

The viscous inversion of the corner-flow blowup geometry fires at the collapse transient's turnaround, at every viscosity tested

Justin Hill
Independent Research
justin@epagoge.io

Draft 0.1 — 2026-08-02*

Abstract

Paper [P1] measured σ_Λ , the scaling exponent of the scale-invariant direction-regularity observable $\Lambda = \sup_P |\nabla \xi| |\omega|^{-1/2}$, on the Boussinesq corner-flow blowup scenario: $+1.00 \pm 0.03$ inviscid (calibrated where blowup is proven), inverting to a deep-collapse value near -1.2 under viscosity at $\nu \in \{10^{-4}, 10^{-3}\}$. This paper asks the questions that result forces: how the inversion depends on viscosity, where in the collapse it happens, and what sets that location. Within the tested scenario family (Section 1): (1) *No resolved viscosity dependence*. At the five tested viscosities $\nu \in \{10^{-5}, 3 \times 10^{-5}, 10^{-4}, 3 \times 10^{-4}, 10^{-3}\}$ the deep-collapse exponent shows no trend (chord estimator -1.12 to -1.29 , every cross-grid spread under the pre-registered 0.15 bar; two-slope-form asymptote -1.07 ± 0.03); the two previously certified viscosities replicated 4/4 against independent runs. Against the window-matched inviscid anchor ($+0.589/ +0.585$) the data show no approach toward the inviscid value down to $\nu = 10^{-5}$: any crossover in ν lies below the tested range. The data are consistent with a $\nu \rightarrow 0^+$ discontinuity of the exponent; they cannot distinguish it from a crossover below 10^{-5} . (2) *A fixed trigger*. The switch sits at $A^* = 52.7\text{--}54.2 \times A_0$ in twelve measurements at $\nu \in \{10^{-4}, 10^{-3}\} \times A_0 \in \{50, 100, 200\}$, both grids, and is bracketed inside the same band at the three other viscosities ($A_0 = 100$ only). Since the initial corner-gradient amplitude is $37.7 \times A_0$ (a geometric constant of this IC family), the trigger is a growth factor 1.43 (spread 1.40–1.44) above the initial value: about a third of an e-fold. The switch is nearly a step, completing within $\sim 0.2\text{--}0.3$ e-folds (measured at $\nu = 10^{-4}$ only). (3) *The two instruments meet*. An empirical clock fit ($\gamma_A = 0.77\text{--}1.11$) places the switch at $s^* = 0.52 \pm 0.05$ in rescaled time (per-run degeneracy widens this to $[0.29, 0.70]$), consistent with the corner-frame march's fast-transient turnaround at $s = 0.5$ (resolved to the $ds = 0.25$ step); the pre-registered lock-on hypothesis ($s \sim 3.7\text{--}13.5$) is refuted by more than a factor of five. A large-deviation march closes the linearity gap as far as the march reaches: the $s = 0.5$ turnaround is amplitude-invariant across 4.5 decades of perturbation amplitude, and the march's basin-or-instrument edge lies between amplitudes 3×10^{-2} and 10^{-1} along the one direction tested. The licensed mechanism statement, exactly bounded: viscosity is necessary for the inversion (inviscid runs never invert), and at the two viscosities where the clock was fit the inversion occurs at a rescaled time consistent with the transient's turnaround, well before self-similar settling; the switch location tracks a ν -independent, A_0 -linear stage of the collapse whose dynamical identity remains open.

*Sequel to [P1]. Every number regenerates from bytes on disk (Section 9). This paper was drafted after, not before, a two-lens verification pass over its own claim records: roughly 230 numeric claims recomputed independently, twelve scoping corrections applied, and four of our own estimators and readings killed en route. The corpses are in Section 8.

1 Scope, stated first

Everything in this paper is measured on one scenario family: the axisymmetric Boussinesq corner flow of the Luo–Hou type, quartic initial condition, corner scale $r_0 = 0.4$, **zpow** 1, free-slip viscosity in 5D-Laplacian form (the no-slip results of Barker–Prange [8] concern a different boundary condition and are cited as prior art for the criteria, not as the same setting). Amplitudes are in code units of the vector- $|\omega|$ maximum; grids are 128×384 and 256×768 (one refinement doubling; the 0.15 cross-grid bar tests consistency under that doubling, not asymptotic grid convergence); the march runs seed-0 perturbations at $ds = 0.25$ with the M2 stop rule, and every accepted march step converged the true nonlinear residual below 10^{-10} .

