domain is the annulus $r \in [0.4, 1]$, z -periodic with period $L = 1/6$ (lines 26 to 36 and 74; recorded per run under `geometry` in every `runs/*.json`). The annulus is a localization, not a shortcut: the singularity forms on the solid wall at the corner ring $(r, z) = (1, 0)$, and the initial data decays like $\exp(-30(1-r^2)^4)$, already $\sim 10^{-11}$ at $r = 0.4$ (line 29), so the excluded coordinate axis plays no role and the $3/r$ coefficient is smooth on the domain. The initial condition is the Luo-Hou PNAS eq. 3a data

$$u_1(z, r, 0) = 100 e^{-30(1-r^2)^4} \sin(2\pi z/L), \quad \omega_1 = \psi_1 = 0 \quad (5)$$

(lines 18 and 74 to 86; the fourth power is load-bearing, the first power confines the swirl to a 25x thinner layer and is a different problem, per the correction recorded at lines 76 to 85 and `HANDOFF.md` sec. 2).

Boundary conditions. No penetration is $\psi_1 = 0$ at $r = 1$ and $r = 0.4$ on both the inviscid and viscous paths. The viscous walls are free-slip: zero tangential stress $\tau_{r\theta} \sim r \partial_r(u^\theta/r) = 0$ gives $\partial_r u_1 = 0$, and $\tau_{rz} = 0$ with $u^r = 0$ on the wall gives $\omega^\theta = 0$, hence $\omega_1 = 0$, at both walls (lines 222 to 228 and 247 to 250). Free-slip keeps the inviscid corner driver (no penetration) while adding real viscosity; no-slip is run only as a separate probe (the NS grid families in Table 2). The quartic initial data satisfies $\partial_r u_1 = 0$ at $r = 1$ exactly, so it is compatible with the Neumann condition to machine precision (lines 227 to 228). The free-slip choice is not a detail: Section 7 shows it manufactures every apparent viscous anomaly in this collapse.

Numerics. RealFourier in z , ChebyshevT in r (wall-clustered), IMEX RK443 with implicit diffusion, full 3/2 dealiasing, the Hou-Li exponential filter [11], fp64 throughout, built in Dedalus [4] (lines 46 to 49 and 230 to 231). The elliptic constraint is solved in the r -multiplied form whose only non-constant coefficient is r itself, exact in Chebyshev (lines 38 to 44).

2.2 The transported invariant and the trust referees

The angular momentum $\Gamma = r^2 u_1$ is exactly transported by the inviscid flow: $D\Gamma/Dt = 0$ pointwise for any smooth state, verified as gate G3 on manufactured data to 2.0×10^{-14} (`dedalus_axisym.py` lines 421 to 446; `HANDOFF.md` sec. 2). Under viscosity the budget closes entirely in the swirl channel, $D\Gamma/Dt = \nu r^2 \Delta_5(r^{-2}\Gamma)$, and ω_1 diffusion does not enter it at all (channel scope F2g, `out/flux/FLUX_RESULT.md`). This theorem-backed null is the foundation of the flux observables in Sections 5 and 6.

In-run trust is enforced by pointwise referees, not by volume integrals: the drift of $\sup |\Gamma|$, which sees the collapsing peak, is gated at `GAMMA_TRUST` = 10^{-4} (line 93), and resolution by the spectral tail fraction above $(2/3)k_{\max}$ at `TAIL_TRUST` = 10^{-6} (line 89). One qualification is owed here and stated in full in the corrections ledger (Section 8, item 1): on viscous runs the logged circulation-drift channel was clipped to zero by a `max(signed,0)` at line 653, so the in-run gate never bit on viscous data; the signed values were recovered post hoc from the stored fields and the regate sidecars, and every number in this paper uses the recovered values.

2.3 The certified profile

The blowup exponent of the corner singularity is

$$\alpha = -0.34240 \pm 3 \times 10^{-5}, \quad (6)$$

my independent reproduction by a corner-regularized sparse spectral Newton solver that shares nothing with either existing method: quadratic $\varepsilon_b \rightarrow 0$ extrapolation -0.34240089 against the cross-method reference -0.34240009 , itself known to 3.4×10^{-7} from the Chen-Hou adaptive-mesh line [5] and the PINN Gauss-Newton line [18, 19] (all numbers: `boussinesq/ALPHA_RESULT.md` secs. 1 and 5). The bar is the linear/quadratic extrapolation spread; the corner-degree check moves α by 1.9×10^{-5} , inside the bar. The self-similar transport speed is $c_l = -1/\alpha = 2.9206$ (`boussinesq/README.md` line 43). The solve is answerable to a residual it does not impose: the free gauge identity $c_l = 2\theta_{xx}(0)/\omega_x(0)$ marched to 1.8×10^{-6} (`ALPHA_RESULT.md` sec. 1), and every convincing-but-false root the campaign ever produced violated exactly this functional by orders of magnitude (`ALPHA_RESULT.md` sec. 7). The same substrate carries my sealed direction-regularity measurements (σ_Λ program, `boussinesq/NOVELTY.md`), whose definitional prior art is the sharp $1/2$ -Hölder direction class of Beirão da Veiga and Berselli [3] within the Constantin-Fefferman direction-criteria line [7]; the type-I direction exclusions of Giga-Miura [10] and Barker-Prange [2] are the qualitative results that program quantifies.

2.4 The linear-stability certificate

The linearized self-similar evolution about the profile, realized as a descriptor pencil restricted to $\ker(C_g)$ (quotient dimension $n = 4758$), is Hurwitz at every resolution computed, certified without trusting any eigenvalue: the Lyapunov equation $L^T P + PL = -I$ is solved by Bartels-Stewart with relative residual 4.97×10^{-16} on the shipped record and `chol(P)` succeeding (`ext_inertia_t2/artifacts/pnorm_P.npz`, fields `lyap_residual_rel` and `cholesky_ok`; the three-root sweep reports 4.68×10^{-16} to 4.88×10^{-16} , `ALPHA_RESULT.md` sec. 8). The certified decay margin is

$$m_{\text{cert}} = \frac{1}{2\lambda_{\max}(P)} = 8.8232737664 \times 10^{-6}, \quad \lambda_{\max}(P) = 5.66683085 \times 10^4 \quad (7)$$

(`pnorm_P.npz`). The scope is stated with the number, not after it: Hurwitz means no right-half-plane spectrum in the corner-clamped admissible class (the corner pin deletes the continuum corner ODE and the axis column is an extra Dirichlet condition on the perturbation), it does not mean decay at the Schur abscissa, which sits 3.6×10^4 above the certified rate; and transient growth is real, $89.7 \leq \sup_t \|e^{tL}\| \leq 5807.6$ in the raw collocation norm, which is not similarity-invariant and is labelled as such (`ALPHA_RESULT.md` sec. 8; the upper bound is $\sqrt{\kappa(P)}$ from `pnorm_P.npz`). The shipped weight P is computed at the $\varepsilon_b = 10^{-4}$ working root of the same ladder ($\alpha = -0.34471229$, free identity residual -1.06×10^{-3} ; `pnorm_P.npz` fields `alpha`, `h_id`, `cfig_eps_b`), the discretization family whose $\varepsilon_b \rightarrow 0$ extrapolation is Eq. (6).

family	grid $N_z \times N_r$	ν rungs	unique writes	role
D2048	2048×768	inviscid $\times 2$; $3 \cdot 10^{-5}$ to 10^{-6}	invB 202; 10^{-5} : 198	production pair
T2048	2048×1536	inviscid; 10^{-6} to 10^{-9}	10^{-6} : 197; 10^{-7} : 212; 10^{-8} : 209	deep- ν ladder
H1024	1024×768	inviscid; 10^{-3} to 10^{-9} (9 rungs)	stream record	extended ladder
H1536	1024×1536	10^{-8} , 10^{-9} ; mid 10^{-6} , $3 \cdot 10^{-6}$	stream record	radial refinement
C128/ C256	128×384 , 256×768	$3 \cdot 10^{-6}$ to 10^{-3} (7 each)	C256 10^{-5} : 93; 10^{-6} : 94	cross-grid controls
NS128/ NS256	128×384 , 256×768	$3 \cdot 10^{-6}$ to 10^{-3} (7/6)	stream record	no-slip wall swap
F384	384×1152	10^{-4} , 10^{-6} , $3 \cdot 10^{-6}$	10^{-6} : 46	spot checks

Table 2: The run archive. Grids and initial-condition parameters from the per-run records `runs/*.json` (all runs: $A = 100$, quartic IC, $r_0 = 0.4$, $L = 1/6$); snapshot write counts, after in-memory de-duplication, from `out/flux/fx_inventory.json` (field `nkeep`) and gate G-0 of `out/flux3/F3_RESULT.md`. Runs without a writes entry carry stream records only. The archive directory also holds earlier lattice and amplitude studies not used in this paper.

2.5 The run archive

The measurement archive (Table 2) consists of an inviscid replicate pair and viscous ladders on three primary grids, spanning $\nu \in [10^{-9}, 10^{-3}]$ across families, with supporting cross-grid, radial-refinement, and no-slip families. Two independent inviscid realizations on the production grid (D2048_inv, D2048_invB) provide the replicate floor for every flux observable. The five production runs with full snapshot series store fields over $t \in [0, t_{\max}]$ with $t_{\max}$ between 2.9846×10^{-3} and 2.9997×10^{-3} (`out/flux/fx_inventory.json`); the remaining rungs carry the full stream record of scalar meters at every log step.

Cross-run comparisons are made at matched collapse depth, not matched time. The frozen anchor is $\Omega^* = 1618.3879571332704$, the value of $\sup |\omega_1|$ at the terminal stored write of D2048_1e-5 ($t = 2.999557 \times 10^{-3}$); each run is read at the stored write whose $\sup |\omega_1|$ is closest to Ω^* , and the selected writes land within 0.85 percent of the anchor on each of the four runs recorded in the sealed selection (worst 0.848 percent, D2048_invB; `out/lane2_matched.json`); the later T2048_1e-6 run enters at the same anchor through the four-rung re-derivation, which left Ω^* unchanged (`out/bchan/bchan_ladder.json`, `out/bchan/BCHANNEL_RESULT.md` sec. 2). The finer Ω -matching instrument (cubic interpolation in $\log \Omega$) is introduced with the instrument arc in Section 4.

2.6 The t^* discipline

Collapse depth is indexed by $\Omega = \sup |\omega_1|$ wherever possible. Where a time convention is unavoidable I use $s = -\log(1 - t/t^*)$ with the free-fit singularity time

$$t^* = 3.6185 \times 10^{-3}, \quad (8)$$

which fits the reference run’s own $\Omega(t)$ with SSR 5.93×10^{-5} against 0.1229 for the nominal TSTAR = 0.0030, a factor of 2070 (`ext_inertia_t2/FLUX_SPEC.md` line 542). The nominal 0.0030 is a depth label, not the singularity time, and no collapse time in this paper uses it (`out/flux/FLUX_RESULT.md` depth convention G-F10; `out/flux3/F3_RESULT.md` gate G-D). It remains hard-coded in four sealed lane files, noted there and not edited (Section 8, item 2). For context, the pre-registered reference time from the Luo-Hou MMS table is $\text{TS_REF} = 0.0035056$ (`dedalus_axisym.py` line 87); the free fit sits 3.2 percent above it, and I attach no physics claim to either constant: both are conventions, and every verdict in this paper carries its depth convention explicitly.

2.7 Data hygiene

Two storage defects exist in the archive and both are handled in memory, with the stored trees left read-only. The T2048_1e-7 snapshot series carries 17 duplicate writes re-emitted across its two restarts (229 written, 212 unique, 16 distinct duplicated instants) and its stream carries 301 rows of which 273 are distinct; D2048_invB carries one duplicate write (203 written, 202 unique) (out/flux/fx_inventory.json; out/flux/FLUX_RESULT.md gate G-F0). The duplicates are not bit-exact: the maximum payload relative difference between re-emissions is 2.9×10^{-13} , and an earlier campaign note claiming bit-exactness was corrected on measurement (out/flux3/F3_RESULT.md gate G-0). Every quadrature in this paper consumes the de-duplicated series; the de-duplication counts are re-verified exactly by the F3 gate battery (gate P-1, F3_RESULT.md).

3 No finite viscous anomaly

There is no finite viscous anomaly in this collapse. The quantity I registered as one, a saturating deformation asymptote $d_\infty \approx 0.19$, decomposed under a structural measurement into a free-slip boundary pedestal, at least 83 percent of the full-lattice value at $\nu = 10^{-8}$, plus a remainder that is zero at $\nu = 10^{-7}$ and 10^{-8} within certified bounds. This section states the registered question, why the magnitude axis could not answer it, the estimator discipline that failure forced, and the vetoed measurement that did answer it. The pedestal is instance 1 of 3 on this geometry (catch 1 of Table 5; the three wall laws sit side by side in Fig. 3).

The question. $d_q(\nu)$ is the deep-collapse deformation of the corner profile of the certified Luo–Hou-type substrate [14] away from its inviscid shape, measured by the branch-constrained corner-frame tracker at matched collapse depth, with the inviscid-to-inviscid mapper floor 0.124 ± 0.004 as the reference (boussinesq/DEFORMATION_CURVE.out, re-gated values). The question, posed in the vanishing-viscosity tradition where such limits are classically sensitive to the wall condition [12] and where genuinely viscosity-surviving effects are the defining anomaly [9, 16]: does $d_q \rightarrow \text{floor}$ as $\nu \rightarrow 0^+$, or does it saturate above the floor, a finite viscous anomaly? The saturating fit $d_q = d_\infty + b\nu^p$ over the seven converged plateaus gave $d_\infty = 0.190$, $b = 1.76$, $p = 0.187$, maximum residual 0.0062, robust to including or excluding the demoted 10^{-4} point ($d_\infty = 0.190/0.196$), against the floor 0.124 (boussinesq/DEFORMATION_CURVE.out, 1e-8 discriminator and re-gate sections). That fit is the registered $d_\infty \approx 0.19$ this section retires.

The magnitude axis fails, honestly. I pushed the ν -ladder as far as it could decide. The pre-registered $\nu = 10^{-9}$ decider (grid control: H1536_1e-8 plateau 0.245 ± 0.003 against H1024's 0.250 ± 0.004 , cross-grid gap 0.005 under the 0.015 bar) measured 0.217 ± 0.005 , inside the pre-declared overlap [0.211, 0.227] of the rival bands (pure power 0.207, band 0.187..0.227; saturating 0.231, band 0.211..0.251): non-decisive by the frozen rule, and honored as such (boussinesq/DEFORMATION_CURVE.out, 1e-9 decider section). Adding the point moved the fitted asymptote from 0.196 to 0.166; an asymptote that drifts with the last datum is not pinned. A negative control with rules frozen before it ran (negctl.py; record in boussinesq/DEFORMATION_CURVE.out) then measured the pipeline itself: the pure power law is dead by two independent routes (win rate 0.0 percent whenever either surviving form is true; $\chi^2/\text{dof} = 9.75$ on the bias-corrected refit, against a reject bar of 3), but the two survivors, saturating versus log-to-floor, are mutually confused in 20.8 to 41.5 percent of trials and worst-case recovery is 58.5 percent, under the control's own 60 percent bar. By its own pre-registered rule the campaign's data cannot perform model selection between the survivors; every claim preferring one over the other is unsupported, and I stopped making them. Completing the deep ladder afterwards changed nothing: T2048_1e-7 = 0.3016 ± 0.0028 and T2048_1e-8 = 0.2688 ± 0.0030 (τ -extrapolated) left all fits inside the measured confusion

band, with the pinned log-to-floor form AIC-ahead by only 2.7 and the saturating form's d_∞ endpoint-unstable (sealed viscous-anomaly ledger, `project_parzival_viscous_anomaly.md`). The magnitude curve alone is undecidable by construction, and that is a measured statement, not a mood.

Estimator discipline. The same negative control exposed the campaign's plateau estimator. On every viscous run the deepest-4 mean sits $+0.016$ to $+0.031$ below the $\tau \rightarrow 0$ intercept while the inviscid floor moves only -0.004 : the bias is asymmetric exactly where the anomaly question is posed, since the floor is the reference the question is asked against (`boussinesq/DEFORMATION_CURVE.out`, negative-control Control B). I replaced deepest-4 with the τ -extrapolated intercept as the reported estimator, tightened the convergence bar from $|\text{slope}| \leq 0.05$ to ≤ 0.01 , and required drift-inclusive uncertainties. The campaign thereafter carries three plateau estimators per track (deepest-4, deep-window, τ -extrapolated). They target different objects, and the known extrapolation systematic is large: switching linear to quadratic in τ shifts intercepts by 5 to 13 times the bootstrap standard errors (sealed ledger, `project_parzival_viscous_anomaly.md`). The rule, applied everywhere in this paper: never mix estimators across the points of one curve. Every d_q and E value below names its estimator.

The pedestal mechanism. The decisive move was structural, not statistical. The free-slip condition clamps $\omega_1(r=1) = 0$, while the inviscid reference's vorticity peak sits on the wall itself: argmax at $y = 2.51 \times 10^{-6}$, $|\omega_1| = 1986$ (`boussinesq/noslip/WALL_SWAP_SPEC.md`). A Dirichlet clamp against a reference peaked on the boundary manufactures a full-amplitude, ν -independent difference for free, and the tracker reads it through wall-clustered Chebyshev nodes: the layer occupies $\sim \nu^{1/4}$ of the nodes and enters the unweighted node-lattice norm as $\nu^{1/8}$, a pure node-density exponent (sealed ledger, `project_parzival_viscous_anomaly.md`). Measured under free-slip: the clamped ω_1 wall amplitude is flat in ν (ratio 1.009 ± 0.019 over $\nu = 10^{-3}..10^{-6}$, exponent -0.0003) while the unclamped u_1 channel vanishes (exponent 0.4049 , ratio 15.6), and inside the layer the deformation is pinned at $0.47..0.55$ of the local inviscid amplitude (`boussinesq/noslip/WALL_SWAP_SPEC.md`, measured free-slip section). Two independent controls check the mechanism. First, the wall-condition swap: the frozen prediction was that the flat channel follows the Dirichlet clamp. The no-slip ladder returned the formal W3 verdict (undecided under the frozen rule) because my spec contained a design error, recorded in the spec itself: no-slip against this initial condition ($u_1(r=1) = A = 100$) manufactures its own $\nu^{-1/2}$ Rayleigh sheet, so the clause requiring the unclamped ω_1 to vanish was unsatisfiable. The diagnostic content survived the verdict: the ν -independent wall amplitude moved with the clamp on both grids, both channels, and both implementations (spec and confound: `boussinesq/noslip/WALL_SWAP_SPEC.md`; verdict: sealed ledger). Second, `boussinesq/budget.py`, term-by-term attribution of the vorticity equation in the collapse frame with spectral evaluation and a Poisson re-solve per snapshot: outside the viscous layer nothing ν -independent survives, with measured decay exponents 0.4158 ± 0.0070 (collapse-frame window) and 0.4435 ± 0.0102 (identical-domain wall excision); a collapse-scale persistent residue is excluded at about 59σ (`boussinesq/budget.py` docstring; `ext_inertia_t2/MEANING_MAP_2026-08-11.json`, node B7).

