

Viscosity inverts the direction-regularity scaling of the corner-flow blowup mechanism: a measured exponent crosses the type-I exclusion threshold

Justin Hill
Independent Research
justin@epagoge.io

Draft 0.1 — 2026-08-02*

Abstract

For vorticity ω with direction $\xi = \omega/|\omega|$ and peak set $P(t)$, the quantity $\Lambda = \sup_P |\nabla \xi| |\omega|^{-1/2}$ is scale invariant under the Navier–Stokes rescaling, and its exponent $\sigma_\Lambda = d \ln \Lambda / d \ln \|\omega\|_\infty$ turns the direction-coherence regularity criteria of Constantin–Fefferman and Beirão da Veiga–Berselli into a single measurable number. We measure it. On the inviscid Luo–Hou corner flow, where blowup is proven in the Boussinesq analogue by Chen–Hou, the instrument returns $\sigma_\Lambda = +1.00 \pm 0.03$, anchored by an exact amplitude symmetry (predicted Λ ratio 0.707107, measured 0.706891) and a 2×2 off-ray grid factorial (spread 0.029). Under free-slip viscosity, $\sigma(A)$ is not one number: every viscous run crosses over from inviscid-like scaling at low amplitude to deep depletion at high amplitude; window-averaged slopes are regime mixtures and we retract our own earlier quotes of that kind. Measured in the asymptotic window (top half of the cross-grid amplitude overlap, symmetric midpoint cut, pre-registered), the deep-collapse exponent is $\sigma_\Lambda(10^{-4}) = -1.12 / -1.24$ (grids $128 \times 384 / 256 \times 768$, spread 0.121) and $\sigma_\Lambda(10^{-3}) = -1.25 / -1.29$ (spread 0.032), against a window-matched inviscid $+0.589 / +0.585$ (spread 0.005). Viscosity does not merely damp the corner mechanism’s alignment failure: it inverts it, by roughly 1.8 powers of $\|\omega\|_\infty$, to a value far below both thresholds of the exclusion chain ($-1/6$ for onset of geometric depletion, $-1/2$ for exclusion of type-I blowup by geometry plus energy, under one structural hypothesis with a measured constant, Section 4). As far as a live literature kill-search reaches, no direction-regularity exponent has previously been measured as a function of viscosity on any blowup candidate. On the dynamical side, a Lyapunov-certified march in the rescaled frame (definiteness by Cholesky only; no eigenvalue enters any decision) shows the collapse profile is an attractor: perturbed orbits settle with $\cos_{\text{step}} = 0.986$ and contraction rate 0.270 ± 0.003 per rescaled time unit, unanimous across five independent seeds at the largest basin-validated amplitude.

1 Introduction

Two problems sit on one function. If a viscous singularity inheriting the corner mechanism exists, it must keep $\sigma_\Lambda(\nu)$ above $-1/6$ as the collapse deepens. If viscosity forces σ_Λ below $-1/2$

*Every number in this paper regenerates from bytes on disk with one command (Section 8). Claims are graded against a frozen standard; the grading ledger, the confound engine, and the failed versions of our own claims are part of the repository, not part of a memory hole.

in the high-vorticity regime, the geometric route closes the mechanism without touching the supercritical energy problem: geometry, not energy, kills the singularity. Nobody had $\sigma_\Lambda(\nu)$ that our kill-search could find (Section 7). This paper reports it, for the mechanism class where the inviscid answer is proven, at two viscosities, on two grids, with the disagreement between grids quantified and under a bar.

What is not new here, stated first. Direction-coherence regularity criteria originate with Constantin and Fefferman [2] and were sharpened to the $1/2$ -Hölder class by Beirão da Veiga and Berselli [4, 5]: the pairing of $|\nabla\xi|$ with the $1/2$ power of amplitude lives in their sharp exponent before it lives in our observable. Qualitative exclusion of type-I blowup under direction continuity is likewise prior art: Giga–Miura [6] for the infinite-energy class, and Barker–Prange [7] with scale-invariant alignment estimates in the half-space with no-slip boundary, which is the Luo–Hou geometry itself. The identity behind our instrument check (Corollary 0) is Constantin’s classical decomposition of the vorticity equation into magnitude and direction [3]. A 2026 analytical line by Grujić [8] pursues logarithmic depletion through direction classes; it defines no observable and measures nothing.