Novelty scoping. The prior-art search of [P1] (a one-session web sweep; “not found by this search,” never “proven absent”) covers the viscosity-resolved exponent curve: no measurement of a direction-regularity exponent as a function of viscosity was found, and Section 2 is that curve. The trigger invariant, switch width, clock comparison, and basin bound are reported as measurements with no novelty claim attached. Adjacent art sighted in a supplementary search: analytical viscous geometric depletion [9, 10], and nearly self-similar viscous blowup in *generalized* axisymmetric Navier–Stokes with dimension as a parameter [11]; our measurements concern the standard equations at fixed dimension, on this scenario family, at the viscosities stated.

2 The curve: no resolved ν -dependence

Ten new DNS runs (five viscosities, both grids, doubled snapshot cadence, fresh tags) went through the certified estimator of [P1] verbatim: signed circulation-drift gate, peak-box Λ recipe, cross-grid common-amplitude overlap, top half, symmetric midpoint cut, OLS slope with 10^4 -resample bootstrap intervals.

ν	128×384	256×768	spread
0 (anchor)	+0.589	+0.585	0.005
10^{-5}	−1.174	−1.190	0.016
3×10^{-5}	−1.271	−1.156	0.115
10^{-4}	−1.119	−1.187	0.068
3×10^{-4}	−1.286	−1.204	0.082
10^{-3}	−1.207	−1.231	0.024

Every spread passes the pre-registered 0.15 bar. The two viscosities certified in [P1] replicated 4/4 against these independent runs (worst $|\Delta| = 0.055$, within tolerance); the three new viscosities are single run-pairs with cross-grid agreement as their check. The deep windows carry $n = 20$ –27 snapshots (the doubled cadence fixed [P1]’s thinnest tables). A two-slope-form fit (Section 3) puts the asymptotic deep slope at -1.07 ± 0.03 across all six fitted runs, systematically shallower than the chord values; both are quoted and neither replaces the other. The chord-vs-form gap (~ 0.1) is the current estimator systematic on the exponent’s absolute value; the ν -flatness conclusion is insensitive to it.

What the table does not show: any approach toward the inviscid anchor. If $\sigma_\Lambda^{\text{deep}}(\nu)$ returns to +0.59 as $\nu \rightarrow 0$, it does so below $\nu = 10^{-5}$. Within the tested range the inversion is viscosity-blind. The Corollary-0 reading of [P1] (the viscous cap on direction-gradient growth exists for

every $\nu > 0$ and is absent at $\nu = 0$ [4]) is consistent with a genuine 0^+ discontinuity; that reading is interpretation, and the measurement alone cannot exclude a crossover below the tested range.

3 The trigger: a fixed growth factor, and a near-step switch

At finite ν , $\sigma(A)$ is not one number: every viscous run crosses from inviscid-like slope (+0.39 to +0.81) at low amplitude to the inverted value at high amplitude. [P1] adjudicated window-averaged slopes as regime mixtures; here the crossover itself becomes the observable.

The estimator lineage is part of the record. A pre-registered chord-intersection estimator for the crossover amplitude *failed* its split-invariance test (15–42% drift tracking the window split, monotone in nine of ten runs) and its outputs are void: $\sigma(A)$ curves smoothly rather than breaking, so a two-chord intersection follows the split point. Its successor, the first zero crossing of a k -row sliding-window local slope, passes the same class of test that killed its predecessor: k -drift 0.9–2.1% on the discovery runs and 0.7–1.0% on the cross-term runs, cross-grid agreement under 1%.