The de-pedestalled ladder. The verdict statistic is E , the de-pedestalled viscous excess: floor and viscous runs scored on identical node sets, with the wall strip $y < 3\delta$, $\delta = \sqrt{\nu t}$, vetoed, at matched collapse depth. The verdict rests on the zeros of this vetoed statistic, not on any three-point fit. Before any verdict number was read, the regenerated matched-grid floor `D2048_invB` reproduced the certified floor to 0.0004 , and the veto scorer reproduced a stored track bit-identically (maximum relative deviation 0.000) (sealed ledger, `project_parzival_viscous_anomaly.md`). The ladder, all values from the sealed ledger:

ν	E ($y < 3\delta$ veto)	reading	vetoed gap $d_q - \text{floor}$
10^{-5}	0.2248	detection: real bulk deformation	survives every veto
10^{-6}	0.0588	upper bound: $E \leq 0.0947$ (2σ)	+0.0123 (+1.4 sd)
10^{-7}	0	zero	-0.0073 ± 0.0059
10^{-8}	0	zero: $E \leq 0.037$ (2σ)	-0.0019 ± 0.0037

Table 3: De-pedestalled viscous excess $E(\nu)$, free-slip, matched collapse depth, floor and viscous runs on identical node sets. Provenance: sealed viscous-anomaly ledger `project_parzival_viscous_anomaly.md` (verdict cells sealed 2026-08-11, the 10^{-6} rung 2026-08-12 from T2048_1e-6, 197/197 snapshots globally clean).

At $\nu = 10^{-8}$ the full-lattice excess is 0.2200 ± 0.0065 while the vetoed gap is -0.0019 ± 0.0037 : at least 83 percent of the registered $d_\infty \approx 0.19$ was boundary-condition pedestal (the round figure carried across the campaign records is about 85 percent, `ext_inertia_t2/MEANING_MAP_2026-08-11.json`, node B5), and the remainder is consistent with zero. Same at 10^{-7} . At 10^{-5} the excess $E = 0.2248$ survives every veto: real bulk finite- ν deformation. The 10^{-6} rung landed outside every cell of its frozen outcome matrix (the monotone decay refuted the resonance cells while the zero clause of the vanishing cell failed at +1.4 sd); I recorded it UNDECIDED-BY-CELL per the matrix’s own out-of-range rule and report it under the one license the matrix grants: as a rung of the ladder, an upper bound and not a detection. The ladder is monotone, $0.2248 \rightarrow 0.0588 \rightarrow 0 \rightarrow 0$, and the measured decay from 10^{-5} to 10^{-6} , $E \sim \nu^{0.582}$, is at or faster than the independent `budget.py` collapse-frame line $\nu^{0.416 \pm 0.007}$ (predicted 0.0863 at 10^{-6} , measured 0.0588, below the line; sealed ledger). The remainder is a vanishing remainder, not a saturation.

Growth is never a falsifier. One discipline governs how every number above may be read. Any blowup solution of 3D Euler is hydrodynamically unstable by theorem [17], and the instability holds within the axisymmetric class itself [13]. Growth of a deviation from the profile near collapse is therefore predicted, not disconfirming: no amplitude-growth statistic on this archive can falsify the collapse, and none is used to. The campaign froze this as a design rule before its stability planning (`ext_inertia_t2/R2_PLAN/R2_RESEARCH_NOTES.md`, guard F7): discriminators must be structural, not amplitude-based. The anomaly verdict conforms. It is carried by a structural decomposition, wall strip against bulk on identical node sets, and by the zeros of a vetoed statistic; the amplitude curve, whose growth with ν is real and was never in dispute, was measured to be unable to decide (58.5 percent recovery against a 60 percent bar) and decides nothing here.

Verdict and reach. No finite viscous anomaly: the registered $d_\infty \approx 0.19$ was about 85 percent free-slip boundary pedestal, and the de-pedestalled excess is $E = 0$ at $\nu = 10^{-7}$ and 10^{-8} (2σ bound $E \leq 0.037$), with $E(10^{-5}) = 0.2248$ real and vanishing by 10^{-7} through the monotone ladder of Table 3. Reach: $\nu \in [10^{-8}, 10^{-5}]$, four rungs with the 10^{-6} rung entering as an upper bound; free-slip walls on this annulus geometry; matched collapse depth; the $y < 3\delta$ vetoed statistic on identical node sets; τ -extrapolated and vetoed-plateau estimators named per point and never mixed. Not claimed: anything below $\nu = 10^{-8}$; anything under no-slip walls, where the wall swap showed the wall physics is manufactured differently (a $\nu^{-1/2}$ Rayleigh sheet) and transfer is forbidden; anything pointwise; anything about generic Navier–Stokes blowup. The same free-slip wall reappears quantitatively in the flux channels, as the swirl-channel circulation depletion $\sup |\Gamma| \sim \nu^{0.51}$ with its argmax pinned to the wall (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 6.1) and F3’s wall-dominated flux share (Fig. 3): three independent instruments, one wall. That is why the linear-response claims of this paper are stated on flux observables whose inviscid null is a theorem, and why any viscous magnitude

read on this geometry must have its wall share checked first.

4 The instrument arc: how a deflated field statistic fails

This section is the methods core of the paper. It reports, in compressed chronology, how the campaign’s field-statistic instrument for the mechanism question failed, how each failure was caught, and what the root cause turned out to be. Every number below is read from a named on-disk artifact; paths are relative to the campaign root. All four instruments in this arc operate on the same ladder, $\nu = 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}$ (quadruples below are always in that order), against the inviscid reference run, at matched collapse depth $\Omega^* = 1618.3879571332704$ defined as the minimum over runs of $\max \sup |\omega_1|$ (the $\sup |\omega_1|$ matching convention, called V1 below), on the certified corner-collapse profile of the Luo and Hou type [14] described in Section 2. Each instrument was pre-registered in a frozen spec before its numbers were computed, and each re-registration inherited the verdict thresholds verbatim.

4.1 The question and the frozen rules

The T2 question asks whether the vanishing viscous remainder is linear response: whether the measured deformation of the collapse fields aligns, as $\nu \rightarrow 0$, with the constructive linear-response solution Δ_1 obtained by solving $L \Delta_1 = -(\text{viscous forcing})$ on the certified profile. The instrument was a deflated field overlap: map the inviscid reference and the viscous snapshot into the corner polar frame, fit a per-member amplitude λ and frame scale L against the root profile, form the pointwise difference dm , project out the two-dimensional gauge basis $G = [S_\lambda, S_L]$ (amplitude and dilation modes), and report the normalized overlap of the deflated dm against the deflated Δ_1 . The frozen rules, written in `ext_inertia_t2/INERTIA_T2_SPEC.md` Lane II and inherited unchanged by every later spec: T2a, overlap monotone increasing in falling ν and ≥ 0.80 at 10^{-8} , fires “the remainder IS linear response”; T2b, overlap saturating ≤ 0.60 , fires “the remainder is NOT linear response”; T2c, anything else, report the curve, no promotion. VOID beats shaky.

The sealed record held a three-rung A-channel T2b, the campaign’s mechanism verdict: overlaps $-0.7770 / -0.2813 / -0.1807$ at $\nu = 10^{-5} / 10^{-7} / 10^{-8}$ (`ext_inertia_t2/out/lane2_final.json`). The original B-channel run under `GAUGE_DRIVE_SPEC.md` had been VOIDed on its own support gate, and the cause was my error: my spec named the measured comparison field u_1 while the packed B slot holds θ -like data, even and quadratic near the corner. I re-registered the corrected object $\theta_{\text{meas}} = (R^2 u_1)^2$ (`ext_inertia_t2/BCHANNEL_SPEC.md`) and added a fourth rung, the new T2048 run at $\nu = 10^{-6}$. The rerun’s reproduction gate passed: the harness reproduced the sealed A-channel curve to a maximum deviation of 4.91×10^{-5} , the four-decimal rounding of the sealed quote, and Ω^* was unchanged by the fourth run (`ext_inertia_t2/out/bchan/BCHANNEL_RESULT.md` §2, `out/bchan/bchan_ladder.json`). The four-rung ladders came out as follows (θ : `out/bchan/bchan_ladder.json`, field `resultsB`; A_{sealed} : the sealed-forcing four-rung ladder tabulated in `out/bchan/BCHANNEL_RESULT.md` and reproduced to deviation 0.0 in `out/bchan_repair/REPAIR_RESULT.md` §1, G-R1):

$$\theta : +0.2521, +0.0567, +0.2066, +0.7681 \quad A_{\text{sealed}} : -0.7770, -0.5004, -0.2813, -0.1807.$$

Under the frozen rules the θ channel returned T2c (non-monotone at the first step, $+0.7681 < 0.80$ at the last), and the extended A-channel ladder was monotone with shrinking gaps $+0.277 / +0.219 / +0.101$, apparently strengthening the sealed T2b.

The noise gates said otherwise, in the same document. Every θ rung was SOFT: signal-to-floor ratios $2.88 / 1.82 / 0.72 / 1.64$ against the adjacent-write floor pair, with the 10^{-7} rung below its own noise floor (`out/bchan/BCHANNEL_RESULT.md` §2, G-B3). And the floor pair itself, which carries zero viscosity contrast, scored a deflated overlap of $+0.9165$ at the 10^{-5}

rung (out/bchan/bchan_ladder.json, field resultsB/D2048_1e-5/floor_overlap), larger than three of the four signal rungs. A statistic whose null control outcores its signal is not yet a measurement. The next three subsections are the anatomy of why.

4.2 Catch 6: the frame fit decided the verdict

The frame scale L was fitted independently for the two members of each pair by maximizing correlation against the root profile on a geometric grid with step ratio 1.028286. At the $\theta/10^{-8}$ rung the two fitted frames landed exactly one grid step apart, and scanning L_i at fixed L_x exposed the defect: across the scan the fit objective moves by 3.56×10^{-4} while the verdict statistic moves by 1.594 and flips sign, a sensitivity ratio of 4.5×10^3 (ext_inertia_t2/out/bchan/bchan_lscan.json, fields corr_span_1e8 = 3.559×10^{-4} and overlap_span_1e8 = 1.594). A uniform multiplier of 0.97 on L_i , invisible to the fit (the correlation moves by 1.6×10^{-4} over the 0.94 to 1.02 band), moves the θ ladder to +0.1091/ - 0.2454/ - 0.6224/ - 0.3370, which fires T2b (bchan_lscan.json, uniform_ladders). T2c and T2b for the θ channel were separated by less than one grid step of a fit four orders of magnitude less sensitive than the statistic it fed. Nor was the fit internally consistent: the fitted ratios L_i/L_x across the ladder are 1.29994/1.12648/1.00000/1.02829 (bchan_lscan.json, fitted_frames), so the near-coincidence at 10^{-8} was one point of a scattered set, and the 30 percent split at 10^{-5} was being read as if it were a viscous difference. One provenance correction belongs here in the main text: the scan behind these numbers was originally computed but not saved; I regenerated it into bchan_lscan.json, which reproduces every quoted digit and self-asserts against them. The frozen rules were applied once, before the scan, and were not touched.

Verdict, honest and with its reach: the θ channel is T2c under the frozen rules on $\nu \in [10^{-8}, 10^{-5}]$ at matched depth in convention V1, and I did not promote it, because the deciding rung sits inside the instrument's own indifference band.

4.3 The frame repair, and what it exposed

The repair spec (ext_inertia_t2/FRAME_REPAIR_SPEC.md) named the structural defect before any repaired number existed: refining the fit would have delivered a precise value of an undetermined quantity. Two defects, both visible in the sealed artifacts. First, the difference was taken across two evaluation loci: with independent per-member fits, the two members are sampled at different physical radii wherever $L_i \neq L_x$, so dm acquires a term proportional to $(L_x - L_i) \partial(\text{field})/\partial(\text{radius})$, which is large near a collapsing corner; at $\theta/10^{-8}$, tuning the two frames into agreement removes 82 percent of the difference norm. Second, the reference frame carried a manufactured ν dependence: the same inviscid snapshot, refitted under each rung's ν -dependent window, drifted 6 percent in fitted L . The repair: RF-1, a common evaluation frame, both members mapped at one L_c ; RF-2, that frame frozen once on the reference under a ν -independent window ($L_{\text{ref}}(A) = 6.927471 \times 10^{-4}$, $L_{\text{ref}}(B) = 5.340674 \times 10^{-4}$); RF-3, every overlap quoted as its range over the correlation-indifference band, never as a point; RF-4, continuous (Brent) refinement recorded explicitly as hygiene that would not by itself have repaired anything. The reimplementaion reproduced the sealed ladders in legacy mode to deviation exactly 0.0 before anything repaired was quoted (ext_inertia_t2/out/bchan_repair/REPAIR_RESULT.md §1, G-R1).

The structural part of the repair worked. The frame-sensitivity ratio fell from 4.5×10^3 to 34 to 294 across all twelve rung-channel cells, and no overlap interval straddles a decision threshold over a band two to three times wider than the one that broke the legacy fit (REPAIR_RESULT.md §1, G-R2). The shift-artifact accounting confirmed the diagnosis where its evidence lived: at $\theta/10^{-8}$, 88.5 percent of the legacy difference norm was frame mismatch (the independent scan estimated 82 percent), and the verdict-blocking +0.7681 became $[-0.0916, -0.0746]$ at a signal-to-floor ratio of 0.04 (REPAIR_RESULT.md §1 and §2). At $\theta/10^{-5}$ and 10^{-6} the repaired dm is

larger, not smaller: the legacy per-member fits had also been absorbing genuine inviscid-viscous scale difference. The legacy instrument erred in both directions, rung by rung.

What the repair exposed, it could not fix. The calibration gate pushed the adjacent-write floor pair, zero viscosity contrast, through the identical repaired instrument, and the null came out at $-0.9516/-0.9685/-0.9733/-0.9747$ against the sealed A-channel forcing, exceeding the signal magnitude at every rung, with size signal-to-floor ratios $4.23/2.76/2.11/1.72$ (REPAIR_RESULT.md §1, G-R3). Eleven of twelve rung-channel cells were UNCALIBRATED; the only calibrated cell (θ at 10^{-5}) passed its null clause by 0.027. Under the frozen rules all three channels landed R-VOID: a rule cannot be carried by rungs the instrument cannot calibrate, and a caveat does not repair a measurement.

The registered prediction P-R4, that the sealed A-channel T2b would survive the repair, failed, and per the spec that failure is reported above everything else in the repair record. The sealed curve moved by up to 0.238 ($-0.7770/-0.5004/-0.2813/-0.1807$ became $-0.6391/-0.5878/-0.5188/-0.3299$), and the saturating shape inverted: gaps $0.277/0.219/0.101$ became $0.051/0.069/0.189$ (REPAIR_RESULT.md §0). The saturation the sealed record described was a property of the legacy per-member frame fits, not of the pair of flow fields. I withdrew the sealed A-channel T2b. Reach of the withdrawal: $\nu \in [10^{-8}, 10^{-5}]$, matched depth, convention V1; the repaired instrument returns R-VOID, which withdraws the verdict without refuting its content, because a design that fails null calibration in every channel cannot support T2b, T2a, or T2c.

4.4 The null was never a null

The disqualifying number had been sitting in the campaign’s own sealed file. At the $T2048\ 10^{-7}$ rung, `ext_inertia_t2/out/lane2_final.json` records `overlap = -0.2813` beside `floor_overlap = -0.9534`: the zero-viscosity adjacent-write pair reproduced the predicted forcing 3.4 times better than the real viscous pair did. I had labeled that column a noise floor, and the label is what let the number sit unread for the length of the campaign. It is not a floor. On the floor support the primary depth-response mode D is the adjacent pair’s own dm divided by $d \log \Omega$, so the deflated cosine of D against Δ_1 must equal minus the recorded null overlap, and it does, in all twelve cells, to a maximum deviation of 2.22×10^{-16} (the $A/10^{-5}$ cell reads `cos = 0.9515680489408568` against recorded null -0.9515680489408567 ; `ext_inertia_t2/DEPTH_MATCH_SPEC.md` §0, Defect 2). The recorded “noise floor” was a direct, already-sealed measurement of the angle between pure collapse-depth response and the predicted viscous response. I retired the word floor for it and relabeled the adjacent-write pair the depth positive control; it stays in every report as the permanent record of what the sealed instrument was measuring.

4.5 Depth matching, and what it exposed

The third instrument attacked the defect the repair revealed. Nearest-stored-write matching leaves residual depth offsets $d\Omega_{\text{pair}} = -0.008411/-0.001922/-0.013851/-0.004351$ against the adjacent-write offset -0.017960 , and the finest stored write spacing near the matched depth (0.01796 down, 0.01882 up) exceeds every one of the four residuals: no choice of stored write can depth-match anything, so interpolation is a precondition, not an option (`DEPTH_MATCH_SPEC.md` §0; `out/bchan_depth/route1_budget.log`). The cost of ignoring this was measured: a one-write depth displacement of the reference moves the A_{sealed} overlap by $+0.1118/+0.1447/+0.2051/+0.2875$ across the ladder, while the entire span of that ladder is 0.309 ; the offset at the 10^{-8} rung alone reproduces 93 percent of the ladder’s range (`DEPTH_MATCH_SPEC.md` §0, Defect 1). The A-channel ladder shape as sealed was, to leading order, the depth-mismatch signal.