What is new, correspondingly narrowed. First, the observable program: treating the scale-forced pairing as a measurable scalar with an exponent, and calibrating the instrument on a case where blowup is proven [11, 12], so that its verdict means something where blowup is unknown. Second, the numbers: $\sigma_\Lambda = +1.00 \pm 0.03$ on the inviscid corner mechanism, the first direction-regularity exponent measured on any blowup candidate as far as our search reaches. Third, the viscous inversion: $\sigma_\Lambda(\nu)$ measured at two viscosities with cross-grid certification, crossing both thresholds deep in the collapse. Fourth, the quantified exclusion chain (Section 4): the known qualitative principle [6, 7] sharpened to an explicit threshold on the measured exponent by feeding the energy dissipation budget through Constantin’s kernel split, with the structural hypothesis carrying a measured constant.

The laboratory operates under a discipline that we consider part of the result (Section 8): a backwards proof of fifteen checks that failed its own authors until the fit protocol was encoded, a confound engine enumerating the ways each claim could be false, exact-symmetry anchors with zero free parameters, and a standing watcher that reruns the certification every twenty minutes. Several of our own earlier numbers died by these instruments and are retained in the repository as corpses with names.

2 The observable

For a Navier–Stokes (or Euler, or Boussinesq) vorticity field, let $\xi = \omega/|\omega|$ and $P(t) = \{x : |\omega(x, t)| \geq \frac{1}{2}\|\omega\|_\infty\}$, and define

$$\Lambda(t) = \sup_{P(t)} |\nabla\xi| |\omega|^{-1/2}, \quad \sigma_\Lambda = \frac{d \ln \Lambda}{d \ln \|\omega\|_\infty}. \quad (1)$$

Under $u \rightarrow s u(sx, s^2t)$ the field $|\omega|$ carries s^2 and $|\nabla\xi|$ carries s , so Λ is exactly scale invariant; the $1/2$ power is forced by the group, and $|\xi| = 1$ makes the direction field globally controlled for free. The same pairing appears, in criterion form, as the $\beta = 1/2$ sharpness in [4, 5]; the observable is that criterion turned into a number a simulation can return.

Corollary 0 (of Constantin’s $|\omega|$ -equation [3]; stated as a corollary, used as an instrument check). *The magnitude of vorticity obeys $\partial_t |\omega| + u \cdot \nabla |\omega| = \alpha |\omega| + \nu (\Delta |\omega| - |\omega| |\nabla\xi|^2)$, where*

$\alpha = \xi \cdot S\xi$ is the stretching rate. At any growing spatial maximum of $|\omega|$ the Laplacian term is nonpositive, hence

$$\nu |\nabla \xi|^2 \leq \alpha \quad \text{at the maximum.} \quad (2)$$

This is a one-step corollary of a classical identity and we claim no novelty for it; a literature kill-search retracted our earlier framing of it as a theorem (Section 7). It earns its place as a validation: our data satisfies it with three orders of magnitude of margin at every gated snapshot (campaign ledger), which any correct direction-field computation must.

3 Instruments

The DNS side is a Dedalus-based axisymmetric Boussinesq solver in the Luo–Hou corner scenario, with free-slip viscosity in 5D-Laplacian form ($\nu \Delta_5$ on u_1 and ω_1 , $\Delta_5 = \partial_{rr} + (3/r)\partial_r + \partial_{zz}$), implicit in the IMEX stepper, and a quartic initial condition exactly compatible with the Neumann wall condition. Trust gates per snapshot: spectral tail at or below 10^{-6} and signed drift of the conserved circulation invariant $\sup |r^2 u_1|$ (decay under viscosity is physics; growth is a violation). Λ is evaluated on the $2 \times$ -HWHM box at the vorticity peak; runs are trimmed at the first $10 \times$ amplitude jump. The exponent is the slope of $\ln \Lambda$ against $\ln \|\omega\|_\infty$ over gated snapshots, with pair-resampled bootstrap confidence intervals (10^4 resamples) everywhere a slope is quoted.

The rescaled-frame side is a corner-frame s -march on the certified self-similar profile: backward Euler on the index-2 DAE with the algebraic manifold re-imposed at every iterate, a frozen reduced Jacobian whose exact validity on $\ker(C_g)$ is a closed-form identity (verified at 4.6×10^{-14} against the exact marcher before being trusted), and a Lyapunov P -norm gate in which definiteness is decided by Cholesky factorization only. No eigenvalue of any operator enters any decision anywhere in this paper; that refusal is standing and was never violated under pressure.