By that estimator, the crossover amplitude is A_0 -linear with measured power +1.00 (ratios 0.50/0.49 at $A_0 = 50$ and 2.01/2.00 at $A_0 = 200$ against the amplitude-symmetry prediction), and the twelve measurements at $\nu \in \{10^{-4}, 10^{-3}\} \times A_0 \in \{50, 100, 200\}$ give $A^*/A_0 = 52.7$ –54.2. At $\nu \in \{3 \times 10^{-4}, 3 \times 10^{-5}, 10^{-5}\}$ ($A_0 = 100$ only) the crossover is bracketed inside the same fixed band: the lower half of every overlap window reads inviscid-positive and the upper half reads inverted, so the crossover cannot have moved with ν within the tested window. The tested region of the (ν, A_0) plane is a cross, not a tile; the plane statement is exact for that cross.

The initial corner-gradient amplitude is $37.7 \times A_0$ in every run, a geometric constant of the quartic IC family (early inviscid snapshots grow $37.71 \rightarrow 37.80 \rightarrow 37.93$: the value at the first trusted snapshot is essentially the initial value). The trigger is therefore a growth factor of 1.43 above the initial amplitude (twelve-run spread 1.40–1.44; the tanh-form estimator places the crossing systematically $\sim 8\%$ lower, so the absolute factor is estimator-conditional at that level while A_0 -linearity and ν -independence are estimator-robust). Three dependencies are declared: the factor is a property of this IC family, of the vector- $|\omega|$ observable, and of the zero-crossing estimator.

The switch is nearly a step. A two-slope blend fit is residual-adequate on all six runs at $\nu = 10^{-4}$ but parameter-degenerate (the optimizer drives it into an exponential-transient limit; its center, width, and low- A slope are not individually physical and are not quoted). The robust extract is the transient decay scale: the local slope completes its transition from 0 to the inverted value within ~ 0.2 –0.3 e-folds of amplitude. The width is measured at $\nu = 10^{-4}$ only.

One confound was tested and killed before these claims were written: across the switch, the $|\omega|$ argmax stays at the corner ($z^* = 0$, $r^* = 1.000$) at every gated snapshot, and the observable tracks the swirl-gradient structure throughout (ω_θ vanishes at the corner by symmetry). The diagnostic was run at $\nu = 10^{-4}$ on the 128-grid and on the matched inviscid run; we treat it as excluding a structure handoff for the runs where the switch is certified. The inversion is a change in the direction-field geometry at one fixed structure, not the instrument changing targets.

4 The clock: two instruments meet at the transient’s turnaround

[P1]’s rescaled-frame march certified the corner profile as an attractor: perturbations peak at $s = 0.5$ (raw factor 2.389 at linear amplitude) and contract at 0.270 per s -unit. The DNS switch

can be placed on the same s -axis by fitting the collapse clock $\ln A = -\gamma_A \ln(1 - t/T^*) + \ln A_i$ per run, T^* scanned, with $s = -\ln(1 - t/T^*)$.

run	γ_A	s^*	s^* over $2\times$ -residual T^* set
OR 128×384 ($\nu = 0$)	0.77	0.61	[0.42, 0.70]
OR 256×768 ($\nu = 0$)	0.94	0.54	[0.41, 0.58]
C 128×384 (10^{-4})	0.86	0.55	[0.38, 0.61]
C 256×768 (10^{-4})	1.11	0.48	[0.35, 0.51]
C 128×384 (10^{-3})	1.04	0.47	[0.30, 0.53]
C 256×768 (10^{-3})	1.03	0.49	[0.31, 0.55]

$s^* = 0.52 \pm 0.05$ (run-to-run scatter; the $T^*-\gamma_A$ degeneracy widens the honest per-run interval to $[0.29, 0.70]$), consistent with the march turnaround at $s = 0.5$, itself resolved only to the $ds = 0.25$ step. The value is conditional on the clock form and on the $54 \times A_0$ switch input: the tanh-implied $49\text{--}50 \times A_0$ shifts s^* by about -0.09 , inside the degeneracy band. γ_A spans $0.77\text{--}1.11$, consistent with the corner frame’s one-e-folding-per- s -unit and measured rather than assumed. The pre-registered hypothesis space put lock-on at $s \sim 3.7\text{--}13.5$; even the widest degeneracy bound clears that window by more than a factor of five. Within the pre-registered space and clock family, the lock-on hypothesis is refuted: the switch does not wait for the collapse to settle onto the profile.