Route I rebuilt the reference at each viscous member’s own stored Ω by four-point Lagrange interpolation in $u = \log \Omega$, anchored by Ω value and never by index, on the single stencil $\Omega = \{1577.3040, 1602.8023, 1632.1156, 1662.8364\}$, after de-duplicating restart re-emissions by running maximum (1 + 4 writes dropped); the viscous members were never interpolated (out/bchan_depth/dmx_depth_gates.json, G_D0, G_D1). The machinery is good: leave-one-out reconstruction error is 0.006 to 0.498 percent of the signal size at every cell against a 5 percent bar (dmx_depth_gates.json, G_D3 and DM_8_LOO), the predicted label-residual suppression in channel A is $92\times$ to $668\times$ (out/bchan_depth/route1_budget.log), and the direct gate measured the depth pair collapsing by $745\times$ to $848\times$ through the depth-matched instrument (810.7 to 848.4 in channel A, that is, the depth-label residual fell by 99.88 percent; dmx_depth_gates.json, G_D2).

What depth matching exposed is that the labels were not the disease. The fitted depth-equivalent shift in the signal, $\varepsilon_{\text{eff}} = \langle dm_{\text{defl}}, D_{\text{defl}} \rangle / \|D_{\text{defl}}\|^2$, fell only 14.7/12.6/44.1/31.6 percent in A_{sealed} (from $-0.061299/-0.052700/-0.037167/-0.018645$, out/bchan_depth/depth_mode.json, to $-0.052306/-0.046040/-0.020778/-0.012756$, dmx_depth_gates.json AN-GLC cells) while the label residual itself fell by 99.88 percent. The depth-like content of the viscous difference is intrinsic to the fields; removing the depth-label mismatch barely touches it.

Two further findings closed the design. First, the depth gauge is ambiguous by more than the residual Route I removes: matching on $\sup |\omega_1|$ (V1) versus matching on the fitted collapse scale L (V2) are both defensible definitions of “same collapse depth,” and they disagree by $0.4\times$ to $45\times$ the label residual and by $240\times$ to $4100\times$ the post-interpolation residual (DEPTH_MATCH_SPEC.md §0, Defect 3). The stencil-legality gate passes at all four rungs under V1 and fails at four of eight cells under V2 (the V2 target at $A/10^{-5}$ is not even bracketed by the reference; the $A/10^{-7}$, $B/10^{-5}$ and $B/10^{-8}$ targets fall outside the fractional-position window; dmx_depth_gates.json, G_D1), so the matching is not well posed across the gauge orbit. Second, the validation battery blocked any re-adjudication before a verdict could form (dmx_depth_gates.json, _blockers, SCOPE): the within-run depth-recovery control PC-1 fails its $[0.90, 1.10]$ window (A ratios 0.845 to 0.922, one cell of eight passing; θ ratios 0.661 to 0.808, none passing); the injected cross-run depth control PC-2 passes in channel A at the two gate amplitudes $|\varepsilon| = 0.014$ and 0.005 (recovery 0.858 to 0.930) and fails at all four θ rungs (0.684 to 0.819 against a $[0.80, 1.20]$ window); the specificity control PC-4 fails on the full pipeline; the disjoint-stencil spread N2 exceeds the interpolation null N1 at every cell (ratios 2.33 to 3.86), so the stencil is not converged at the operating point; and PC-3, the control that converts a null result into a measured upper bound, could not run because it requires a second inviscid replicate (invC) that does not exist. I formed no depth-matched overlap and applied no verdict rule; the run is validation gates only, and it is recorded that way in its own artifact.

4.6 The root cause

Table 4 is the arc’s autopsy, computed cell by cell from ext_inertia_t2/out/bchan_depth/dmx_depth_gates.json (ANGLE block). For each channel and rung: ρ_D is the magnitude of the cosine between the deflated depth mode D and the deflated prediction Δ_1 ; the raw column is the same cosine before gauge deflation; the surviving prediction fraction is $\|\Delta_{1,\text{defl}}\|^2 / \|\Delta_1\|^2$, the energy fraction of the prediction that survives projecting out $[S_\lambda, S_L]$; the span column is the energy fraction of D inside the gauge span; VIF is the variance inflation factor of the prediction regressor in the four-term amplitude regression $dm = c_\lambda S_\lambda + c_L S_L + c_D D + c_P \Delta_1 + r$.

The table says three things, and together they are the root cause. The gauge deflation retains 0.96 to 1.51 percent of the θ prediction’s energy, 13.3 to 28.2 percent of A_{sealed} ’s, and about 5 percent of A_{comb} ’s: the statistic deletes 71.8 to 99.0 percent of the very response

channel	ν	ρ_D (defl.)	ρ_D (raw)	surv. pred. [%]	D in span [%]	VIF
θ	10^{-5}	0.8427	0.9645	0.959	95.54	169.8
θ	10^{-6}	0.7114	0.9342	1.506	96.33	64.9
θ	10^{-7}	0.7872	0.9457	1.441	97.18	88.4
θ	10^{-8}	0.7856	0.9432	1.396	97.52	90.2
A_{sealed}	10^{-5}	0.9469	0.0096	13.34	90.23	56.8
A_{sealed}	10^{-6}	0.9560	0.2582	23.69	88.41	42.8
A_{sealed}	10^{-7}	0.9693	0.2968	26.70	88.39	54.5
A_{sealed}	10^{-8}	0.9651	0.3314	28.21	88.55	45.8
A_{comb}	10^{-5}	0.5935	0.8659	4.949	90.23	11.1
A_{comb}	10^{-6}	0.5296	0.8482	5.134	88.41	10.1
A_{comb}	10^{-7}	0.6370	0.8722	5.039	88.39	12.6
A_{comb}	10^{-8}	0.6399	0.8705	4.925	88.55	13.0

Table 4: The root-cause table. All values from `ext_inertia_t2/out/bchan_depth/dmx_depth_gates.json` (ANGLE cells; full precision in the file), on the per-rung support with the frozen frames of Section 4.3. ρ_D : alignment of the deflated depth mode with the deflated prediction; raw: the same alignment before deflation; surv. pred.: prediction energy surviving the two-mode gauge deflation; D in span: depth-mode energy inside the gauge span; VIF: variance inflation of the prediction amplitude c_P .

it is asked to detect before it looks at anything. The depth mode sits 88.4 to 97.5 percent inside the gauge span, so the deflation nearly deletes it too, and the two surviving slivers are nearly parallel: in A_{sealed} , ρ_D after deflation is 0.947 to 0.969 against a raw alignment of 0.0096 at 10^{-5} (rising to 0.331 at 10^{-8}). The degeneracy between “viscous linear response” and “the two snapshots sit at different collapse depths” is not in the fields; it is manufactured by the deflation. And the regression that tries to separate the two directions pays for it with variance inflation factors of 10 to 170. Every upstream symptom follows from this table: the 4.5×10^3 frame hypersensitivity, the -0.9534 non-null null beside the -0.2813 signal, the SOFT rungs, the uncalibrated cells. In a self-similar collapse the frame is not bookkeeping to be quotiented away; the sealed Lane G result identifies frame drive as how viscosity acts (`ext_inertia_t2/out/laneG_verdict.json`), so a statistic that quotients out the collapse frame is structurally blind to exactly the response it was built to find, and what it then measures is the angle between two deflation residues.

One route was measured and forbidden rather than taken. Projecting D out as a third deflation mode would delete 90.5 to 95.0 percent of the remaining A_{sealed} prediction energy and drive the null residual to 2.5×10^{-16} , so every rung would pass calibration by construction (DEPTH_MATCH_SPEC.md §1). I executed that counterfactual once, recorded it, and barred the route: D is measured and regressed, never deflated.

4.7 Verdicts

Stated once, honestly, each with its reach; all on $\nu \in [10^{-8}, 10^{-5}]$ at matched depth $\Omega^* = 1618.3879571332704$, convention V1.

- θ channel: T2c on the sealed instrument (`out/bchan/bchan_ladder.json`); R-VOID on the repaired instrument, deciding gate G-R3, rungs $10^{-6}, 10^{-7}, 10^{-8}$ uncalibrated (`out/bchan_repair/repair_repaired.json`). The sealed T2c is retired as an instrument artifact: 88.5 percent of the deciding rung’s difference norm was frame mismatch.
- A_{sealed} : the sealed T2b is withdrawn; R-VOID, all four rungs uncalibrated (null -0.95 to -0.97 exceeds the signal everywhere). Withdrawal is not refutation: no channel supports T2a, T2b, or T2c under the repaired instrument.

- A_{comb} : R-VOID, same structure.
- Depth-matched re-adjudication: blocked at its own validation gates (PC-1, PC-4, N2, gauge orbit; PC-3 unrunnable without `invC`); no overlap formed, no rule applied (`out/bchan_depth/dmx_depth_gates.json`, SCOPE).

The structural result, which is a result about the design and not a failure to reach a threshold: a matched-depth pair comparison against an inviscid reference cannot separate viscous linear response from collapse-depth response at these depths. The pure-depth control scores 0.9516 to 0.9747 against the prediction while the real viscous pairs score 0.33 to 0.64 (`DEPTH_MATCH_SPEC.md §0`; `REPAIR_RESULT.md §1`). That is why the campaign’s answer to T2 does not come from a field statistic at all. The repair that worked is ontological, not statistical: measure fluxes of material invariants, whose inviscid null is zero by theorem, leaving nothing for a frame quotient to delete. The flux observables of Sections 5 and 6 carry a measured gauge share of 2.0×10^{-4} to 0.21 percent of signal, against the 71.8 to 99.0 percent deleted here (`ext_inertia_t2/out/flux/FLUX_RESULT.md`).

4.8 The seven catches

Across the campaign, seven times, an instrument reported itself and was caught. Table 5 is the taxonomy. Every catch reduces to the same two questions, asked of a gate that was passing: what data actually enters this number, and can this gate fail? A gate that cannot fail is not a gate, and a statistic whose input contains its own output is not a measurement. None of the seven was found by re-deriving mathematics; all seven were found by tracing data flow.

5 Flux observables with theorem-backed nulls: the swirl channel (F2)

Two scope statements come first because the pre-registration (`FLUX_SPEC.md`, rule F2g) makes them mandatory in the first paragraph of any write-up that quotes these numbers. First, the channel: the Γ budget below closes entirely in the swirl channel, $D\Gamma/Dt = \nu L_{\text{sw}}\Gamma$, and ω_1 diffusion does not enter it at all, while the frozen T2 forcing was explicitly the ω_1 Laplacian (`lane2_solve.py`, declaration D2: forcing support on ω_1 rows only). F2 and T2 probe different forcing channels; nothing in this section resolves T2, and T2 remains VOID and unanswered until Section 6 takes up its channel directly. Second, the reach: every fixed-time exponent quoted here lives at $\nu \in [10^{-5}, 10^{-3}]$; at $\nu \leq 10^{-6}$ the archive cannot decide the question because the response falls to the archive’s own inviscid floor (about 1.8×10^{-10} , `out/flux/FLUX_RESULT.md §6.2`). Both statements are load-bearing and neither is softened below.

5.1 The observable and its theorem null

The instrument arc of Section 4 ended in a diagnosis, not a repair: in a self-similar collapse, a field value at a point is almost entirely frame motion, so any frame-quotiented field statistic deletes its own prediction. The move I made is to stop measuring state and measure the drift of an exact material invariant, whose inviscid null is zero by theorem, so that there is no frame in the observable and nothing to deflate.

For the solver’s own system the swirl angular momentum $\Gamma := r^2 u_1$ obeys, exactly,

$$\frac{D\Gamma}{Dt} = \nu L_{\text{sw}}\Gamma, \quad L_{\text{sw}} := \partial_{rr} - \frac{1}{r}\partial_r + \partial_{zz}, \quad (9)$$

so at $\nu = 0$ the quantity Γ is transported unchanged (`boussinesq/dedalus_axisym.py:422`, with the conservation rationale at lines 52 and 93). The meridional flow is divergence free with

#	what the instrument reported	what it actually was	deciding number	artifact
1	a finite viscous anomaly, $d_\infty \approx 0.19$	$\sim 85\%$ free-slip wall pedestal; de-pedestalled excess $E = 0$ at $\nu = 10^{-7}, 10^{-8}$ (2σ bound $E \leq 0.037$)	85%	out/qm/EVAL_POINT_AUDIT.md §4
2	corner gauge response dcl as physics	node-density divergence: $dcl \sim deg_0^{+4.0399}$; the $q \cdot m$ factorization gives $q=2$ (Chebyshev) times intrinsic $m \approx 2.02$; a designed $q=1$ family halves the exponent to $+1.9891$, a $q=0$ family flattens it to -0.2151 ; invariant $dcl \xi_1^2 = 26.375$ to 27.124 (2.80%)	$+4.0399$	out/s1_report.txt; out/qm/ANALYSIS.txt
3	corner gauge response dcw as physics	same node-density class	$+4.1737$	out/s1_report.txt
4	an apparent $\nu^{1/8}$ scaling law	the sampler's own geometry: density exponent 0.125 against the sealed physical schedule exponent 0.5 ; same variable, apparatus origin	0.125 vs 0.5	JUMP/chains/C12/journal.jsonl
5	the $\sigma_{\min}$ direction as 99.989% gauge	self-parallelism: the projection basis was not independent of the vector under test; against the declared one-mode basis (root scaling) the gauge fraction is 5.2×10^{-6}	fraction 0.99989 vs 5.2×10^{-6}	out/vmin/vmin.json; out/vmin/vmin_corrected.json
6	a T2c verdict for the θ channel	the frame fit decided a verdict its objective could not see: fit moves 3.6×10^{-4} , statistic moves 1.594	4.5×10^3	out/bchan/bchan_lscan.json
7	gate G-D2 PASS: depth mismatch removed	a self-comparison: the gate holds out a reference write and rebuilds it from the reference's own disjoint writes, so both members of the differenced pair are the reference run; the PASS certifies interpolation machinery, not cross-run matching, and the cross-run controls (PC-1, PC-2 in θ , PC-4) failed beside it	PASS at $811 \times$	out/bchan_depth/dmx_depth_gates.json G_D2

Table 5: The seven instrument catches, in campaign order. Paths relative to `ext_inertia_t2/` except where given in full. Catch 1 is pedestal instance 1 of 3 (Section 3); catches 2 to 4 are the node-density class, diagnosed by the frozen band ($|d \ln X / d \ln deg_0| < 0.2$ flat, > 0.5 artifact) and closed by the $q \cdot m$ factorization; catch 6 is Section 4.2; catch 7's number is the $A/10^{-5}$ cell's collapse factor 810.7.

respect to the measure $r dr dz$, so every functional of the distribution of Γ has identically zero inviscid drift. The primary observable is the $p = 2$ Casimir

$$C_2(t; \nu) := \int_0^L \int_{r_0}^1 |\Gamma|^2 r dr dz, \quad D_2 := \frac{C_2(t) - C_2(0)}{C_2(0)}, \quad a_2 := \frac{|D_2| - F_2}{\nu}, \quad (10)$$

with F_2 the inviscid floor, taken per gate G-F5 as the maximum over inviscid realizations on that grid. The exact single-snapshot budget derived from (9) is

$$\begin{aligned} \frac{dC_2}{dt} &= \nu (\text{BULK}_2 + \text{WALL}_2), & \text{BULK}_2 &= -2 \int |\nabla \Gamma|^2 r dr dz, \\ \text{WALL}_2 &= 2 \int_0^L [|\Gamma(1, z)|^2 - |\Gamma(r_0, z)|^2] dz, \end{aligned} \quad (11)$$

and the linear-response prediction operator is the explicit quadrature on the inviscid run's own stored fields,

$$a_2^0(t) = -\frac{1}{C_2(0)} \int_0^t S_2[\Gamma^0(t')] dt', \quad (12)$$

with no operator inverse, no linearized solve, no fitted frame, and no gauge basis, so 100% of the prediction enters the comparison (FLUX_SPEC.md §1.1–1.2, §3). Quasi-static linear response, $\Gamma^\nu = \Gamma^0 + \nu \Gamma^{(1)} + o(\nu)$, then predicts $a_2(t; \nu) = a_2^0(t) + O(\nu)$; the discriminating statistic is the deviation exponent α in the fit $a_2(\nu) = A + B\nu^\alpha$ ($\alpha = 1$ linear response, $\alpha = 1/2$ square-root singular), because the naive exponent $q = d \log |D_2| / d \log \nu$ tends to 1 for every $\alpha > 0$ and cannot return the unwelcome answer. That degeneracy is not a conjecture: over all 191 adjacent-rung pairs across every ladder, depth, and convention, $q_2 \in [0.9960, 1.0351]$, and 191 of 191 pairs sit inside $[0.95, 1.05]$ (out/flux/judge/fx2_q2.json). This construction is in the spirit of the conservation and dissipation-anomaly literature, where integrated invariants rather than pointwise fields carry the sharp statements [8, 9, 16].

5.2 Nothing to deflate: proof, and then measurement

C_2 is a functional of the distribution of Γ under the transport measure. It contains no coordinate, no length, no amplitude scale, no fitted λ or L , no lattice, and no reference field, so the two-mode gauge basis that deleted 71.8–99.0% of the sealed field-statistic prediction has no object to act on. Stronger, and this part is a theorem rather than an absence: the collapse frame's entire motion is generated by the inviscid flow, which is volume preserving with $D\Gamma/Dt = 0$, so every frame motion the collapse performs produces exactly zero drift in C_2 . One honest limit is carried with the claim: for a per-particle Lagrangian flux observable the statement is false as an absolute (a prior pass measured 19.3–27.6% of that observable's energy inside the same two-mode gauge span), and that measurement applies to the Lagrangian instrument, not to C_2 (FLUX_SPEC.md §2.1; out/flux/FLUX_RESULT.md §3, C0).

