A probabilistic illustration of the non-uniqueness mechanism

Forced three-dimensional Navier-Stokes, a compact ring and self-similar variables

Main result. In the simulations performed, the final separation between the branches remains close to 295 while their initial physical distance falls as low as 1.52×10^{-45}. The signal persists as the number of samples increases. These are nonlinear, three-dimensional axisymmetric computations on a truncated domain, not two certified Leray solutions on $\mathbb{R}^3$.

1. An explicit profile chosen for this experiment

Set $r = \sqrt{(x_1^2 + x_2^2)}$, $z = x_3$ and $\rho^2 = (r - 12)^2 + z^2$. We define

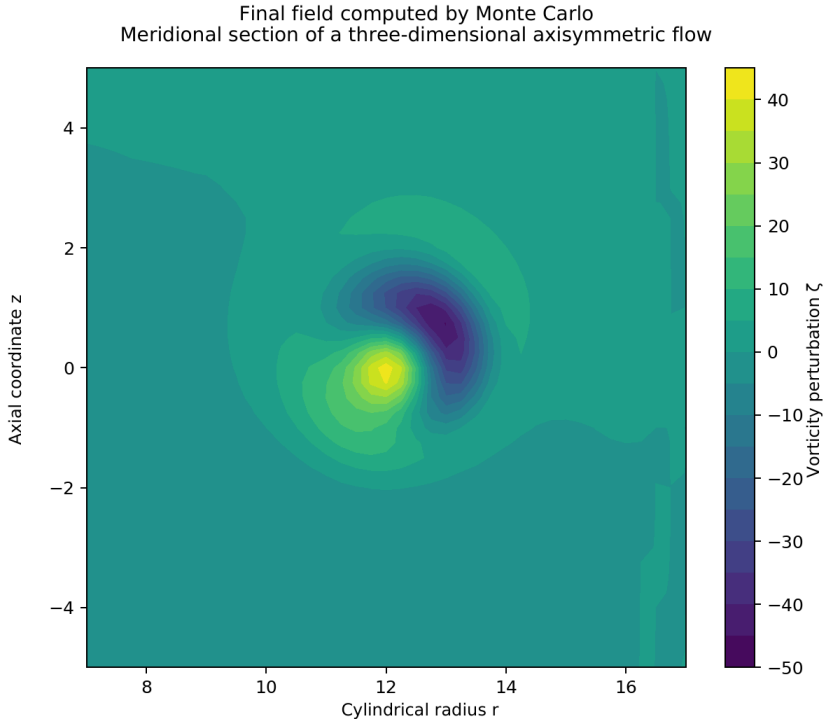
$$\bar{\Phi}(r,z) = 200 e^{-\rho^2/2}\chi(\rho), \qquad \bar{U} = -\partial_z\bar{\Phi}\, e_r + \left(\partial_r\bar{\Phi} + \frac{\bar{\Phi}}{r}\right)e_z.$$

The smooth cutoff χ equals 1 for $\rho \le 3.5$ and 0 for $\rho \ge 4.5$. Its full formula is given in the code and README. The profile is smooth, compactly supported, divergence-free and zero near the axis. It is not the Vishik profile from the theoretical manuscript: it is a candidate whose instability is tested numerically.

We choose the common forcing

$$\bar{F} = -\tfrac{1}{2}(\bar{U} + \xi\cdot\nabla\bar{U}) - \Delta\bar{U} + (\bar{U}\cdot\nabla)\bar{U},$$

which makes U stationary in self-similar variables, with viscosity $\nu = 1$.



Final perturbed field computed by Monte Carlo, in a meridional section. Dependence on (r, z) represents a three-dimensional axisymmetric flow with no azimuthal component.

2. The equation and probabilistic method actually used

We use the convention $\zeta = \partial_r U_z - \partial_z U_r = -(\nabla \times U)_\theta$, and set $q = \zeta - \bar{\zeta}$. The perturbation velocity v is reconstructed from

$$D\Phi_q = q, \qquad D = \partial_r^2 + \tfrac{1}{r}\partial_r + \partial_z^2 - \tfrac{1}{r^2}, \qquad v = \left(-\partial_z \Phi_q, \partial_r \Phi_q + \tfrac{\Phi_q}{r}\right).$$

The solver integrates every term in

$$\partial_\tau q = Dq + q + \tfrac{1}{2}(r\partial_r + z\partial_z)q - \bar{U}\cdot\nabla q + \tfrac{\bar{U}_r}{r}q$$
$$- v\cdot\nabla\bar{\zeta} + \tfrac{v_r}{r}\bar{\zeta} - v\cdot\nabla q + \tfrac{v_r}{r}q.$$

The forcing is identical for both branches and cancels algebraically in this difference equation. The geometric terms in 1/r and $1/r^2$ and the weight $2\pi r$ are retained. The stream function enforces the divergence constraint.

Discretization. The main domain is $0 < r < 24$, $|z| < 12$, with a grid of 95×95 interior points, second-order centred differences, a sparse LU factorization for D, Strang splitting and RK4. The outer conditions $q = \Phi_q = 0$ are artificial; their influence is tested by enlarging the domain. At the axis they are compatible with regularity of the variables. Velocity is not prescribed to vanish on the outer boundary.

The Monte Carlo component

For the discrete Cartesian Laplacian $\Delta_h = \partial_{rr,h} + \partial_{zz,h}$, the Dirichlet frequencies are

$$\kappa_j = \frac{2}{h}\sin\frac{j\pi}{2(n+1)}.$$

A diffusion step of duration s uses the exact identity

$$e^{-s(\kappa_j^2+\kappa_k^2)} = \mathbb{E}\left[\cos(\sqrt{2s}\,\kappa_j Z_1)\cos(\sqrt{2s}\,\kappa_k Z_2)\right], \quad (Z_1, Z_2) \sim N(0, I_2).$$

The expectation is replaced by m independent Gaussian vectors at each half-step. The product of cosines incorporates four antithetic reflections. The geometric diffusion correction belongs to the RK4 part.