The comparison identifies matching timescales of one operator, not coincident trajectories: the march transient is measured for perturbations about the settled profile, while the DNS approach is a large deviation. The next section closes that gap as far as the march can reach.

5 The large-deviation march

The certified march ran the same seed-0 perturbation at amplitudes 3×10^{-3} , 10^{-2} , 3×10^{-2} , and 10^{-1} (the largest previously validated amplitude was 10^{-3}). At the first three, all 240 steps are valid and the raw-norm peak sits at $s_{\text{peak}} = 0.5$ exactly, at step resolution. Combined with [P1]’s M2 ladder, the transient turnaround is amplitude-invariant across 4.5 decades of relative deviation (10^{-6} to 3×10^{-2}). The $s = 0.5$ timescale that the DNS switch lands on is a property of the operator’s nonlinear neighborhood, not merely of its linearization.

Around the pinned turnaround, genuine nonlinear structure appears for the first time in this campaign: the transient’s end stretches ($s_{\text{transient}} = 13.5 \rightarrow 14.0 \rightarrow 14.75$, +9%) and the peak factor grows ($2.4068 \rightarrow 2.4504 \rightarrow 2.5884$ within this ladder; M2’s linear value is 2.389) while the peak position does not move. The Lyapunov P -norm contracts monotonically from $V_0 = 77$ through 13.7 decades with zero violations above the pre-registered V -floor artifact threshold: the certificate’s contraction holds at 3% deviation.

At amplitude 10^{-1} the march stalls at step one (raw growth $3.3\times$ in one step; the P -norm nevertheless decreased on that step, a mixed signal). The pre-registered protocol does not separate dynamical escape from solver stall, so the honest statement is a basin-or-instrument edge: along the seed-0 direction, $3 \times 10^{-2} < r < 10^{-1}$ in the march’s norm. This also bounds the remaining daylight in Section 4’s comparison: the DNS initial deviation exceeds this basin, and its approach enters the basin during the transient. Full closure would need DNS-state initialization of the march, which is a different instrument build.

6 Mechanism, stated within bounds

For every viscosity tested, the inversion is present only with $\nu > 0$: inviscid runs on identical amplitude windows never invert. At the two viscosities where the clock was fit, the switch occurs at a rescaled time consistent with the corner-frame transient’s turnaround, well before self-similar settling, and the turnaround timescale is amplitude-robust across every amplitude the march can hold. Viscosity is necessary for the inversion; the location of the switch tracks a ν -independent, A_0 -linear stage of the collapse. What sets that stage dynamically is open, and Section 2’s gap below $\nu = 10^{-5}$ is open with it. The strongest reading consistent with all measurements: the geometric depletion that [P1] measured deep in the collapse is already fully switched on by the end of the approach transient, at every viscosity and amplitude tested, and the deep-collapse exponent it switches to does not depend on how much viscosity did the switching.

7 Relation to prior art

The criteria territory is [P1]’s: Constantin–Fefferman [3], Beirão da Veiga–Berselli [5, 6] at the definition of Λ , Giga–Miura [7] and Barker–Prange [8] for qualitative type-I exclusion under direction coherence, and Constantin’s identity [4] behind the Corollary-0 diagnostic. Analytical viscous geometric depletion is developed in [9] and, through logarithmically weighted direction classes, in [10]; neither defines a measured exponent or an onset. Nearly self-similar viscous blowup is reported in *generalized* axisymmetric Navier–Stokes with dimension as a parameter [11]; the present measurements concern the standard equations on this scenario family and are not in tension with results for modified systems. The measured objects of this paper (an exponent-vs-viscosity curve, a trigger invariant, a cross-instrument clock comparison) were not found in the [P1] search; only the curve carries a novelty claim, under that search’s stated horizon.

8 What died on the way

Four of our own results were killed by our own instruments during this campaign, and the corpses are retained:

- The chord-intersection crossover estimator: failed split-invariance (15–42% drift); all its outputs, including a fitted crossover-vs-viscosity slope, are void.
- The tanh-form parameters: residual-adequate fit, degenerate parameters; only the decay scale and asymptote survive.
- The “four e-folds of dynamical age” reading of the trigger: killed by the normalization audit (the initial corner-gradient amplitude is $37.7 \times A_0$, not A_0); the honest trigger is $1.43\times$ above the initial value, essentially at collapse onset.
- The lock-on bridge hypothesis: the clocks were compared and it lost, by a factor of five beyond the widest degeneracy bound.