4 The exclusion chain, quantified

Constantin’s identity for the stretching rate $\alpha = \xi \cdot S\xi$ consumes the Biot–Savart kernel geometry: misalignment is the only way to stretch. Splitting the integral at radius ρ , the near field is controlled by $\Lambda \|\omega\|_\infty^{3/2} \rho$ and the far field by the energy through Cauchy–Schwarz; optimizing ρ ,

$$\alpha \lesssim \left(\Lambda \|\omega\|_\infty^{3/2} \right)^{3/5} \|\omega\|_{L^2}^{2/5}. \quad (3)$$

Hypothesis (H): the peak set is a single coherent structure on which the gradient bound holds out to the optimal ρ . The hypothesis carries a measured constant ($\lambda_0 = 5.0$) and the chain inequality held at 21 of 21 gated snapshot pairs with the stretching rate inferred independently (campaign ledger, STANDARD.md; that run predates the artifacts retained in this campaign’s outputs). Feeding the only global control Navier–Stokes provides (the energy dissipation identity) through this bound, with blowup rate $\|\omega\|_\infty \sim (T - t)^{-\gamma}$, $\gamma \geq 1$:

$$\sigma_\Lambda \leq -\frac{1}{6} : \text{geometric depletion begins to act at all,} \quad (4)$$

$$\sigma_\Lambda < -\frac{1}{2} : \text{type-I blowup excluded by geometry and energy alone.} \quad (5)$$

The qualitative principle here is not ours: under a type-I condition, direction continuity where vorticity is large already excludes blowup [6], including in the half-space with no-slip boundary

[7]. What the chain above adds is the exponent: an explicit threshold on a measurable number, so that a simulation, or eventually an experiment, can locate a given flow on the right or wrong side of it.

5 The inviscid calibration

On the inviscid corner flow the instrument returns

$$\sigma_{\Lambda} = +1.00 \pm 0.03. \quad (6)$$

Two validations with zero free parameters between them. First, the exact amplitude symmetry $u \rightarrow \lambda u$, $t \rightarrow t/\lambda$ of the Euler equations predicts a Λ ratio of $1/\sqrt{2} = 0.707107$ between paired runs; measured 0.706891 ± 0.001441 , three parts in ten thousand, with the predicted $|\nabla \xi|$ ratio of 1 measured at 0.999746. Second, a 2×2 off-ray factorial refining N_z and N_r independently returns slopes $+1.0159$, $+0.9873$, $+1.0027$, $+0.9978$ (spread 0.029, main effects at or below 0.017), rejecting the axial-mesh and single-ray confounds by test rather than by assertion.

Reading: $|\nabla \xi| \sim \|\omega\|_{\infty}^{3/2}$. Direction regularity fails a full $3/2$ of a power above the exclusion threshold; the Euler mechanism blows up because its geometry is maximally non-depleting. The observable returns the correct sign on a case where blowup is proven in the Boussinesq analogue [11, 12], which is precisely what licenses pointing it at cases where the answer is unknown.

Refinement recorded during the viscous campaign: the full-window $+1.00$ is amplitude-composite. Fit on matched cross-grid amplitude windows, the inviscid slope runs $+1.42/+1.36$ in the lower half and $+0.589/+0.585$ in the deep half (cross-grid spread 0.005, the tightest certification in this paper; retained artifact `M4_SIGMA_DEEP_NU0.out`). The symmetry and factorial validations certify the instrument, not slope constancy across amplitude; window-matched comparisons below therefore use $+0.59$ as the inviscid contrast value.

6 The viscous measurement

Runs at $\nu \in \{10^{-4}, 10^{-3}\}$ on grids 128×384 and 256×768 with inviscid references on both grids. Three findings, in discovery order.

(i) *The common-amplitude overlap window reproduces the calibration.* Fit on the same amplitude range, the two inviscid grids give $+1.016/+1.033$ (spread 0.017) against the calibrated $+1.00 \pm 0.03$. The cross-grid disagreement of naive full-window fits (0.147) was a window artifact, not a physics disagreement.