The archive then makes the proof a measurement, because its inviscid runs are pure-gauge experiments: their entire state change is frame motion with zero physics in it, by theorem. On D2048_invB the frame runs the full campaign window: $\sup |\omega_1|$ goes from 0 to 1758.5, the corner-box field moves by 439% of $\|\Gamma\|$, and $\sup |\nabla \Gamma|$ sharpens by a factor 25.1, while the measured inviscid drift of C_2 at the matched depth $\Omega^* = 1618.3879571332704$ is 1.820×10^{-10} (out/flux/judge/fx2_gauge.json). The four inviscid realizations give floors 1.798×10^{-10} (D2048_inv), 1.820×10^{-10} (D2048_invB), 1.802×10^{-10} (T2048_inv), and 6.919×10^{-10} (H1024_inv); the one genuine replicate pair, D2048_inv versus D2048_invB, agrees to 1.2% at matched depth, and the two different grids D2048_inv versus T2048_inv agree to 0.22% (out/flux/FLUX_RESULT.md C0, R8). Against viscous signals of 8.75×10^{-8} to 9.05×10^{-5} on the same grids, the gauge share of the signal is $2.0 \times 10^{-4}\%$ to 0.21%, versus 71.8–99.0%




Figure 1: Gauge share of the observable, per rung, log scale. Field statistics: share of the predicted response deleted by the two-mode gauge deflation, 100% minus the surviving prediction fraction, for the θ , A_{sealed} , and A_{comb} channels (`ext_inertia_t2/out/bchan_depth/dmx_depth_gates.json`, ANGLE cells). Flux observable C_2 : the measured inviscid floor as a share of the viscous signal on the same grids (`ext_inertia_t2/out/flux/judge/fx2_gauge.json`, `gauge_pct`). The field statistics spend 71.8 to 99.0% of their prediction into the gauge quotient; the flux observable’s gauge share is $2.0 \times 10^{-4}\%$ to 0.21%.

for the deflated field statistic it replaces (`out/flux/judge/fx2_gauge.json`; Figure 1). The improvement is a factor 4.8×10^2 at the worst rung and 4.9×10^5 at the best, and it is not an argument: it is the inviscid run’s own measured response put through the identical pipeline.

5.3 The exact budget closes on the discrete fields

Gate G-F1c is the gate whose failure would have voided the whole design: the theorem null is worthless if the budget (11), free-slip boundary term included, does not hold on the discrete fields. Computed on each viscous run’s own stored snapshots, the closure residual $|\Delta C_2^{\text{meas}} - \nu \int S_2 dt| / |\Delta C_2^{\text{meas}}|$ over the reported window ($t \geq 2.0 \times 10^{-3}$) is 6.2×10^{-6} at $\nu = 10^{-5}$, 3.2×10^{-5} at 10^{-6} , 3.1×10^{-4} at 10^{-7} , and 3.4×10^{-3} at 10^{-8} , against a frozen bar of 10^{-2} (`out/flux/flux_ladder.json`, `gates.G-F1c_closure`). One disclosure belongs beside that: the whole-window maximum always sits at the shallowest evaluable time $t = 2.04 \times 10^{-4}$, where the closure numerator is ν -independent while the denominator scales with ν , and at $\nu = 10^{-8}$ that whole-window value is 2.37×10^{-2} , above the bar. The gate is written at every reported t and is applied as written; both numbers are given and the bar is not moved (`out/flux/FLUX_RESULT.md` §1.0).

The sign check closes the same loop from the other side. D_2 is negative at every viscous rung and positive on the inviscid null, recovered directly from stored fields on all five snapshot series; the budget requires this, since $\text{BULK}_2 < 0$ strictly and the measured wall share $|\text{WALL}_2| / |\text{BULK}_2| = 8.1 \times 10^{-3}$ is far too small to flip it (`out/flux/FLUX_RESULT.md` §2.1).

Table 6: F2 matched-time deviation exponent on the H1024 ladder, G-F5-gated rungs ($\nu \in [10^{-5}, 10^{-3}]$ at the labelled times). Fit $a_2(\nu) = A + B\nu^\alpha$; profile interval at $\text{RSS} \leq 2\text{RSS}_{\min}$; jk is the drop-one jackknife spread, defined only for $n \geq 4$. The deterministic instrument budget propagated through the fit is below the 5×10^{-4} scan step at every quoted row, so the jackknife dominates and is the uncertainty to quote. Source: `ext_inertia_t2/out/flux/flux_ladder.json` (ladders block); prose values `out/flux/FLUX_RESULT.md` §4.2.

t	s	α_2	profile	jk	label
1.0×10^{-3}	0.323	2.5000 (edge)	[0.643, 2.500]	2.4500	NO-FIT
1.5×10^{-3}	0.535	0.9780	[0.957, 1.000]	0.2550	VOID-3
2.0×10^{-3}	0.805	0.8760	[0.876, 0.876]	n/a	F2a-BLOCKED
2.2×10^{-3}	0.936	0.8435	[0.844, 0.844]	n/a	F2c
2.4×10^{-3}	1.088	0.8125	[0.813, 0.813]	n/a	F2c
2.5×10^{-3}	1.174	0.7970	[0.797, 0.797]	n/a	F2c
2.6×10^{-3}	1.268	0.7745	[0.770, 0.780]	0.0520	F2c
2.7×10^{-3}	1.371	0.7575	[0.753, 0.763]	0.0485	F2c
2.8×10^{-3}	1.486	0.7380	[0.734, 0.743]	0.0410	F2b
2.85×10^{-3}	1.549	0.7265	[0.723, 0.730]	0.0360	F2b
2.9×10^{-3}	1.617	0.7130	[0.710, 0.716]	0.0290	F2b

The recovery was necessary because the solver clips the logged circulation drift to $\max(\text{signed}, 0)$ at `boussinesq/dedalus_axisym.py:653`, which is identically zero on viscous data; the correction ledger (Section 8) states the consequence for prior gating claims. The restart re-referencing defect in the logged Casimir stream was repaired by the anchored splice of gate G-F6 (a_2 : 2.69782 $\rightarrow$ 8.80469 on T2048_1e-7, 6.18292 $\rightarrow$ 8.75404 on T2048_1e-8, unrestarted runs unmoved), and the spliced values were then verified against full snapshot recomputation to 0.216% and 0.099% of signal, against a 1% bar (G-F6b, `out/flux/flux_ladder.json`). The one grid change along the campaign ladder is controlled, not acknowledged: the full 2×2 grid cell at $\nu = 10^{-6}$ spreads by 7.2×10^{-5} relative, and the same- ν cross-grid cell differs by 2.7×10^{-5} , while one decade of ν moves a_2 by 6.2×10^{-3} (`out/flux/FLUX_RESULT.md` C4).

5.4 The ladder at fixed time, and the verdict under the frozen rules

Fixed t is the derivation-forced primary convention, because the expansion behind (12) is at fixed time. At fixed time the G-F5 floor gate reads on the deviation $|a_2 - A|\nu$, not on the drift, and that reading is what removes the deep rungs: T2048_1e-8 reaches deviation SNR 4.0–6.7 against a bar of 20 and never clears at any reported time, and no rung at $\nu \leq 10^{-6}$ clears on any grid against a well-constrained reference (H1024_1e-6 1.2–2.0, D2048_1e-6 0.33–5.7, H1024_1e-7 2.0–3.1; `out/flux/FLUX_RESULT.md` §6.2). The spec’s designated primary ladder, T2048, therefore never fields three clearing rungs at any reported t : VOID-2 fires on it and no α exists there at matched time. T2048_1e-7 does clear its own gate reading, but only against a reference level A that is itself an unconstrained three-point extrapolation, and I do not lean on it (`out/flux/FLUX_RESULT.md` §4.2). The deepest rung that ever clears the deviation bar anywhere is $\nu = 3 \times 10^{-6}$, on D2048 at $t \geq 2.85 \times 10^{-3}$.

The one matched-time ladder that meets every structural conjunct is H1024 (seven rungs, five decades, single grid, zero restarts, so no splice enters it anywhere). Table 6 gives its $\alpha_2(t)$ with the profile interval $\{\alpha : \text{RSS}(\alpha) \leq 2\text{RSS}_{\min}\}$ and the drop-one jackknife spread. All collapse times $s = -\log(1 - t/t^*)$ use the free-fit $t^* = 3.6185 \times 10^{-3}$ (gate G-F10; the nominal 0.0030 is not the singularity time, Section 8).

Under the frozen rules, applied once at the end: F2b (not linear response) is met at $t \geq 2.8 \times 10^{-3}$, with $\alpha_2 = 0.7380/0.7265/0.7130$ and jackknife 0.041/0.036/0.029 against a bar

of 0.15, every entering rung clearing G-F5, over $\nu \in [10^{-5}, 10^{-3}]$. The sub-label square-root is not earned: $|\alpha - 0.5| = 0.213$ at the deepest quoted time, outside the 0.15 band. F2a is unavailable at every depth for a reason that is disclosed rather than argued away: F2a carries the conjunct that all gates pass, and the must-fail control G-F4a missed its frozen bar (inviscid enstrophy drift 0.342 against a required 1.0), while still discharging its purpose (9.3 orders of separation between the non-conserved enstrophy and the conserved Casimir on the same run, same quadrature). The bar was not moved; my registered prediction P2 for that control is scored falsified; and the asymmetry (a G-F4a failure blocks F2a but not F2b) is a property of my frozen text that did no work here, since α never entered $[0.85, 1.15]$ at a quotable depth (`out/flux/FLUX_RESULT.md` §1.1, §4.5, §8). Three controls protect the trend itself: refitting on a fixed rung set reproduces the monotone decline (H1024 fixed set $1.1165 \rightarrow 0.7130$; the independent D2048 fixed set gives the same shape, $1.3500 \rightarrow 0.6935$, over a disjoint decade of ν); removing floor subtraction moves α by at most 0.074 at the shallowest point and 0.0015 at depth; and swapping the floor realization moves it by at most 0.0005 (`out/flux/FLUX_RESULT.md` C1–C3). The pre-registered suspect rung T2048_1e-9 behaves exactly as predicted (P9): adding it takes the jackknife to 2.42 and breaks monotonicity at every depth, so it is excluded as floor-contaminated (`out/flux/FLUX_RESULT.md` C5).

5.5 The retracted crossing

The lane’s result document claims that four independent routes put the depth at which α_2 crosses 0.75 in a 3% window (`out/flux/FLUX_RESULT.md` §4.2). I retract that sentence; the lane’s own ladder file refutes the independence. Of the four routes in `out/flux/flux_ladder.json` (`alpha_crosses_0.75`): the gated H1024 route and its fixed-rung control return the bitwise-identical crossing $x_{\text{cross}} = 0.0027384615384615416$, because at the bracketing times the gated rung set equals the fixed set and the control is the same fit, not an independent route; the pooled route’s bracket endpoints are a VOID-3 fit and an F2c fit with every intermediate time NO-FIT, so its crossing rests on fits that are not quotable; and the D2048 fixed-set route contains a rung failing G-F5 with an endpoint jackknife of 0.3245, above the 0.25 VOID bar. Four routes were one. What survives is a single-route measurement: the gated H1024 ladder crosses 0.75 at $t = 2.7385 \times 10^{-3}$, $s = 1.414$, and the D2048 fixed-set control, on a different grid over a different decade of ν , reproduces the monotone decline and brackets its own crossing at $t = 2.660 \times 10^{-3}$, $s = 1.328$. I therefore quote the breakdown location as $s = 1.33$ to 1.41, one quotable route plus one consistency control, and not as four independent measurements. The depth-localised statement itself stands on Table 6: α_2 is theorem-backed near 1 at shallow depth (the inviscid solution is analytic in ν on a smooth interval, so the shallow rows validate the instrument and are not physics confirmations) and falls monotonically to 0.7130 at the deepest time clearing G-F5.

5.6 The leading coefficient is confirmed; what breaks is the correction

Prediction P-i is the sharpest positive statement in this lane, and it must not be skipped: the explicit inviscid quadrature (12), evaluated on the inviscid run’s own stored fields with nothing subtracted, reproduces the measured $\nu \rightarrow 0$ limit A of every matched-time fit to between 4.1×10^{-5} and 1.9×10^{-4} relative, on two grids, over the whole window (`out/flux/judge/fx2_pi.json`; `out/flux/FLUX_RESULT.md` C7). Set against the statistic it replaces, where an injected pure dilation was recovered at -0.52 to $+0.39$ instead of 1.00, this is the instrument working at the level the theorem promises. Linear response is therefore not wrong here: its leading coefficient is confirmed to one part in 6000 to 25000, and F2a’s monotonicity conjunct, $|a_2(\nu) - a_2^0|$ decreasing as ν falls, holds at every quoted depth on both ladders. What the α measurement detects is the exponent of the correction: $a_2(\nu) - a_2^0$ should be $O(\nu)$ and is

measured $O(\nu^{0.71})$ at the deepest times reached. The verdict is a statement about the second term of the expansion, never to be quoted as “linear response fails” without that qualifier.

One correction to my own spec was found while building (12). The closed wall form in (11) rests on $\Gamma_r|_{r=1} = 2\Gamma/r$, which the spec asserted as exact; that condition belongs to the viscous problem and degrades to nothing as $\nu \rightarrow 0$: the measured residual at $r = 1$ is 3.5×10^{-8} at $\nu = 10^{-5}$, 8.9×10^{-8} at 10^{-6} , 1.4×10^{-5} at 10^{-7} , 5.1×10^{-3} at 10^{-8} , and 5.6×10^1 on the inviscid run, which has no diffusion and hence no such condition. The practical consequence runs opposite to the obvious guess: the prediction operator must keep the closed form, because that is the boundary condition the viscous problem actually imposes; substituting the inviscid field’s own measured Γ_r degrades the agreement with the measured limit from 1.6×10^{-4} to 5.9×10^{-2} , a factor of 360. Both variants are computed and stored (`out/flux/judge/fx2_pi.json`; `out/flux/FLUX_RESULT.md §5`).

5.7 The conventions fork

Matched depth $\Omega = \sup |\omega_1|$ is the campaign convention and is computed as the declared second reading. It returns a different world: every single-grid matched-depth fit lands in a narrow band around $\alpha_2 \approx 0.4$ with the tightest jackknives in the lane (D2048: 0.4375 at $\Omega = 900$ falling to 0.3980 at Ω^* , jackknife 0.000–0.014; H1024: 0.4035 to 0.3195; pinned-model comparison $\text{RSS}(\alpha=1)/\text{RSS}(\alpha=0.5) = 26$ to 58 on D2048; `out/flux/FLUX_RESULT.md §2.4, §4.3`). By the letter of F2b these rows are labelled square-root, and they are reported and not promoted, because the spec pre-registered exactly this artifact of the convention: matching on depth compares runs at times offset by the viscous schedule shift, independently measured as $ds \sim \nu^{0.498}$, and a regular field response on a $\nu^{1/2}$ -shifted clock presents as $\alpha \approx 0.5$ on the depth convention (`FLUX_SPEC.md §3`). The two conventions straddle the 0.75 decision bar in opposite directions at shallow depth (0.9780 versus 0.4035) and differ by 0.394 at depth, so VOID-5 fires and the frozen outcome class is V-4: the measurement is sound, no single scalar α may be quoted across conventions, and the fork itself is the result (`out/flux/FLUX_RESULT.md §4.4`). I also record, for whoever writes the next spec, that rule F2e’s named orientation appears inverted relative to the mechanism it targets; the rule was applied as frozen and not rewritten to fit the data.

Scope is applied before any mechanism language (rule F2f). At the times where α is quoted, the free-slip wall share of the budget is $W_2 = 8.1 \times 10^{-3}$ against the 0.5 wall trip, and the fixed lab-frame corner-box share is $\text{box}_2(0.02) = 0.078$ to 0.131, above the 0.05 swirl-sheet floor, so neither clause strikes the result; at the shallowest reported depth ($\Omega = 900$, $\text{box}_2 = 0.0416$) the corner clause does fire and the corner question is declared unanswered by this instrument there (`out/flux/FLUX_RESULT.md`, F2f header). Two honest demotions follow. First, my spec’s sealed claim that a third of the Casimir dissipation sits in the fixed 2% corner box is not reproduced: ten box variants were probed and none returns 0.325–0.340; the supported share at the quoted depths is 8–16% (exposure R2, `out/flux/FLUX_RESULT.md §7`). Second, the p -ladder that was to localise the mechanism cannot be certified on this machine (C_p dealias doubling for $p = 8$ needs 35–89 GB against 17.2 GB physical; gate G-F1b is VOID for $p = 8$ everywhere and $p = 4$ is uncertifiable on the T grid), so no α_4 or α_8 is quoted; what the certified D-grid pass does show is structural and runs against my registered reasoning in P10: higher p is less corner-localised and more wall-weighted ($\text{box}_p(0.02) = 0.156/0.061/0.013$ and $W_p = 0.008/0.025/0.061$ for $p = 2/4/8$), so $p = 2$ is the most corner-weighted member of the family and the right one to quote (`out/flux/judge/fx2_p48.json`; `out/flux/FLUX_RESULT.md C6`).

5.8 The $\nu^{0.51}$ circulation depletion is a wall law

Recovering the signed circulation drift from the raw streams (the logged value is the clipped zero of Section 8) yields the cleanest power law in the campaign, and its geometry disqualifies it as a corner statement. The signed drift of $\sup|\Gamma|$ is negative and monotone at every ν on every grid, as the maximum principle requires, and scales as $\nu^{0.49-0.53}$ across the archive (FLUX_SPEC.md §0.2): adjacent-rung exponents on the certified grids are 0.532, 0.512, 0.500 from $\nu = 10^{-5}$ down to 10^{-8} (out/flux/fx_supgamma.json, nu_scaling), the excess-over-floor exponents at depths $\Omega = 800$ to 1400 sit at 0.508–0.517 (out/flux/critic/supgamma_ladder.json), and the rungs span $\nu = 10^{-9}$ to 3×10^{-5} , 4.5 decades, across the three grid families H1024, D2048, T2048 (out/flux/f0_supgamma.json; out/flux/f0_signed_drift.json; the 4.5-decade statement is also the sealed summary in F3_SPEC.md §0.1). Grid effects are far below the law: at $\nu = 10^{-6}$ the two certified grids agree on the signed drift to 2.0×10^{-4} relative while one decade of ν moves it by a factor 3.24 (out/flux/f0_supgamma.json). But $\arg \max |\Gamma|$ is pinned at $(r, z) = (1.000000, 0.041667)$, the wall at a quarter period, at every snapshot of every one of the five stored series, and it never moves; that is 2.08 times the frozen corner ball $r_c = 0.02$ away from the corner, while $\arg \max |\omega_1|$ sits at $r = 0.9993$ – 0.9999 (FLUX_SPEC.md §0.3; out/flux/FLUX_RESULT.md §6.1). The free-slip condition $\partial_r u_1 = 0$ at $r = 1$ forces $\Gamma_r|_{\text{wall}} = 2\Gamma/r \neq 0$, so Γ genuinely owns a wall layer there. The $\nu^{0.51}$ depletion law is a wall-layer measurement, the second of the three free-slip wall laws assembled in Section 7, and it must never be quoted as a corner result.