[P1]’s standing retractions (window-averaged viscous slopes; the amplitude-composite reading of the full-window inviscid calibration) are inherited: no full-window σ is quoted anywhere

in this paper, and every viscous contrast is drawn against the window-matched inviscid +0.589, never against +1.00.

9 Reproducibility

The DNS campaign is 18 runs (about 90 minutes of laptop wall time) plus the march ladders (~ 90 seconds each), all regenerable: run scripts, pre-registration (every estimator, bar, and hypothesis space in this paper was written down before the data it judges existed; the spec file’s section dates interleave with the run logs), extraction code, and per-claim records with their adjudication addenda are in the repository alongside [P1]’s verification suite (15/15 standing, watcher live). The two-lens audit that preceded this draft (a numbers tracer recomputing ~ 230 claims from the raw arrays and snapshots; an adversarial referee producing the twelve scoping corrections now baked into the text) is itself part of the record.

10 Tool and computational resource disclosure

As in [P1], and following the Leiden Declaration [1]:

Tools. Dedalus for the simulations, NumPy/SciPy for the estimators, no proof assistant. An AI assistant (Anthropic Claude, used interactively) wrote the run drivers, the tracking and estimator code, and this paper’s prose, and executed the supplementary prior-art search of Section 7.

Division of labour. The code and prose are machine-assisted; the protocol is the author’s, and in this paper the protocol is most of the method: every estimator, bar, and hypothesis space in Sections 2–5 was written down before the data it judges existed (the specification file’s section timestamps interleave with the run logs), and two pre-registered routes were stopped by their own tests rather than rescued. The four corpses in Section 8 are machine-produced results that the author’s protocol killed.

What the author can and cannot defend at expert level. This paper is almost entirely measurement, and its exposure is correspondingly narrower than [P1]’s. The author can defend the estimators, their invariance tests, the pre-registration record, and every retraction. The one inherited analytical dependency is [P1]’s exclusion chain, which is cited here for context and carries the same caveat stated in [P1]’s disclosure: it is machine-derived, conditional on a structural hypothesis with a measured constant, and unverified. No measurement in this paper depends on it.

Responsibility. Rests with the human author, including for the machine-assisted portions.

References

- [1] Leiden Declaration on Artificial Intelligence and Mathematics. 2 June 2026. doi:10.5281/zenodo.20302944. <https://leidendeclaration.ai>
- [2] J. Hill. Viscosity inverts the direction-regularity scaling of the corner-flow blowup mechanism: a measured exponent crosses the type-I exclusion threshold. 2026. <https://epagoge.github.io/parzival/sigma/> (paper, LaTeX, data, and claims ledger; repository EPA-GOGE/parzival).

- [3] 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).
- [4] P. Constantin. Geometric statistics in turbulence. *SIAM Review* **36**, 73–98 (1994).
- [5] H. Beirão da Veiga, L.C. Berselli. On the regularizing effect of the vorticity direction in incompressible viscous flows. *Differential Integral Equations* **15** (2002).
- [6] H. Beirão da Veiga, L.C. Berselli. Navier–Stokes equations: Green’s matrices, vorticity direction, and regularity up to the boundary (2009).
- [7] Y. Giga, H. Miura. On vorticity directions near singularities for the Navier–Stokes flows with infinite energy. *Comm. Math. Phys.* (2011).
- [8] 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.
- [9] N. Ju. Geometric depletion of vortex stretch in 3D viscous incompressible flow. *J. Math. Anal. Appl.* **321**(1), 412–425 (2006). doi:10.1016/j.jmaa.2005.08.048.
- [10] Z. Grujić. Logarithmic depletion of vortex stretching and singularity evasion in the 3D Navier–Stokes equations. arXiv:2607.08866 (2026).
- [11] T.Y. Hou. Nearly self-similar blowup of generalized axisymmetric Navier–Stokes equations. arXiv:2405.10916 (2024, revised 2025). [Verified 2026-08-02: generalized equations, dimension as parameter; preprint.]