(ii) *At finite ν , $\sigma(A)$ is not one number.* Every viscous run crosses over from inviscid-like slope ($+0.39$ to $+0.81$) at low amplitude to deep depletion near -1.2 at high amplitude, on both grids, at both viscosities. A window-averaged slope is a regime mixture; our own earlier quotes of that kind (window averages near -0.5 , and a single-grid ladder -0.37 to -0.57) are retracted as measurements of a mixture, and the retraction is part of the repository record.

(iii) *Measured where the asymptotic regime actually lives* (top half of the cross-grid overlap in $\ln A$, symmetric midpoint cut, pre-registered before the cut-sensitivity scan), the deep-collapse exponent is

	128×384	256×768	spread
$\nu = 10^{-4}$	-1.116	-1.237	0.121
$\nu = 10^{-3}$	-1.254	-1.285	0.032
$\nu = 0$ (matched window)	+0.589	+0.585	0.005

with bootstrap 95% intervals of half-width ± 0.14 to ± 0.30 on each fit and cut-position sensitivity reported at 0.4/0.5/0.6 of the overlap span: the magnitude sits between -0.99 and -1.31 at every cut, on every grid, at both viscosities. Against the window-matched inviscid $+0.59$, viscosity inverts the direction-regularity scaling by roughly 1.8 powers of $\|\omega\|_\infty$, to a value below both thresholds of Section 4.

Honest bounds on what this shows. Two viscosities, one mechanism class, one scenario geometry, laptop resolution; the deep windows carry 11 to 14 snapshots each; the crossover amplitude $A_c(\nu)$ is visible but not yet resolved as its own observable; and the exclusion chain runs through Hypothesis (H) with its measured constant. What the measurement supports, within those bounds, is exact: on the strongest known blowup mechanism of this class, the geometric quantity that must stay above $-1/6$ for the mechanism to survive viscosity is measured at -1.2 and below, deep in the collapse, grid-certified, at both viscosities tried.

7 The kill-search

A live literature search (2026-08-02; arXiv, Springer, journal indices; queries on direction-coherence criteria, Hölder-1/2 direction results, scale-invariant alignment estimates, type-I exclusion via direction, Luo–Hou alignment measurements, and name collisions for σ_Λ) graded every novelty-bearing claim. One claim died: our Theorem 0 is a one-step corollary of Constantin’s $|\omega|$ -equation and is demoted to Corollary 0 above. Theorem-1-as-principle died too, as it should: the qualitative type-I exclusion belongs to [6] and [7]; what survives is the quantification. The observable program, the inviscid measurement, and the viscous inversion survived the search outright, with the standing caveat that survived means *not found by this search*, not *proven absent*. Closest adjacent art in the numerical literature: Hou’s dynamic-depletion diagnostics (alignment with strain eigenvectors, vortex-line regularity; qualitative) [9, 10], and Kerr’s numerical tests of direction-derived regularity conditions on antiparallel vortex tubes [13].

8 Reproducibility and discipline

```
python verify.py      # the backwards proof: 15 checks, 3 tiers
python mythos.py      # the confound engine
./done.sh             # the milestone ladder, M1-M5, all PASS
cat VERIFY_STATUS     # what the standing watcher currently holds
```

`verify.py` holds at 15/15 with EXACT-tier worst deviation 3.05×10^{-4} against references forced by the equations before any measurement existed; input checksums are in the verify output; a watcher reruns the full certification every twenty minutes and files a regression report if anything drifts. The laboratory was built by a human directing an AI system; the AI invented a branch that did not exist, fit through windows that moved, and glossed a mechanism its own certification refuted within the hour. Every failure was caught the same day by exact symmetry

anchors, context-free review, or the backwards proof, and the corpses are kept on the page. The discipline that survives contact with an AI that launders inference is part of what this repository demonstrates.

9 Open

The strip $-1/2 \leq \sigma \leq 0$: Corollary 0 forbids the top under viscosity at a growing maximum; the chain of Section 4 needs the bottom. Close the strip analytically and type-I Navier–Stokes blowup is finished for this route. The crossover amplitude $A_c(\nu)$ and the $\sigma_\Lambda(\nu)$ curve beyond the two certified viscosities are unmeasured (stream data exists at $\nu = 3 \times 10^{-4}, 3 \times 10^{-5}, 10^{-5}$, extending coverage one further decade; snapshots were not retained). The wandering sector remains the measured frontier: generic near-wall data dies too young at laptop resolution to read, and that failure is the recorded reason the sector is unexplored.