5.9 Reach, and the deliverable debts

Reach statement, in full. The F2 verdict (F2b, depth-localised, crossing at $s = 1.33$ to 1.41) is decidable at fixed time for $\nu \in [10^{-5}, 10^{-3}]$, matched-time convention, on the H1024 ladder with the D2048 fixed-set control; the deepest rung ever clearing the G-F5 deviation bar is 3×10^{-6} (D2048, $t \geq 2.85 \times 10^{-3}$); at $\nu \leq 10^{-6}$ the archive cannot decide, because the deviation from the linear-response limit falls to the archive’s own inviscid floor (1.80×10^{-10} on the D and T grids, 6.92×10^{-10} on H1024), which is a property of the inviscid runs and not of the estimator. Calling the quoted α a statement about the vanishing viscous remainder would be an extrapolation of two to three decades and is not made; in the Kato sense of the vanishing-viscosity limit [12], the limit itself is probed here only through the confirmed leading coefficient, not through the deep- ν correction. This is pre-registered outcome V-2 in its exact form: the instrument is validated and the archive is the ceiling. On the certified substrate of Section 2 [14], that ceiling is $\sup |\omega_1| \approx 1800$ of an unbounded collapse.

Two deliverable debts of this lane are disclosed here and were repaired in the F3 lane (Section 6; ledger item 4, Section 8): the machine-readable verdict was not emitted (out/flux/flux_ladder.json carries empty verdict, predictions, and exposures keys, a disclosed debt; the ladder, gate, control, and crossing blocks are populated and are the blocks every number in this section was verified against), and the lane’s log file is a byte-copy of its result document (out/flux/flux_ladder.log is byte-identical to out/flux/FLUX_RESULT.md). Neither debt touches a number; both are the kind of defect this campaign’s own gate taxonomy exists to catch, and they are recorded rather than repaired in place because the flux sandbox is sealed.

6 The ω_1 channel: F3 and the joint license

F3 measures the viscous enstrophy flux of the ω_1 channel as an integrated functional over $\nu \in [10^{-8}, 10^{-5}]$ within the stored collapse window; it does not resolve the pointwise T2 question, which remains open. That sentence was frozen in the pre-registration (ext_inertia_t2/F3_SPEC.md, section 5.6) before any verdict number existed, and by that spec’s own rule

it opens every document that quotes the numbers below. F2 (Section 5) confirmed the leading linear-response coefficient of the swirl Casimir, but every joule of its closed budget was swirl dissipation: the ω_1 channel, the one the frozen mechanism question T2 actually asked about, was untouched. F3 executes the same zero-by-theorem construction one channel over, one subtraction deeper. All F3 artifact paths in this section are relative to `ext_inertia_t2/out/flux3/` unless another tree is named; every headline number is re-derivable from `f3_ladder.json` alone.

6.1 The observable, and why its null is a theorem

Enstrophy built from ω_1 is not conserved inviscidly, but its inviscid production is an exact functional of the stored fields. With the solver’s own integration measure (weight r),

$$\begin{aligned} Z(t) &:= \int_{r_0}^1 \int_0^L \omega_1^2 r \, dz \, dr, & P(t) &:= 2 \int \omega_1 \partial_z (u_1^2) r \, dz \, dr, \\ R(t) &:= [Z(t) - Z(t_0)] - \int_{t_0}^t P \, dt. \end{aligned} \tag{13}$$

The budget is derived term by term from the equations the solver actually integrates (`boussinesq/dedalus_axisym.py`: inviscid ω_1 equation at line 211; the viscous system adds $\nu(\partial_{zz} + \partial_{rr} + \frac{3}{r}\partial_r)\omega_1$ and lift tau terms at lines 243 to 246, with boundary conditions $\omega_1 = 0$ at both walls at lines 249 and 250; $u_r = -r \partial_z \psi_1$, $u_z = 2\psi_1 + r \partial_r \psi_1$ at lines 189 and 190). Two exact statements carry the design. First, advection contributes exactly zero to dZ/dt in this measure: the meridional velocity satisfies $\frac{1}{r}\partial_r(ru_r) + \partial_z u_z = 0$ pointwise for any ψ_1 , so the advective term is a pure divergence, and its boundary flux vanishes because $\psi_1 = 0$ on both walls forces $u_r = 0$ there. Hence $dZ/dt = P$ identically on the inviscid system and $R = 0$ by theorem; the measured inviscid R is the instrument floor, nothing else. Second, on the viscous system one integration by parts in the r measure gives

$$2\nu \int \omega_1 \Delta_5 \omega_1 r \, dz \, dr = \nu (V_{\text{bulk}} + V_{\text{wall}}), \quad V_{\text{bulk}} := -2 \int ((\partial_r \omega_1)^2 + (\partial_z \omega_1)^2) r \, dz \, dr,$$

with $V_{\text{wall}} = 2[\oint (r \omega_1 \partial_r \omega_1 + \omega_1^2) dz]_{r_0}^1$. Under the viscous boundary conditions $\omega_1 = 0$ at both walls every term of V_{wall} vanishes exactly: the boundary flux is zero by proof, not by neglect. The measured check, per viscous run: $\max |V_{\text{wall}}|/|V_{\text{bulk}}|$ is 8.9×10^{-13} at $\nu = 10^{-5}$ and rises to 2.3×10^{-8} at 10^{-8} , with the wall $|\omega_1|$ at 6.6×10^{-8} rising to 6.8×10^{-4} , against interior values of order 1618 (`derive_floor.json`, `Vwall_share_max` and `wall_w_hi_max`; the sealed spec’s blanket “ 8.9×10^{-13} on viscous runs” in `F3_SPEC.md` section 0.4 is the 10^{-5} rung’s value, and I state the full range instead). So on viscous runs $R = \nu \int (V_{\text{bulk}} + V_{\text{wall}}) dt$ exactly, up to tau projections and discretization bounded by the closure gate, and R is sign-definite negative wherever significant (gate G-S: zero violations on the ladder, both schemes, `f3_gates.json`).

This construction is the repair Section 4 demanded. R contains no coordinate frame, no fitted scale, no lattice, no reference field: the collapse frame’s entire motion is generated by the inviscid flow, and every frame motion produces exactly zero R . There is nothing to deflate. The measured floor traverses a window in which $\sup |\omega_1|$ runs from 0 to 1758 while the trapezoid floor holds at 1.20×10^{-6} of Z (`derive_floor.json`).

Two independent derivation passes disagreed on the radial weight of Z : the solver measure r above, versus r^3 , the ω^θ component of the physical enstrophy. I froze the r measure as primary before the ladder ran, for three reasons recorded in the spec (section 0.3), none aesthetic: the entire measured feasibility record (floors, closures, the decidable rung set) was produced in the r measure; the r measure has strictly fewer artifact channels (no streamfunction recovery per snapshot, no stretching term, advection identically zero, and a term that does not exist

cannot leak); and both weights are functionals of ω_1 alone, so the channel scope holds either way, the fork being a bounded radial factor $r^2 \in [0.16, 1]$ on this annulus, not a channel change. The r^3 reading was registered as a non-verdict secondary and was not computed this phase: it requires per-snapshot ψ_1 recovery across roughly 800 stored writes, and its absence removes a cross-check, not a verdict input (`f3_ladder.json`, `secondary_r3_reading`). The disagreement between the two derivation passes is disclosed here rather than silently resolved.

6.2 Scheme discipline

The scheme history, told straight. The derive phase registered trapezoid time quadrature, and its closure prediction failed at the two deep rungs: 2.07×10^{-2} and 6.57×10^{-2} against the 10^{-2} bar at $\nu = 10^{-7}$ and 10^{-8} , with the deficit shown to be exactly the trapezoid floor share ($-0.7630 = -0.7791 + 0.0161$ at 10^{-7} , `derive_floor.json`). Simpson quadrature on the same stored nodes drops the inviscid floor by a measured factor 8349 (from $+1.62 \times 10^{-2}$ absolute, 1.20×10^{-6} of Z , to -1.94×10^{-6} absolute, -1.4×10^{-10} of Z). Adopting Simpson after seeing the trapezoid failure is a post-registration scheme change, open to a scheme-shopping objection, so the spec froze the resolution before any verdict number: Simpson is primary; trapezoid is computed everywhere as the mandatory scheme-and-cadence control; both are reported at every number; if the two schemes' verdict classes disagree after each is quoted against its own floor, the fork itself is the result and the weaker verdict is reported (VOID-5); no further scheme change is permitted for any reason (VOID-7). The registered trapezoid failure is preserved unrewritten as F3-P3: the live demonstration that closure can fail on this data, and the reason Simpson is primary.

The floor is a measurement, not self-cancelling algebra, by demonstrated can-fail: stride-2 cadence inflates the inviscid floor by a median factor 3.91 (trapezoid, an $O(\Delta t^2)$ signature) and 21.2 (Simpson, $O(\Delta t^4)$) against the must-move bar of 2 (gate G-K); a 10^{-3} perturbation of P moves R by exactly $-10^{-3} \int P dt$ to 9.0×10^{-14} relative arithmetic agreement, and the scheme swap moves the floor by the factor 8349 already quoted (gate G-A, `f3_gates.json`). The Simpson closure $|R - \nu \int (V_{\text{bulk}} + V_{\text{wall}}) dt| / |\nu \int (V_{\text{bulk}} + V_{\text{wall}}) dt|$ at window end is $5.8 \times 10^{-7} / 1.1 \times 10^{-6} / 8.6 \times 10^{-6} / 1.45 \times 10^{-4}$ for $\nu = 10^{-5} / 10^{-6} / 10^{-7} / 10^{-8}$, two to four orders under the 10^{-2} bar (gate G-C). Floors are depth-local always: each excess is quoted against its own window's floor, both realizations (inviscid replicate and the rung's own full-budget residual), factor-3 agreement required, no global floor anywhere.

6.3 The ladder

The primary statistic is $\rho(\nu, W) := |R(\nu, W)| / (2\nu \int_W D_{\text{inv}} dt)$ with $D_{\text{inv}} := \int ((\partial_r \omega_1)^2 + (\partial_z \omega_1)^2) r dz dr$ evaluated on the inviscid fields: the denominator is itself the quasi-static linear-response prediction of R , so LR predicts $\rho \rightarrow 1$. Windows are depth-matched in $\Omega = \sup |\omega_1|$ through the certified cubic-in-log Ω interpolator (`out/bchan_depth/`, leave-one-out error 0.006% to 0.5% of signal); every collapse time uses the fitted $t_* = 3.6185 \times 10^{-3}$, never the nominal 0.0030. Because the free-slip inviscid wall carries the supremum of $|\omega_1|$ itself (wall value 1758.530 against domain supremum 1758.475, `derive_floor.json`) while every viscous run clamps $\omega_1 = 0$ there, a reconciling wall layer of width $O(\sqrt{\nu t})$ is guaranteed viscous physics in this budget [12], and scoring it as a bulk-channel violation is the named misread the spec closed in advance. The frozen split: $s_{\text{wall}}(\nu, W)$ is the share of $\nu \int_W (-V_{\text{bulk}}) dt$ inside wall strips of frozen width $w_s = 5\sqrt{\nu t_*}$ (multiplier forks {3, 8} reported); ρ_{bulk} excises the same strips from both numerator and denominator. Both statistics are verdict-bearing, each in its own scope: ρ_{tot} adjudicates the global ω_1 -channel flux, and ρ_{bulk} carries the T2-relevant sentence whenever $s_{\text{wall}} \geq 0.5$ at the deepest decidable rung. The split is well-posed on this data: s_{wall} moves by less than 0.02 across the {3, 5, 8} forks at every quoted rung and window, and the scope assignment is identical under every fork (check F-4, `F3_RESULT.md` section 4).

Table 7: The F3 ladder at the deepest common window $[1200, \Omega_*)$, $\Omega_* = 1618.3879571332704$, both schemes, wall strip multiplier 5. S = Simpson (primary), T = trapezoid (control). σ_{rel} is the deterministic relative budget on ρ_{tot} (trapezoid); σ_{abs} on ρ_{bulk} (trapezoid) adds the $\{3, 8\}$ strip-fork half-spread. Under Simpson this window is floor-realized only on the 10^{-7} rung (see Section 6.4); the Simpson verdict fit runs at $[200, 400)$. Source: `ext_inertia_t2/out/flux3/f3_ladder.json`.

ν	run	ρ_{tot}^S	ρ_{tot}^T	σ_{rel}	ρ_{bulk}^S	$\rho_{\text{bulk}}^T \pm \sigma_{\text{abs}}$	s_{wall}
10^{-5}	D2048	3.911	3.909	0.0013	1.243	1.242 ± 0.005	0.704
10^{-6}	T2048	9.021	9.012	0.073	1.095	1.094 ± 0.087	0.882
10^{-7}	T2048	25.35	25.26	0.082	1.039	1.036 ± 0.11	0.959
10^{-8}	T2048	76.95	76.04	0.119	1.029	1.017 ± 0.2	0.987

Table 7 is the result. The global statistic ρ_{tot} rises as ν falls, 3.911 to 76.95, and the wall share rises with it, 0.704 to 0.987: the excess is a wall number. The power-law fit $\rho_{\text{tot}} = c\nu^{\beta-1}$ returns $\beta = 0.521$, leave-one-rung-out jackknife CI $[0.511, 0.532]$ (Simpson, fit window $[200, 400)$; trapezoid at $[1200, \Omega_*)$ gives 0.569, CI $[0.537, 0.595]$), echoing the derive-phase $\nu^{0.538}$ and F2’s swirl wall law $\nu^{0.51}$: the campaign’s third free-slip pedestal, treated together in Section 7. The verdict on ρ_{tot} is F3b on both schemes, and by the frozen scope rule it is a wall-layer result in the verdict string itself: it does not adjudicate T2’s bulk question. The matched-time cross-check tells the same story at every fixed τ on the fitted- t_* grid, and the inviscid wall functional B_{inv} integrates to -1011 at Ω_* (trapezoid, `f3_ladder.json`): the mismatch layer is real inviscid-wall physics, not an artifact of the split.

The scope-carrying statistic is ρ_{bulk} , and it converges to 1 from above: $|\rho_{\text{bulk}} - 1|$ falls monotonically from 0.242 to 0.017 across the trapezoid ladder. The frozen one-correction fit $\rho(\nu) = 1 + b\nu^q$, q free in $[0.4, 1.0]$, gives intercept $\rho(\nu \rightarrow 0) = 1.0035$ (leave-one-rung-out jackknife spread 9.9×10^{-3} , trapezoid, four rungs) and 0.9999 (jackknife spread 9.8×10^{-7} , Simpson, three rungs): quasi-static linear response confirmed in the leading coefficient, F3a on both schemes, no scheme fork. Anomalous-flux forms in the sense of a non-vanishing dissipative remainder [9] are excluded in the bulk over the measured range: the competing family $c\nu^{\beta-1}$ fitted to ρ_{bulk} returns β within 0.03 of 1 on both schemes. Two disclosures belong to this verdict, not to a footnote. First, model separability: VOID-4 (the inseparability void) did not fire, but on the primary scheme the two families differ by more than 1σ at exactly one rung of three, $\nu = 10^{-5}$, with margin $1.145\times$ (model separation 1.47×10^{-4} against $\sigma = 1.28 \times 10^{-4}$); the trapezoid fit separates at the same single rung at $7.19\times$ (`f3_ladder.json`, `rules`). The family discrimination therefore rests on the shallowest rung, where the budget is tightest. Second, the corner: the largest fixed corner-box share of the dissipation is 0.412 ($b = 0.10$, $\nu = 10^{-8}$, $[1200, \Omega_*)$) against the 0.5 gate, so corner language is not licensed anywhere in this lane, the gated reduced-solve route was correctly not evaluated, and every verdict sentence says channel, never corner. The largest single error term on every T-grid rung is the grid term 0.068, the measured floor-transfer spread standing in for the uncomputable same- ν cross-grid cell (exposure E-2; `snap_D2048_1e-6` does not exist). Figure 2 shows $\rho_{\text{bulk}}(\nu)$ with its depth-local floors, both schemes.

6.4 Reach: the honest ruling on the 10^{-8} rung

The verdict lines carry their reach. On the primary scheme: F3a on ρ_{bulk} over $\nu \in [10^{-7}, 10^{-5}]$ at depth window $[200, 400)$, Ω -matched through the certified interpolator on fitted t_* ; s_{wall} at the deepest Simpson-decidable rung (10^{-7}) is 0.959. On the control: F3a on ρ_{bulk} over $\nu \in [10^{-8}, 10^{-5}]$ at $[1200, \Omega_*)$; $s_{\text{wall}} = 0.987$ at 10^{-8} (`f3_ladder.json`, `verdict`).

The 10^{-8} rung needs its mechanics stated, because the primary scheme does not decide it.

■ Simpson (primary) □ open marker: floor-unrealized, not quotable under that scheme (gate G-F)
 ● trapezoid (control)




Figure 2: $\rho_{\text{bulk}}(\nu)$ at the deepest common window $[1200, \Omega_*)$, wall-strip multiplier 5, both schemes, with the frozen F3a band $[0.8, 1.25]$ shaded. Error bars are the deterministic instrument budget (σ_{rel} times the value) plus the $\{3, 8\}$ strip-fork half-spread. Open markers are floor-unrealized under that scheme (gate G-F) and not quotable there; at this window the Simpson column is quotable only on the 10^{-7} rung, and the Simpson verdict fit runs at $[200, 400)$ (Section 6.4). Bottom panel: the depth-local floor share $|\text{floor}|/|R|$ per scheme. All values from `ext_inertia_t2/out/flux3/f3_ladder.json`.

Under Simpson, that rung’s only verdict window $[1200, \Omega_*)$ is floor-unrealized: the two floor realizations, both at roundoff scale, disagree by the factor 12.73 against the frozen factor-3 bar (gate G-F, `f3_gates.json`), and by the frozen quotability rule nothing in that window may be quoted under that scheme. The Simpson floor is so small that its cross-run realization ratio is noise on noise; the check degrades exactly where the scheme is best. Under trapezoid the same window is floor-realized (ratio 1.048) with windowed SNR 82.1 and per-interval median SNR 50.6 over 22 snapshot intervals, and the rung is decided there: $\rho_{\text{bulk}} = 1.017 \pm 0.2$, inside the F3a band, continuing the monotone approach to 1. The scheme classes agree wherever both quote, so VOID-5 is silent, and the ruling I adopt is: primary reach $[10^{-7}, 10^{-5}]$, the 10^{-8} rung decided by the trapezoid control in class agreement.

Two corrections attach here, stated in full in the corrections ledger (Section 8, items 5 and 7). First, the spec contains an internal reach contradiction, and I wrote it: section 1.7 declared the reach as $\nu \in [10^{-8}, 10^{-5}]$ and named $[1200, \Omega_*)$ the primary 10^{-8} verdict window, anticipating Simpson decidability there from the derive-phase SNR; but the same spec froze Simpson as primary (0.5) and the depth-local floor-realization rule (1.4), and on the measured data those clauses jointly remove the 10^{-8} rung from the primary scheme. The clauses cannot all be satisfied, and I resolve the contradiction the conservative way, as above. Consequently the sealed record’s campaign outcome label “A-1 on ρ_{bulk} ” (F3_RESULT.md section 0; `f3_ladder.json`, `campaign_outcome`), whose defining clause requires both 10^{-7} and 10^{-8} inside the verdict, overstates the primary reach: I retract that label. The measurement stands; the label does not. Second, the phase left three latent code defects, found in post-run audit, none outcome-changing, all disclosed: the verdict assembler tests the F3a criteria before the VOID-4 inseparability check, an ordering that never executed because VOID-4 fired nowhere; the G-D convention gate’s pass boolean was keyed to the median sup-grid offset (4.0×10^{-5}

on the 10^{-8} rung) while the same record carries a maximum of 1.06×10^{-2} at one of 42 depth nodes against the gate’s stated over- 10^{-3} can-fail clause; and the G-N(a) must-fail metric is degenerate on this initial condition ($Z(t_0) = 0$ exactly, so the relative-drift reading saturates at 1.00 and cannot grade by degrees, though the null still fails as the theorem demands).

6.5 Verification record

The battery ran before any verdict quantity, and everything below is on disk. All eleven gates pass (`f3_gates.json`, `all_gates_pass`). Of eleven predictions registered in the frozen spec before the ladder ran, ten were confirmed and one was uncomputable on a declared branch (P-9, the missing cross-grid run); P-8, the only prediction registered with low confidence, is the one that landed F3a. The ladder harness reproduced all eight sealed derive-phase anchors from scratch through a new code path, worst relative miss 1.88×10^{-6} against the 10^{-4} bar (gate G-R). The wall pass then recomputed every dissipation total a third way, by Chebyshev-antiderivative quadrature on a resolution-doubled grid against the sealed spectral values: self-check 3.0×10^{-15} on a known polynomial, inviscid-run agreement to 7.3×10^{-14} , viscous discrepancies tracking each run’s spectral tail, and only same-convention shares entering any verdict quantity (`F3_RESULT.md` section 1). Every anchor quantity in this section is therefore the survivor of three independent computations, the derive pass (`derive_floor.json`), the harness field pass, and the wall pass, in numerical agreement. Data hygiene was enforced in memory only: unique write counts 202/198/197/212/209 reproduced exactly, the 17+1 duplicate writes de-duplicated with a measured splice term of at most 2.18×10^{-6} of the anchor, and the duplicates found not bit-exact (maximum payload difference 2.9×10^{-13} , correcting an earlier campaign note; disk untouched). The read-only campaign tree, 1590 files, 63.93 GB, was md5-hashed at open and close of both the gate battery and the verdict ladder: zero changed, zero deleted, zero appeared, across all four manifest passes (`f3_manifest_diff.json`, `f3_verdict_manifest_diff.json`). Three adversarial review lenses were then run against the sealed record under the campaign’s default-refute protocol (Section 10); none produced a refutation.

6.6 The joint license

F3’s clean F3a on the scope-carrying statistic, joined to F2’s confirmed leading coefficient (4.1×10^{-5} to 1.9×10^{-4} relative, two grids, Section 5), licenses exactly the sentence pre-committed in the spec (5.5) before either verdict existed, quoted here with the spec’s paired-dash punctuation set as parentheses and its capitalized emphasis set in lowercase, no other change:

over ν in [decidable range] and the matched depth window, the vanishing viscous remainder acts as quasi-static linear response at the level of integrated quadratic functionals (the swirl Casimir C_2 in the r -weighted L^2 pairing and the ω_1 -channel enstrophy flux in the r -weighted H^1 -seminorm pairing) in the leading coefficient, with sub-leading corrections vanishing no slower than $\nu^{0.4}$.

The decidable ranges are those measured: $[10^{-5}, 10^{-3}]$ for the swirl channel and, for the ω_1 channel, primary $[10^{-7}, 10^{-5}]$ with the 10^{-8} rung decided by the trapezoid control in class agreement. The two lanes do not jointly license the unqualified statement “the remainder is quasi-static linear response.” Both confirmations are integrated statements on the certified corner-collapse profile [14], and both are silent on the second term of the expansion: F2 measured its swirl-channel correction exponent at 0.71, not the regular 1, and F3’s bulk correction fits $q = 0.425$ to 0.438 (trapezoid jackknife spread 0.088), sitting at the frozen band edge of 0.4: the sub-leading exponent is fitted, not resolved. The wall layers are removed by a frozen split, not explained; they are quantified in Section 7. Nothing here is pointwise: F3 cannot




Figure 3: The three free-slip wall laws on one log-log panel; each law is in its own units, the shared content is the exponent class straddling $\nu^{1/2}$ (dashed guide). Law 1: the de-pedestalled excess $E(\nu)$, the two detections at $\nu = 10^{-5}, 10^{-6}$, with the 10^{-7} and 10^{-8} zeros shown at their 2σ bound $E \leq 0.037$ (sealed viscous-anomaly ledger, Table 3). Law 2: the signed sup $|\Gamma|$ depletion at matched depth Ω_* on the certified grids (`ext_inertia_t2/out/flux/fx_supgamma.json`). Law 3: ρ_{tot} at $[1200, \Omega_*)$, Simpson, $\sim \nu^{\beta-1}$ with $\beta = 0.521$ (`ext_inertia_t2/out/flux3/f3_ladder.json`). Inset: the wall share $s_{\text{wall}}(\nu)$ of the viscous flux (Simpson, strip multiplier 5, same window, same source), rising from 0.704 to 0.987 against the 0.5 scope gate.

address $\omega_{1,\nu} - \omega_{1,\text{inv}} = \nu\delta + o(\nu)$ anywhere, least of all at the corner ring, nothing in sup-norm, nothing below $\nu = 10^{-8}$, and nothing deeper than the stored window ($\Omega \lesssim 1780$, $\tau \gtrsim 6.3 \times 10^{-4}$ of the fitted t_*). Structural blind spots survive even a clean F3a: rearrangements preserving $\int |\nabla \omega_1|^2 r$ are invisible to this functional, and time-local violations can cancel within a window, mitigated but not removed by the nine-window report. Any future quotation of the joint result must carry the two pairings and the ν window, or it is an overclaim. The pointwise T2 question remains open at every ν .

7 The wall layers

Three instruments, built weeks apart to ask three different questions, each returned a viscous magnitude scaling in the $\nu^{1/2}$ class. All three are the same object: the free-slip wall layer. I set the three laws side by side, state the mechanism by which the boundary condition manufactures exactly this class, and then state the corner result that the wall measurements license. Figure 3 shows the three laws on one log-log panel with the wall share $s_{\text{wall}}(\nu)$ as an inset.

7.1 The three laws

Law 1: the deformation pedestal. The registered deformation floor $d_\infty \approx 0.19$ decomposed as roughly 85% free-slip boundary pedestal; the sealed statement is on disk in the flux pre-registration (`ext_inertia_t2/FLUX_SPEC.md`, section 0.3: “the $d_\infty \sim 0.19$ episode that turned out $\sim 85\%$ boundary condition”), and the full estimator discipline and the depedestalled verdict $E = 0$ at $\nu = 10^{-7}$ and 10^{-8} are Section 3 of this paper. What vanishes outside the layer carries a measured exponent: the budget instrument returned $\nu^{0.4158 \pm 0.0070}$ on the collapse-frame window and $\nu^{0.4435 \pm 0.0102}$ under identical-domain wall excision, and the unclamped u_1 wall-strip channel returned $\sim \nu^{0.40}$ (`boussinesq/noslip/WALL_SWAP_SPEC.md`, which froze against those session measurements). Inside the layer the clamped field does not vanish at all: the ω_1 wall amplitude is ν -flat, ratio 1.009 ± 0.019 across $\nu = 10^{-3}..10^{-6}$ with fitted exponent -0.0003 (same record). Flat amplitude on a $\sqrt{\nu t}$ support is the pedestal.

Law 2: the swirl-channel circulation depletion. The second law is the depletion of the wall-crest circulation maximum, $\sup |\Gamma|$ with $\Gamma = r^2 u_1$, measured as the signed end-of-window drift recovered from stored fields (the in-run log clipped it to zero; Section 8, correction 1). The campaign’s own sidecars carry it on every run (`boussinesq/regate_*.json`, field `gamma_drift_signed`, final stored write, $t = 2.97 \times 10^{-3}$ to 3.00×10^{-3} per run):

ν	H1024	H1536	D2048	T2048
10^{-4}	1.462×10^{-3}			
3×10^{-5}	7.582×10^{-4}		7.817×10^{-4}	
10^{-5}	4.283×10^{-4}		4.397×10^{-4}	
3×10^{-6}	2.363×10^{-4}	2.355×10^{-4}	2.375×10^{-4}	
10^{-6}	1.348×10^{-4}	1.345×10^{-4}	1.337×10^{-4}	1.337×10^{-4}
10^{-7}	4.118×10^{-5}			4.259×10^{-5}
10^{-8}	1.172×10^{-5}	1.300×10^{-5}		1.360×10^{-5}
10^{-9}	(1.041×10^{-5})	3.524×10^{-6}		4.054×10^{-6}
inviscid floor	8.31×10^{-7}		$1.38 / 1.40 \times 10^{-6}$	3.56×10^{-7}

(H1536 column: the H1536mid runs at $\nu = 3 \times 10^{-6}$ and 10^{-6} , the H1536 runs at 10^{-8} and 10^{-9} ; D2048 floor: both inviscid realizations, agreeing to 1.7% at the end write.)

Adjacent-decade exponents along the deep ladder are 0.532, 0.512, 0.500 ($\nu = 10^{-5} \rightarrow 10^{-8}$, `ext_inertia_t2/out/flux/fx_supgamma.json`, `nu_scaling`). Endpoint fits over $\nu = 3 \times 10^{-5}$ to 10^{-9} , which is 4.5 decades, give 0.510 (D2048 to T2048), 0.507 (H1024 to T2048) and 0.521 (H1024 to H1536); ordinary least squares over the 19 viscous rungs gives slope 0.514 (computed from the sidecar values above; including the excluded point below it is 0.491). I exclude one point: H1024 at $\nu = 10^{-9}$ sits a factor 2.6 above the same- ν T2048 and H1536 values, which agree with each other to 15%; the 10^{-9} rung was pre-registered as suspect and confirmed floor-contaminated in the F2 lane (control C5, `ext_inertia_t2/out/flux/FLUX_RESULT.md`). The law is not a grid effect: at the $\nu = 10^{-6}$ cell, four runs on four grid configurations agree to 0.84% relative spread while one decade of ν moves the depletion by the factor 3.29 ($10^{-6} \rightarrow 10^{-5}$), a margin of better than 270:1 of law over grid. And it is a wall number: $\arg\max |\Gamma|$ is pinned at $(r, z) = (1.000000, 0.041667)$, the wall crest at $|z| = L/4$, at every write of all five snapshot series, $2.08\times$ the frozen corner ball $r_c = 0.02$ away from the corner (`FLUX_RESULT.md`, gate G-F8 and finding 6.1.7). The sealed spec states the attribution in its own words: the free-slip condition $\partial_r u_1 = 0$ at $r = 1$ makes $\Gamma_r|_{\text{wall}} = 2\Gamma/r \neq 0$, so Γ genuinely has a wall layer, and “the $\nu^{0.51}$ ladder is a wall-layer measurement, not a corner result” (`ext_inertia_t2/FLUX_SPEC.md`, section 0.3). Reach: the ladder spans $\nu \in [10^{-9}, 10^{-4}]$ at fixed end-of-window time; the quoted exponent rests on $[10^{-9}, 3 \times 10^{-5}]$ with the one named exclusion; floors 3.6×10^{-7} to 1.4×10^{-6} per family.

Law 3: the ω_1 -channel total flux. The production-subtracted enstrophy ladder of Section 6 lands F3b on the global statistic ρ_{tot} under both quadrature schemes: at the deepest common window $[1200, \Omega_*)$, $\rho_{\text{tot}} = 3.911/9.021/25.35/76.95$ (Simpson; trapezoid 3.909/9.012/25.26/76.04; relative σ 0.0013/0.073/0.082/0.119) for $\nu = 10^{-5}..10^{-8}$ (`ext_inertia_t2/out/flux3/f3_ladder.json`; `F3_RESULT.md`, section 3). The registered power-law family $\rho_{\text{tot}} = c\nu^{\beta-1}$ fits with $\beta = 0.521$, drop-one jackknife CI $[0.511, 0.532]$, on the primary scheme (Simpson, fit window $[200, 400)$, three rungs, $\nu \in [10^{-7}, 10^{-5}]$), and $\beta = 0.569$, CI $[0.537, 0.595]$, on the trapezoid control (fit window $[1200, \Omega_*)$, four rungs, $\nu \in [10^{-8}, 10^{-5}]$); both CIs exclude the linear-response band $[0.9, 1.1]$, the scheme fork does not fire (`f3_ladder.json`, `family_B`). The verdict scopes itself: the wall share of the flux, measured with strip width $w_s = M\sqrt{\nu t_*}$ at the frozen $M = 5$ (forks $M \in \{3, 8\}$ move it by less than 0.02 everywhere), is

$$s_{\text{wall}} = 0.704 / 0.882 / 0.959 / 0.987 \quad (\nu = 10^{-5}..10^{-8}, [1200, \Omega_*)),$$

rising as ν falls, so by the pre-registered scope rule $s_{\text{wall}} \geq 0.5$ the F3b is a wall-layer result and the bulk statistic ρ_{bulk} carries the channel verdict (F3a, Section 6). Matched-time readings corroborate: at $\tau = 1.4 \times 10^{-3}$, $\rho_{\text{tot}} = 7.0/19.9/60.7/189.7$ with $s_{\text{wall}} = 0.858/0.950/0.984/0.995$ (`F3_RESULT.md`, section 3). Reach: primary scheme $[10^{-7}, 10^{-5}]$, the 10^{-8} rung decided by the trapezoid control in class agreement; depth windows as quoted, Omega-matched via the certified interpolator on fitted $t_* = 3.6185 \times 10^{-3}$, never the nominal 0.0030.

Three laws, three observables, three norms, one exponent class: 0.40..0.57 across every instrument, straddling $\nu^{1/2}$.

7.2 Why free-slip manufactures the $\nu^{1/2}$ class

The free-slip conditions at $r = 1$ are $\psi_1 = 0$, $\partial_r u_1 = 0$, $\omega_1 = 0$ (`boussinesq/dedalus_axisym.py`; `boussinesq/noslip/WALL_SWAP_SPEC.md` records the triple). The certified inviscid profile puts its vorticity maximum on that wall: the inviscid ω_1 peak sits at wall distance $y = 2.51 \times 10^{-6}$ with amplitude $|\omega_1| = 1986$ (`WALL_SWAP_SPEC.md`). The Dirichlet clamp $\omega_1(r=1) = 0$ is imposed exactly at the wall row (a tau condition, satisfied to solver roundoff), so for every $\nu > 0$ the viscous solution must bridge from zero at the wall to the full inviscid amplitude across the diffusion width $\delta \sim \sqrt{\nu t}$. That is a full-amplitude deformation confined to a $\nu^{1/2}$ -thin strip. Every integrated quadratic functional then acquires a wall term of order $\delta \times (\text{amplitude})^2$, which is the $\nu^{1/2}$ class; and a sup-norm instrument pointed at the clamped field reads the depletion of the peak itself and returns $\nu^{1/2}$ as well (the 0.51 ladder). The vanishing-viscosity limit being decided inside a ν -thin wall strip is the classical Kato picture [12]; here the strip is manufactured by a free-slip clamp acting on a profile whose maximum sits on the wall, not by no-slip friction.

The mismatch is a property of the pair of problems, not of the solver. The viscous problem imposes the wall identity $\Gamma_r = 2\Gamma/r$ at $r = 1$ (equivalent to $\partial_r u_1 = 0$), and the measured residual $\max |\Gamma_r - 2\Gamma| / \max |2\Gamma|$ on the wall is 3.5×10^{-8} at $\nu = 10^{-5}$, 8.9×10^{-8} at 10^{-6} , 1.4×10^{-5} at 10^{-7} , 5.1×10^{-3} at 10^{-8} , and 5.61×10^1 on the inviscid run (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 5, third correction): the viscous runs follow their own wall condition ever more loosely as ν falls and the inviscid field does not satisfy it at all, so the layer must be rebuilt at every ν and never converges away. The layer is also real inviscid-wall physics rather than a viscous-solver artifact: the inviscid wall functional B_{inv} bounding the F3 layer term integrates to -1011 at Ω_* (trapezoid, `F3_RESULT.md`, section 3), nonzero and of the signature sign.