10 Tool and computational resource disclosure

This section follows the Leiden Declaration’s first recommendation [1] and the call to normalize responsible disclosure of AI assistance rather than conceal it.

Tools. Dedalus (spectral PDE solver) for the direct numerical simulations; NumPy/SciPy for all estimators; no proof assistant was used. An AI assistant (Anthropic Claude, used interactively throughout) wrote and debugged the solver configuration, the estimator and analysis code, and the corner-frame march; derived the analytical chain of Section 4; drafted this paper’s prose; and executed the prior-art search of Section ??.

Division of labour, stated plainly. The instrument code, the derivation in Section 4, 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 choice of observable and the decision to calibrate it on a case where blowup is proven; the epistemic protocol under which every result here was obtained (pre-registration of every estimator, bar, and hypothesis space *before* the data it judges existed; adversarial re-audit before publication; the standing refusals of Section 3); the decisions of what to measure, what to retract, and what to publish; and the repeated enforcement of scope against machine-generated first drafts. Several results in this paper exist *because* that protocol killed earlier machine-produced claims; those are recorded in Section ??.

What the author can and cannot defend at expert level. The author can defend: the observable’s construction and the scale invariance that forces the $-1/2$ power; the measurement protocol and its failure modes; why window-averaged slopes are regime mixtures; the estimator lineage, including which estimators died and why; and the retraction record. The author *cannot* defend at expert level the analytical steps of Section 4 — the Biot–Savart kernel split, the optimization over ρ , and the passage to the threshold exponents. Those were machine-derived under the author’s direction, are reported as a chain conditional on Hypothesis (H) with its measured constant, and have not been independently verified. Readers should treat Section 4 as a stated derivation awaiting expert review, not as established analysis. **The measurements of Sections 5 and 6 do not depend on it;** they stand or fall on the instrument and the data, both of which 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.

References

- [1] Leiden Declaration on Artificial Intelligence and Mathematics. 2 June 2026. doi:10.5281/zenodo.20302944. <https://leidendeclaration.ai>
- [2] P. Constantin, C. Fefferman. Direction of vorticity and the problem of global regularity for the Navier–Stokes equations. *Indiana Univ. Math. J.* **42**, 775–789 (1993).
- [3] P. Constantin. Geometric statistics in turbulence. *SIAM Review* **36**, 73–98 (1994).
- [4] H. Beirão da Veiga, L.C. Berselli. On the regularizing effect of the vorticity direction in incompressible viscous flows. *Differential Integral Equations* **15** (2002).
- [5] H. Beirão da Veiga, L.C. Berselli. Navier–Stokes equations: Green’s matrices, vorticity direction, and regularity up to the boundary (2009).
- [6] Y. Giga, H. Miura. On vorticity directions near singularities for the Navier–Stokes flows with infinite energy. *Comm. Math. Phys.* (2011). doi:10.1007/s00220-011-1197-x
- [7] T. Barker, C. Prange. Scale-invariant estimates and vorticity alignment for Navier–Stokes in the half-space with no-slip boundary conditions. *Arch. Ration. Mech. Anal.* **235**, 881–926 (2020). arXiv:1906.08225.
- [8] Z. Grujić. Logarithmic depletion of vortex stretching and singularity evasion in the 3D Navier–Stokes equations. arXiv:2607.08866 (2026).
- [9] T.Y. Hou. Blow-up or no blow-up? A unified computational and analytic approach to 3D incompressible Euler and Navier–Stokes equations. *Acta Numerica* (2009).
- [10] G. Luo, T.Y. Hou. Toward the finite-time blowup of the 3D axisymmetric Euler equations: a numerical investigation. *Multiscale Model. Simul.* (2014).
- [11] J. Chen, T.Y. Hou. Stable nearly self-similar blowup of the 2D Boussinesq and 3D Euler equations with smooth data I: Analysis. arXiv:2210.07191.
- [12] J. Chen, T.Y. Hou. Stable nearly self-similar blowup of the 2D Boussinesq and 3D Euler equations with smooth data II: Rigorous Numerics. *Multiscale Model. Simul.* **23**(1), 25–130 (2025). arXiv:2305.05660.
- [13] R.M. Kerr. Bounds for Euler from vorticity moments and line divergence. arXiv:1212.1106.
- [14] (Survey.) A survey of geometric constraints on the blowup of solutions of the Navier–Stokes equation. arXiv:2111.00040.