The wall-condition swap. I tested the clamp attribution directly by changing the one free variable, the boundary condition itself. The experiment was frozen before any no-slip

field existed (boussinesq/noslip/WALL_SWAP_SPEC.md, 2026-08-07): same solver copy, same IC, same grids, same stepper, same ν ladder, only the wall condition swapped from free-slip (ω_1 Dirichlet-clamped, u_1 Neumann-free) to no-slip (u_1 Dirichlet-clamped, ω_1 unclamped and slaved), driven at seven viscosities on each of two grids (128×384 and 256×768 , $\nu = 10^{-3}..10^{-6}$; boussinesq/noslip/logs/wallswap_driver.log). Thirteen run records landed (runs/NS*.json); the 256×768 member at $\nu = 10^{-6}$ went NaN in its resume and carries no record (boussinesq/noslip/logs/NS256_1e-6_resume.log). Under the frozen rules the formal verdict was W3, undecided, and the reason is a design error I made in the spec and record here: rule W2 required the newly unclamped ω_1 wall amplitude to vanish, but a no-slip wall started against a swirl IC with $u_1(1) = A$ manufactures a Rayleigh sheet whose ω_1 scales like $\nu^{-1/2}$, so “unclamped ω_1 should vanish” was never the right null (the confound was declared in the frozen spec and the sealed record names it a design error). The discriminating observation nonetheless held on both grids and both channels: the ν -flat wall amplitude moved with the Dirichlet clamp. Free-slip clamps ω_1 and ω_1 is the flat channel (ratio 1.009 ± 0.019 , exponent -0.0003) while unclamped u_1 vanishes (exponent 0.4049 , ratio 15.6); no-slip clamps u_1 and u_1 becomes the flat channel (the sealed campaign record of 2026-08-11, memory ledger). Persistence tracks the boundary condition, not the physics.

7.3 The reading rule, and the corner result

What the three laws imply for this geometry is a rule I now apply to every viscous magnitude in the campaign: check the wall share first. The archive provides the instruments: the strip share s_{wall} , the boundary-term share $W_p = |\text{WALL}_p|/|\text{BULK}_p|$, the fixed corner-box share box_b , and the level-set localisers. The rule has teeth in both directions. The Casimir functional C_2 passes it: $W_2 = 8.1 \times 10^{-3}$ at the quoted depths, sixty times below the 0.5 wall-instrument gate, so its verdicts in Section 5 are bulk statements (FLUX_RESULT.md, F2f and G-F8). The sup $|\Gamma|$ ladder fails it (argmax on the wall, Law 2), and ρ_{tot} fails it (s_{wall} to 0.987, Law 3); both are quoted in this paper only as wall numbers. The p -ladder gives the rule a direction: on the D grid, $\text{box}_p(0.02) = 0.156/0.061/0.013$ while $W_p = 0.008/0.025/0.061$ for $p = 2/4/8$ (FLUX_RESULT.md, C6), so weighting the field’s peak harder walks an observable toward the wall and away from the corner: sup-adjacent refinements make wall contamination worse, not better.

The corner non-license is the positive form of the same measurement. At the deepest decidable window $[1200, \Omega_*)$ the corner-box share of the ω_1 -channel viscous flux is, for $b = 0.02/0.05/0.10$: $0.135/0.260/0.377$ at $\nu = 10^{-5}$, $0.118/0.265/0.399$ at 10^{-6} , $0.111/0.267/0.408$ at 10^{-7} , and $0.109/0.268/0.412$ at 10^{-8} (ext_inertia_t2/out/flux3/f3_ladder.json; F3_RESULT.md, section 3). The largest corner share any box up to $b = 0.10$ ever reaches is 0.412, against the pre-registered 0.5 gate; the Route-2 corner-share gate fails at 0.408 (Simpson) and 0.412 (trapezoid) and the pushforward coefficient κ_{D1} is therefore not evaluated. Meanwhile 98.7% of the same flux at $\nu = 10^{-8}$ sits in a wall strip of width $5\sqrt{\nu t_*}$. So the licensed statement, and it is a result about where viscosity acts on this collapse rather than an apology about where it does not: within $\nu \in [10^{-8}, 10^{-5}]$ and the stored window, the viscous enstrophy flux is carried by the outer wall layer, its share rising toward unity as ν falls, and no fixed corner neighborhood up to $b = 0.10$ ever carries as much as half of it. Viscosity’s leading interaction with the certified corner collapse [14] is a boundary-layer interaction at the wall, and any corner-mechanism reading of any viscous magnitude on this archive must first survive the wall-share check. Reach of the non-license: the four F3 rungs, $\nu \in [10^{-8}, 10^{-5}]$, depth windows through $[1200, \Omega_*)$, Omega-matched on fitted $t_* = 3.6185 \times 10^{-3}$; the 10^{-8} row is trapezoid-decided as in Law 3.

8 Corrections ledger

This ledger is main text, not an appendix, because the corrections are part of the measurement. Each entry states what was wrong, where the defect lives on disk, what I did about it, and what it changes. Three of the seven (items 1, 4, and 7) are instruments reporting themselves; every one was found by asking what data enters a gate and whether that gate can fail. None changes a headline number of Sections 5 or 6; item 6 removes a headline claim of the earlier field-statistic arc, and item 5 removes one F2 deliverable and one F3 label.

Correction 1: the circulation gate is one-sided, and vacuous on viscous data

The solver logs a circulation-drift monitor `gamma_drift` on every record and gates trust on `gd < GAMMA_TRUST` (10^{-4}). The monitor is one-sided by deliberate design, documented at `boussinesq/dedalus_axisym.py:650--653`: on inviscid runs it gates on `|drift|`, while on viscous runs it is computed as `max(_gd_signed, 0.0)`, because `sup |Γ|` decays physically under viscosity and only an increase is a numerical violation.

I verified the consequence across the whole archive. In all 24 serial recomputation sidecars `boussinesq/regate_*.json`, spanning $\nu = 10^{-9}$ to 10^{-4} and four grid families, the signed drift is negative on every record and positive on none. The clip therefore returns exactly zero on every viscous record, so the circulation half of the trust conjunction is vacuously satisfied there and only the spectral-tail clause constrains a viscous snapshot. On inviscid runs the gate is active and passes with margin: the largest `|drift|` is 1.41×10^{-6} (`regate_D2048_inv.json`) against the 10^{-4} bar.

I record this as a scope statement, not as a defect, and I withdraw a stronger form of it that I circulated internally during the F2 cycle. The gate does what its design says: it is a tripwire for spurious growth of `sup |Γ|`, that failure mode did not occur, and the tripwire did not fire. My earlier internal phrasing, that any statement of the form “gated on tail and circulation” overstates what was enforced, is itself an overstatement and does not apply to the published record: the prior paper states the one-sidedness explicitly in the text that defines the gate, “decay under viscosity is physics; growth is a violation”.

What is true and load-bearing for this paper is narrower. The in-run gate constrained the *sign* of the circulation change and never its *magnitude*, and that magnitude is precisely the quantity the F2 cycle later promoted to a primary observable. Recovering it required recomputing the signed drift post hoc from the stored fields. No information was lost: the 24 sidecars carry `gamma_drift_signed` unclipped, and no clipped logged value was consumed by any verdict quantity in this paper.

Correction 2: the nominal TSTAR is not the singularity time

The campaign’s lane code carries `TSTAR = 0.0030`. A free fit of $\Omega(t)$ gives $t^* = 3.6185 \times 10^{-3}$ with SSR 5.93×10^{-5} against 0.1229 for the nominal value, a factor of 2070 (`ext_inertia_t2/FLUX_SPEC.md:538-546`, gate G-F10); the fit is well conditioned, with the SSR staying within $\times 1.05$ of its minimum only for $t^* \in [3.602653 \times 10^{-3}, 3.634667 \times 10^{-3}]$ (`ext_inertia_t2/out/bchan_depth/tstar_cond.log`). Every collapse time $s = -\log(1 - t/t^*)$ in this paper is on the fitted clock, never the nominal one (`ext_inertia_t2/out/flux3/F3_RESULT.md:169`).

The nominal value remains hard-coded in four sealed lane files: `ext_inertia_t2/lane2_overlap.py:87`, `ext_inertia_t2/laneG_final.py:27`, `ext_inertia_t2/laneG_meas.py:32`, and `ext_inertia_t2/laneG_verdict.py:31`; I verified all four lines in place. Sealed instruments are not edited even to fix known defects (`ext_inertia_t2/F3_SPEC.md:800-801`), so the defect is noted, not repaired: any s quoted from those four files is on the wrong clock.

Correction 3: duplicate writes in two archives

Two stored series carry duplicate records. T2048_1e-7 contains 17 exact restart re-emission duplicates at $t \in [0.002100141, 0.00232528]$; D2048_invB contains one undocumented duplicate write at $t = 1.199604 \times 10^{-3}$ (ext_inertia_t2/F3_SPEC.md:92-93). De-duplication is performed in memory on loaded copies for every quadrature; no file on disk is rewritten. The de-dup is itself gated: gate G-0 freezes the counts that must reproduce exactly, 17 removed from T2048_1e-7 and 1 from D2048_invB, with unique write counts 202 (D2048_invB), 198 (D2048_1e-5), 197 (T2048_1e-6), 212 (T2048_1e-7), 209 (T2048_1e-8) (ext_inertia_t2/F3_SPEC.md:424-429); the F3 ladder records spliced_dups = 17 on the T2048_1e-7 rung and 0 elsewhere (ext_inertia_t2/out/flux3/f3_ladder.json, rungs), and the live de-dup of the inviscid replicate is logged as “dropped 1/203” (ext_inertia_t2/out/flux/logA.txt). The include-versus-exclude difference on every quoted number enters the error budget as the splice term: 0.216% at $\nu = 10^{-7}$, 0.099% at 10^{-8} , zero elsewhere (ext_inertia_t2/out/flux/FLUX_RESULT.md:275).

Correction 4: the F2 deliverable debts

Two F2 deliverables were not what their names promised. First, ext_inertia_t2/out/flux/flux_ladder.json shipped its verdict, predictions, and exposures keys as empty containers (`{}`, `{}`, `[]`; verified by direct read), with the verdict existing only as prose in the result document. Second, ext_inertia_t2/out/flux/flux_ladder.log is not an execution log: it is a byte-copy of FLUX_RESULT.md. I hashed both files for this paper: md5 b6b32ff8f43bc3c0a656d685a387671d on each, and cmp reports them identical.

Both debts were named as mandatory non-repeats in the F3 pre-registration (ext_inertia_t2/F3_SPEC.md:785-791) and were repaired there: ext_inertia_t2/out/flux3/f3_ladder.json carries a populated verdict block, 11 scored predictions, and 6 exposures (verified by direct read), and f3_ladder.log is a real timestamped execution record (first entry 2026-08-13T01:41:38, md5 distinct from F3_RESULT.md). Every F3 headline number is re-derivable from the JSON alone.

Correction 5: two retractions, the F2d crossing and the A-1 label

I retract the F2d crossing claim. F2 reported the depth at which α_2 crosses 0.75 as a result corroborated by four routes (ext_inertia_t2/out/flux/FLUX_RESULT.md:441, headline at line 24). The four routes were one measurement. Two of the four recorded crossings are bitwise identical at $x_{\text{cross}} = 0.0027384615384615416$ (ext_inertia_t2/out/flux/flux_ladder.json, alpha_crosses_0.75, entries H1024|time and C1|H1024_fixed_1e-3..1e-5), and the audit found all four drawing on a single node family under a single dust rule, so on those axes the ensemble is one measurement reported four times (ext_inertia_t2/JUMP/chains/C03/journal.jsonl, entry 82). The crossing survives as a single-route observation; the corroboration structure it was sold with did not exist, and the claim as made is withdrawn.

I also retract the “A-1” campaign outcome label that the sealed F3 record attached to itself (ext_inertia_t2/out/flux3/F3_RESULT.md:17). A-1 requires at least three decidable rungs spanning at least two decades including both 10^{-7} and 10^{-8} (ext_inertia_t2/F3_SPEC.md:758-761). The primary (Simpson) scheme adjudicates $[10^{-7}, 10^{-5}]$; the 10^{-8} rung is decided by the trapezoid control alone, in class agreement, because the Simpson floor is unrealized in that rung’s only verdict window (ext_inertia_t2/out/flux3/F3_RESULT.md:169-171). The label compresses away the one fact a reader needs. The reach is quoted in full everywhere in this paper: primary $[10^{-7}, 10^{-5}]$, the 10^{-8} rung decided by the trapezoid control in class agreement, and never as “A-1”.

Correction 6: the withdrawn sealed A-channel T2b

The sealed A-channel verdict T2b (“the remainder is not linear response”) was the field-statistic arc’s mechanism claim. The repaired instrument reproduces the sealed curve first, to a maximum deviation of 4.91×10^{-5} , four-decimal rounding of the sealed quote (`ext_inertia_t2/out/bchan/BCHANNEL_RESULT.md:68-75`), so what follows is a statement about the same object and not a lookalike. Under the common evaluation frame and frozen reference frame, a zero-viscosity-contrast pair reproduces the T2b-shaped signal and the viscous signal is not separable from it: the verdict on every channel is R-VOID, and P-R4, the registered prediction that the sealed T2b would survive the repair, failed (`ext_inertia_t2/out/bchan_repair/REPAIR_RESULT.md:18,51-53,268-280`). I withdraw T2b. Under the repaired instrument no channel supports T2a, T2b, or T2c; the mechanism question passed to the flux observables of Sections 5 and 6, which is where it was answered.

Correction 7: three latent F3 code defects, disclosed, none outcome-changing

(a) *Verdict ordering.* The verdict routine grants F3a before checking VOID-4: in `ext_inertia_t2/out/flux3/f3_ladder_verdict.py:542-553` the branch `if all(f3a.values())` precedes the `void4_inseparable` test. The frozen discipline is that VOID beats shaky; the code inverts the precedence, so an F3a coinciding with model-inseparability would have been reported as F3a. Not outcome-changing on this data: model separability held wherever fits were quoted, and VOID-4 did not fire on any statistic or scheme (`ext_inertia_t2/out/flux3/F3_RESULT.md:127-130`).

(b) *A disclosure boolean keyed to the median.* Gate G-D’s machine record declares “an offset over $1e-3$ or any use of the nominal clock fails this gate,” yet its pass boolean `sub_permille` was computed from the median offset. On the deepest rung the measured snapshot-grid versus stream-dealias sup offset over 42 exact- t depth nodes reads median 4.00×10^{-5} , p95 7.19×10^{-4} , max 1.06×10^{-2} (`ext_inertia_t2/out/flux3/f3_gates.log:268; f3_gates.json, gates.G-D, rel_offset_max = 0.010595388506380617` with `sub_permille: true`). Read against the max, the clause fails on that rung. Not outcome-changing: the offset is a labeling comparison between two depth conventions, and every windowed quantity in the ladder is keyed to the snapshot-grid convention on both sides of every ratio (`f3_gates.json, conventions_carried`), so the offset never enters a verdict quantity; the sealed gate table already carried the honest wording “sub-permille median” (`ext_inertia_t2/out/flux3/F3_RESULT.md:33`). The boolean is the defect, and it is disclosed here.

(c) *A must-fail null with a degenerate metric.* Gate G-N(a) demands that the raw inviscid Z drift, computed without the production subtraction, be large (≥ 0.1): a pipeline that accidentally conserved raw Z must fail it. On this initial condition $Z(0) = 0$ identically (`ext_inertia_t2/out/flux3/f3_gates.json, gates.G-N, Z0_identically_zero_on_this_IC: true`), so the implemented metric $|Z_{\text{end}} - Z_0|/Z_{\text{end}}$ returns exactly 1.0 for any pipeline with nonzero Z_{end} , including the broken pipelines the null exists to catch, and the drift relative to Z_0 is undefined. The gate’s conclusion (raw Z is massively non-conserved) is true, and F2’s independent full- Z variant measured the same fact non-degenerately at 0.342 (`f3_gates.json, gates.G-N, measured record`), but on this IC the F3 metric could not have concluded otherwise. The gate’s can-fail clause was an assertion, not a measurement, and it is disclosed as such.

9 Scope and open problems

9.1 Pointwise T2

F3’s frozen scope sentence is the boundary and I repeat it verbatim: F3 measures the viscous enstrophy flux of the ω_1 channel as an integrated functional over $\nu \in [10^{-8}, 10^{-5}]$ within the

stored collapse window; it does not resolve the pointwise T2 question, which remains open (`ext_inertia_t2/out/flux3/F3_RESULT.md`, section 6). Nothing in this paper constrains $\omega_{1,\nu} - \omega_{1,\text{inv}} = \nu \delta_1 + o(\nu)$ pointwise, least of all at the corner ring; nothing is in sup-norm; rearrangements preserving the gradient integral are invisible to Z ; window-local violations can cancel inside a window, mitigated but not removed by the nine-window report. Section 4 explains why the pointwise question is hard here rather than merely unfinished: the sealed record has viscosity acting through the collapse frame itself, and a frame-quotiented statistic deleted 71.8 to 99% of its own prediction (the root-cause table, Section 4). A pointwise instrument for this collapse therefore must not quotient the frame, and no such instrument with a theorem-backed null is in hand. The one concrete function-space lead comes from the corner treatment of Chen and Hou, who carry the corner with singular weights and explicit corrections rather than excision [5]: a pointwise corner response coefficient would need that function-space change, not grid refinement.

9.2 The sub-leading exponent

The joint license's floor, corrections vanishing no slower than $\nu^{0.4}$, is a band edge, not a measurement. The F3 bulk correction fit returns $q = 0.438$ (Simpson, window [200, 400]) and $q = 0.425$ (trapezoid, window [1200, Ω_*], jackknife spread 0.088) against the frozen band [0.4, 1.0]: both sit at the edge, so the exponent is unresolved (`ext_inertia_t2/out/flux3/f3_ladder.json`, family_A; `F3_RESULT.md`, section 9). The F2 swirl correction exponent at fixed time is 0.7130, drop-one jackknife spread 0.029, profile [0.710, 0.716], and it lives at $\nu \in [10^{-5}, 10^{-3}]$ (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 4.2). Different channel, different pairing, different ν range: I do not reconcile the two numbers and this archive cannot. Resolving q is a floor problem; see the archive limits below.

9.3 R2: azimuthal stability, parked, with the certified route

Whether the corner ring survives azimuthal symmetry breaking is not addressed by any measurement in this paper. The question is parked with a certified route on disk. The Tier-0 pre-registration (`ext_inertia_t2/R2_PLAN/TIER0_SPEC.md`) froze the route: an m-mode operator about the certified profile, azimuthal coupling in a polynomial normal form $L_m = L_0 + \tilde{k}B_1 + \tilde{k}^2B_2$ at $\tilde{k} = m\ell(t)$, with the certified margin $m_{\text{cert}} = 1/(2\lambda_{\max}(P))$ from the Lyapunov certificate. The execution (`ext_inertia_t2/R2_PLAN/TIER0_RESULT.md`) returned T0c, VOID by the spec's own stop rule, and the void has content. Certified and re-derived from P : $m_{\text{cert}} = 8.8232737664 \times 10^{-6}$, and $\text{Re } \lambda \leq -8.8233 \times 10^{-6}$ for every eigenvalue of the reduced generator, no eigensolve used. Measured obstructions: the 2D certificate cannot host the azimuthal couplings, since 82.71% of its nodes sit at cylindrical radius $R \leq 0$ and those rows carry 91.27% of $\text{tr}(P)$, with the $1/R$ pole interior to the discretization; and my polynomial ansatz was wrong twice over, because the dominant azimuthal channel is the swirl advection entering at $\kappa = m(t_s - t)$, which exceeds $\tilde{k}$ by $\ell^{-0.658}$ at $c_l = 2.9206$, and because pressure slaving makes the m -dependence operator-rational rather than polynomial. The honest normal form is two-parameter, $(\kappa, \tilde{k})$, and the certified band depth for a fixed m scales as m^{-1} (swirl channel), not $m^{-0.342}$. The corner circle, the one locus where the azimuthal coefficients are unambiguous ($R = 1$ exactly at $\rho = 0$), is exactly the ring the certificate deletes. The closing route is itemized in the result file, first item a 3D corner similarity system; the fixed- $\tilde{k}$ ladder remains the right experiment. Two disciplines bind any future R2 run. First, perturbation growth is never a falsifier: blowup solutions of this type are hydrodynamically unstable by theorem [17], so discriminators must be structural. Second, no unconditional linear stability of the profile is claimed here or anywhere in this paper: the certificate is corner-clamped, in the P-norm, with measured transient growth $89.7 \leq \sup_t \|e^{tL}\| \leq 5807$ (Section 2), and that caveat travels with every mention of the certificate.

9.4 Archive limits, and what a purpose-built archive would buy

The deep end of this archive is a floor property, not an estimator property. In the swirl channel the fixed-time deviation falls to the inviscid instrument floor (1.80×10^{-10} on the D and T grids) below $\nu \approx 3 \times 10^{-6}$; the deepest rung that ever clears the deviation gate is 3×10^{-6} , on D2048 at $t \geq 2.85 \times 10^{-3}$, so the F2 exponent lives at $\nu \in [10^{-5}, 10^{-3}]$ and nothing below 10^{-6} is decidable there (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, sections 6.2 and R1). In the ω_1 channel the 10^{-8} rung's deepest window is floor-unrealized under the primary scheme (cross-realization ratio 12.7 on roundoff-scale floors) and is decided by the trapezoid control in class agreement; the grid sigma term, 0.068, forced by the absent `snap_D2048_1e-6` cell, is the largest single error term on every T rung; the floors are single-realization on two of the three grids (in the swirl lane the T and H floors inherit the D-grid replicate fraction, 1.20%, `ext_inertia_t2/out/flux/FLUX_RESULT.md`, section 2.5; in the ω_1 lane no inviscid field replicate exists at all and the floor-realization spread across independent executions is unmeasured, substituted by the 5 to 7% cross-grid floor transfer, `F3_RESULT.md`, exposure E-1); and the thin deep bin $[\Omega_*, 1900)$ holds five intervals at 10^{-8} (`ext_inertia_t2/out/flux3/F3_RESULT.md`, sections 5 and 9, exposures E-1..E-3). A purpose-built archive would need: an inviscid replicate pair on the T grid; snapshot storage, not streams, for every rung entering a fit; denser write cadence in the corner window $\Omega \in [1200, 1900)$; and the $\nu = 10^{-6}$ same-depth D/T cell. That list buys floors and cross-checks. It creates no new physics claims, and I make none from it.

9.5 What is not claimed

Stated explicitly, each with the reach that bounds it.

1. Nothing about generic Navier-Stokes blowup. Every result in this paper is about one certified profile ($\alpha = -0.34240 \pm 3 \times 10^{-5}$, Section 2), one geometry, one boundary family.
2. Nothing about no-slip walls beyond the wall-swap probe of Section 7. The no-slip ladder was a boundary-condition experiment on the clamp hypothesis, thirteen completed short runs, formally W3; the no-slip half-space regularity theory is prior art [2, 10] and this campaign adds nothing to it. A direct no-slip comparison of the direction-regularity exponent was attempted on the stored no-slip family (18 August 2026, pre-registered, census-gated) and refused itself: the impulsive-start Rayleigh sheet attains the vorticity supremum on every gated snapshot of all fourteen runs, at the quarter-period rings $z = L/4, 3L/4$ rather than at the corner, so the certified observable reads a different object under no-slip walls and no cross-boundary-condition comparison is licensed from that data.
3. Nothing pointwise, at any ν . T2 is open (Section 9.1).
4. No single scalar sub-leading exponent. The matching-convention fork fires in the swirl channel (VOID-5 / V-4, Section 5) and the bulk-channel q sits at its band edge (Section 9.2).
5. Nothing below $\nu = 10^{-8}$, and nothing deeper than the stored window ($t \leq 2.985 \times 10^{-3}$, $\Omega \lesssim 1780$). No number in this paper is an extrapolation beyond its measured range.
6. No corner-mechanism attribution. The corner share never reaches its gate (maximum 0.412 against 0.5, Section 7), and corner language is licensed in neither flux lane.
7. No unconditional linear stability of the profile, and no claim about azimuthal stability: R2 is parked (Section 9.3), and the certificate's transient-growth caveat stands.

8. No anomalous-dissipation statement outside this window and geometry. The zero-anomaly verdict of Section 3 concerns this collapse through Ω_* ; it says nothing about Onsager-class dissipation in turbulence [9, 16].

10 Methods

Pre-registration

Every lane ran against a spec frozen before any number existed: `ext_inertia_t2/FLUX_SPEC.md`, `F3_SPEC.md`, `DEPTH_MATCH_SPEC.md`, `BCHANNEL_SPEC.md`, `FRAME_REPAIR_SPEC.md`. A spec fixes the observable, the numeric bars with their provenance stated before the bars are used (`ext_inertia_t2/F3_SPEC.md:539-556`, section 5.1), the verdict classes, and the predictions. Verdict rules are applied once, at the end, on the complete ladder (`ext_inertia_t2/FLUX_SPEC.md:549`; `F3_SPEC.md:536`). VOID beats shaky (`ext_inertia_t2/F3_SPEC.md:826`; `DEPTH_MATCH_SPEC.md:952`): an uncertifiable number is voided with its specific failure, never quoted with a caveat, and “not resolvable” is a result distinct from VOID. Changing any threshold after any number is computed voids the entire run and forces re-registration (`ext_inertia_t2/DEPTH_MATCH_SPEC.md:820-827`). Missed bars are reported as missed, never moved: F2 missed two of its own bars, G-F4a’s 1.0 and P8’s box range, and reports both (`ext_inertia_t2/out/flux/FLUX_RESULT.md:631`). Registered predictions are scored one by one in the sealed result, including the failed and the half-inverted (`ext_inertia_t2/out/flux/FLUX_RESULT.md`, P1..P11; `out/flux3/F3_RESULT.md`, P-1..P-11; `DEPTH_MATCH_SPEC.md`, P-D1..P-D4).

Error budget and jackknife, and why bootstrap is invalid here

Every quoted uncertainty is a deterministic instrument budget: σ is the absolute sum of six measured terms, cadence, spatial quadrature, depth-local floor realization, interpolation, splice, and grid (`ext_inertia_t2/F3_SPEC.md:549-555`; per-window values in `out/flux3/f3_ladder.json`, `sigma_budget`); ρ_{bulk} adds the $\{3, 8\}$ wall-strip fork half-spread (`ext_inertia_t2/out/flux3/F3_RESULT.md:105`). Every fitted coefficient and exponent additionally carries a leave-one-rung-out jackknife (spreads and CIs in `out/flux3/F3_RESULT.md:127-130`).

I do not bootstrap, and the reason is structural, not stylistic. There is one run per rung and the primary observable is a volume functional of stored fields, so no repeated-sampling distribution exists to resample. An earlier cycle’s row-block bootstrap was additionally invalid as executed: 10 to 12 blocks with lag-1 residual correlation 0.61..0.91, and it is not reused (`ext_inertia_t2/out/flux/FLUX_RESULT.md:268`). The frozen discipline states the consequence plainly: any confidence-interval language beyond the deterministic budget plus jackknife is fabricated (`ext_inertia_t2/F3_SPEC.md:554-555`). The one place a genuine ensemble exists, the seeded Lagrangian Γ meter (`SEED = 20260812`, `NPART = 2000`, `nsub = 8`), gates admissibility and never enters the primary (`ext_inertia_t2/out/flux/FLUX_RESULT.md:268`).

Depth matching

Viscous runs are compared at matched collapse depth Ω , not matched wall-clock time. Matching uses the certified 4-point cubic Lagrange stencil in $\log \Omega$ from the depth-matching instrument (`ext_inertia_t2/out/bchan_depth/`). Its leave-one-out error gate measured reference-side E_{LOO} between 0.006% and 0.498% of signal across rungs against a 5% bar (`ext_inertia_t2/out/bchan_depth/dmx_depth_gates.log:108-113`). The depth-collapse gate G-D2 drives the adjacent-write pair, a pure depth displacement carrying zero viscosity contrast, down by factors $\times 745.. \times 848$ against a threshold of 100 (`dmx_depth_gates.log:73-79`); the corresponding residual depth-offset suppression of $92\times$ to $668\times$ was registered in advance as

prediction P-D1 (`ext_inertia_t2/DEPTH_MATCH_SPEC.md:833-838`). The matching anchor $\Omega^* = 1618.3879571332704$ is exactly the final snapshot depth of D2048_1e-5, so that run's matched-depth values sit on the last interpolation node; every use of the upper edge carries the E-4 flag (`ext_inertia_t2/out/flux3/f3_gates.json, gates.G-D`).

Gates and the can-fail requirement

Every gate in this campaign states what data enters it and how it can fail (`ext_inertia_t2/F3_SPEC.md:348-350`). The requirement exists because of a measured failure, not a principle: an earlier depth gate whose held-out-write null compared the reference to itself scored PASS at $\times 811$ (`ext_inertia_t2/F3_SPEC.md:348-349`; the $\times 810.7$ record at `out/bchan_depth/dmx_depth_gates.log:73`). A gate whose failure branch cannot execute on the data at hand is an assertion, and the ledger's items 1 and 7(c) are this campaign's two instances, both found by applying the test to my own gates. Null checks run in polarity pairs: a must-fail null (G-N(a)) beside a must-return-zero null (G-N(b)), cross-Casimir $\leq 1.65 \times 10^{-15}$ against a 10^{-12} bar, with a deliberate parity-break demonstration firing at factor 2079 over bar on the inviscid replicate) (`ext_inertia_t2/out/flux3/F3_RESULT.md:147-148; f3_gates.json, gates.G-N`).

Adversarial refutation

Before a verdict is quoted, the result is attacked through three independent review lenses that answer alone, with three separate verdicts that are never pre-merged; recorded disagreement between lenses is the information, and a lens that always agrees is not earning its slot (`lenses/README.md`). The posture is default-refute: each lens tries to kill the claim, and the claim is quoted only if no attack lands. The on-disk instantiations are the five adversarial probe batteries `adv_stat_probe` through `adv_stat_probe5` and `adv_depth_decomp` in `ext_inertia_t2/out/bchan_depth/`, `av1..av8` in `out/flux/adv_verify/`, and the independent recomputations `a_recompute.json` and `floors_recompute.json` in `out/flux/judge_adv/`. The certified α result underwent the same triple pass before publication: a numbers-tracer lens (63 claims traced, 15 of 15 live re-runs matched, zero material mismatches), an adversarial referee lens (14 overclaims fixed), and a technical lens (published assets hash-identical to sources), recorded at `eja_state/journal.jsonl`, entry 454.

Integrity

All input trees are read-only, enforced by md5 manifests taken at phase open and re-taken at phase close (rule VOID-6). The F3 ladder and verdict phases each manifest 1590 files, 63.93 GB (63,927,727,710 bytes), re-hashed at close with zero files changed, deleted, or appeared (`ext_inertia_t2/out/flux3/f3_manifest_diff.json` and `f3_verdict_manifest_diff.json`); the F2 phase used a `find+stat` manifest over 1394 files, byte-identical in name, size, mtime, and inode at close (`ext_inertia_t2/out/flux/FLUX_RESULT.md:625`); the depth phase manifested 61 run directories (`ext_inertia_t2/out/bchan_depth/dmx_manifest_open_dirs.txt`, 61 entries). Run data is never deleted, moved, renamed, or overwritten, anywhere; de-duplication happens in memory on loaded copies only (`ext_inertia_t2/out/flux/FLUX_RESULT.md:626; F3_SPEC.md:802-804`). No DNS was re-run for any number in this paper: everything is computed from stored snapshots and stream sidecars. No bytecode is emitted into read-only directories (`-B` with `PYTHONDONTWRITEBYTECODE=1`; `ext_inertia_t2/out/flux/FLUX_RESULT.md:628`). Compute hygiene on every analysis process: `nice -19`, single-threaded BLAS pins, and the `.venv-dedalus` interpreter (`ext_inertia_t2/out/flux/FLUX_RESULT.md:629`).

References

- [1] Leiden declaration on artificial intelligence and mathematics, June 2026. URL <https://leidendeclaration.ai>.
- [2] Tobias Barker and Christophe Prange. Scale-invariant estimates and vorticity alignment for Navier-Stokes in the half-space with no-slip boundary conditions. *Archive for Rational Mechanics and Analysis*, 235:881–926, 2020.
- [3] Hugo Beirão da Veiga and Luigi C. Berselli. On the regularizing effect of the vorticity direction in incompressible viscous flows. *Differential and Integral Equations*, 15(3):345–356, 2002.
- [4] Keaton J. Burns, Geoffrey M. Vasil, Jeffrey S. Oishi, Daniel Lecoanet, and Benjamin P. Brown. Dedalus: a flexible framework for numerical simulations with spectral methods. *Physical Review Research*, 2:023068, 2020. doi: 10.1103/PhysRevResearch.2.023068.
- [5] Jiajie Chen and Thomas Y. Hou. Stable nearly self-similar blowup of the 2D Boussinesq and 3D Euler equations with smooth data I: analysis of the velocity field, 2022.
- [6] Jiajie Chen and Thomas Y. Hou. Stable nearly self-similar blowup of the 2D Boussinesq and 3D Euler equations with smooth data II: rigorous numerics, 2023.
- [7] Peter Constantin and Charles Fefferman. Direction of vorticity and the problem of global regularity for the Navier-Stokes equations. *Indiana University Mathematics Journal*, 42(3):775–789, 1993.
- [8] Peter Constantin, Weinan E, and Edriss S. Titi. Onsager’s conjecture on the energy conservation for solutions of Euler’s equation. *Communications in Mathematical Physics*, 165:207–209, 1994. doi: 10.1007/BF02099744.
- [9] Jean Duchon and Raoul Robert. Inertial energy dissipation for weak solutions of incompressible Euler and Navier-Stokes equations. *Nonlinearity*, 13(1):249–255, 2000. doi: 10.1088/0951-7715/13/1/312.
- [10] Yoshikazu Giga and Hideyuki Miura. On vorticity directions near singularities for the Navier-Stokes flows with infinite energy. *Communications in Mathematical Physics*, 303:289–300, 2011. doi: 10.1007/s00220-011-1197-x.
- [11] Thomas Y. Hou and Ruo Li. Computing nearly singular solutions using pseudo-spectral methods. *Journal of Computational Physics*, 226(1):379–397, 2007. doi: 10.1016/j.jcp.2007.04.014.
- [12] Tosio Kato. Remarks on zero viscosity limit for nonstationary Navier-Stokes flows with boundary. In S. S. Chern, editor, *Seminar on Nonlinear Partial Differential Equations*, volume 2 of *Mathematical Sciences Research Institute Publications*, pages 85–98. Springer, New York, 1984.
- [13] Laurent Laffèche, Alexis F. Vasseur, and Misha Vishik. Instability for axisymmetric blow-up solutions to incompressible Euler equations. *Journal de Mathématiques Pures et Appliquées*, 155:140–154, 2021. doi: 10.1016/j.matpur.2021.02.006.
- [14] Guo Luo and Thomas Y. Hou. Potentially singular solutions of the 3D axisymmetric Euler equations. *Proceedings of the National Academy of Sciences*, 111(36):12968–12973, 2014. doi: 10.1073/pnas.1405238111.

- [15] Guo Luo and Thomas Y. Hou. Toward the finite-time blowup of the 3D axisymmetric Euler equations: a numerical investigation. *Multiscale Modeling & Simulation*, 12(4):1722–1776, 2014. doi: 10.1137/140966411.
- [16] Lars Onsager. Statistical hydrodynamics. *Il Nuovo Cimento (Supplemento)*, 6:279–287, 1949. doi: 10.1007/BF02780991.
- [17] Alexis F. Vasseur and Misha Vishik. Blow-up solutions to 3D Euler are hydrodynamically unstable. *Communications in Mathematical Physics*, 378:557–568, 2020. doi: 10.1007/s00220-020-03790-5.
- [18] Y. Wang, M. Bennani, J. Martens, J. Gómez-Serrano, R. Jiang, C.-Y. Lai, et al. Discovery of unstable singularities, 2025.
- [19] Yongji Wang, Ching-Yao Lai, Javier Gómez-Serrano, and Tristan Buckmaster. Asymptotic self-similar blow-up profile for three-dimensional axisymmetric Euler equations using neural networks. *Physical Review Letters*, 130:244002, 2023. doi: 10.1103/PhysRevLett.130.244002.