This is Gaussian quadrature of the discrete heat semigroup, not a fully simulated three-dimensional Lagrangian flow. No noise is added to the physical equation. Each diffusion half-step is unbiased for discrete diffusion; the complete nonlinear evolution need not be unbiased.

Initialization and measurement

We compute $L_h \eta_h = \lambda_h \eta_h$, normalized by $\|v[\mathrm{Re}\ \eta_h]\|_{L^2} = 1$, and then set

$$q(\tau_0) = 300\,\mathrm{Re}\left(e^{\lambda_h \tau_0}\eta_h\right), \qquad \tau_0 = -T.$$

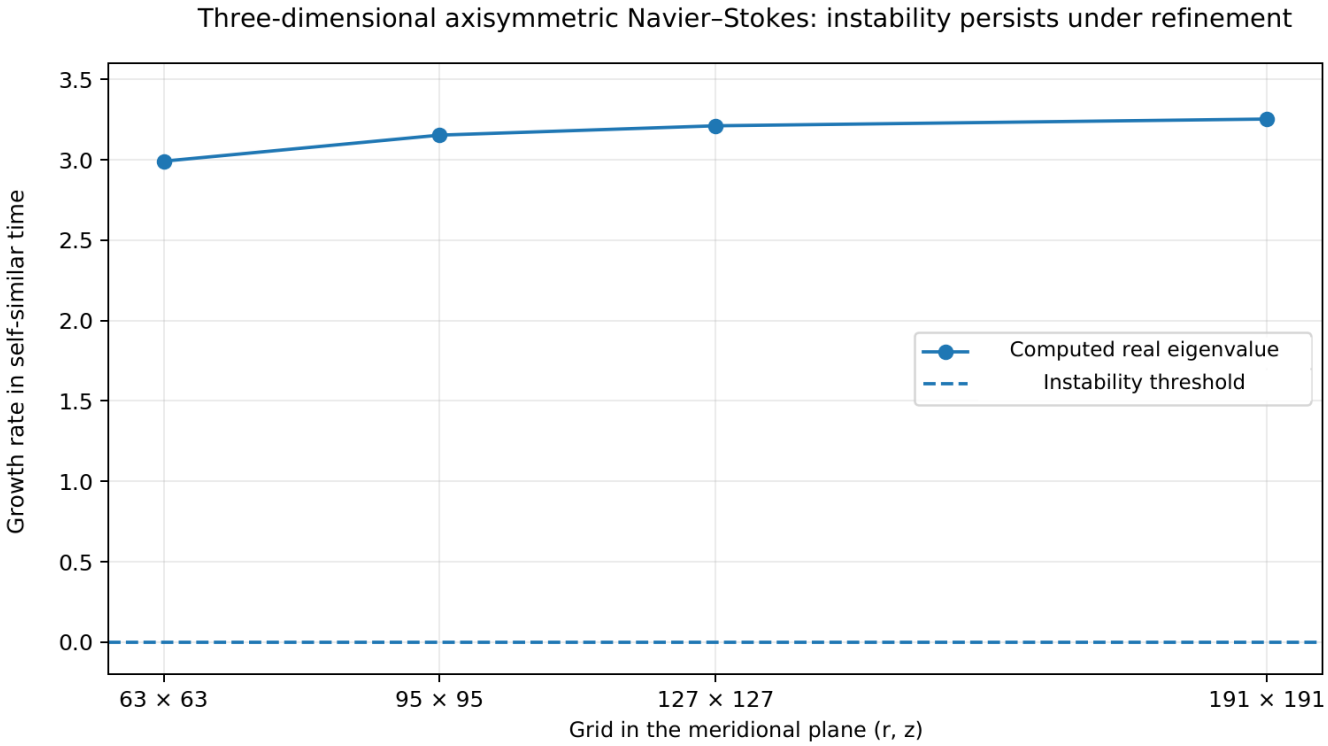
The first branch is $q = 0$. No renormalization is applied during evolution. The second-order nonlinear initial correction is not added; instead, the effect of moving τ_0 backwards is tested. The perturbation is stored separately from the large base profile to avoid loss of significant digits.

3. Instability that survives the numerical checks

The spectral search found a positive real mode that grows faster than the two complex modes examined. The following results concern the main domain. They do not certify the entire spectrum.

Grid	Real λ_h	$\|v(0)\|$	Field difference / 191^2
63 × 63	2.991298	294.6295	7.23%
95 × 95	3.153231	294.9132	2.67%
127 × 127	3.211328	295.0006	1.10%
191 × 191	3.253313	295.0587	0.00%

Nonlinear evolutions from $\tau_0 = -4$ with time step 0.001. Each grid uses its own normalized eigenmode. The field difference is a distance between velocity fields, not merely a difference of norms. The fine reference is not certified.



The positive eigenvalue does not disappear under refinement. When the domain is enlarged at fixed mesh size h = 0.25, the ring centre remains r = 12:

Domain in (r, z)	Real λ_h	Field difference / largest domain
[0, 24] × [−12, 12]	3.153231	1.270%
[0, 32] × [−16, 16]	3.174099	0.498%
[0, 40] × [−20, 20]	3.177716	0.000%

Evolutions from $\tau_0 = -4$ with time step 0.002. Comparisons use the common points of the domains. They do not control the tail of the field at infinity.

The growth rate measured in the linear control is 3.153226834, compared with 3.153231226 in the main spectral calculation. The time evolution therefore recovers the spectral instability.

4. The central test: approaching zero initial time

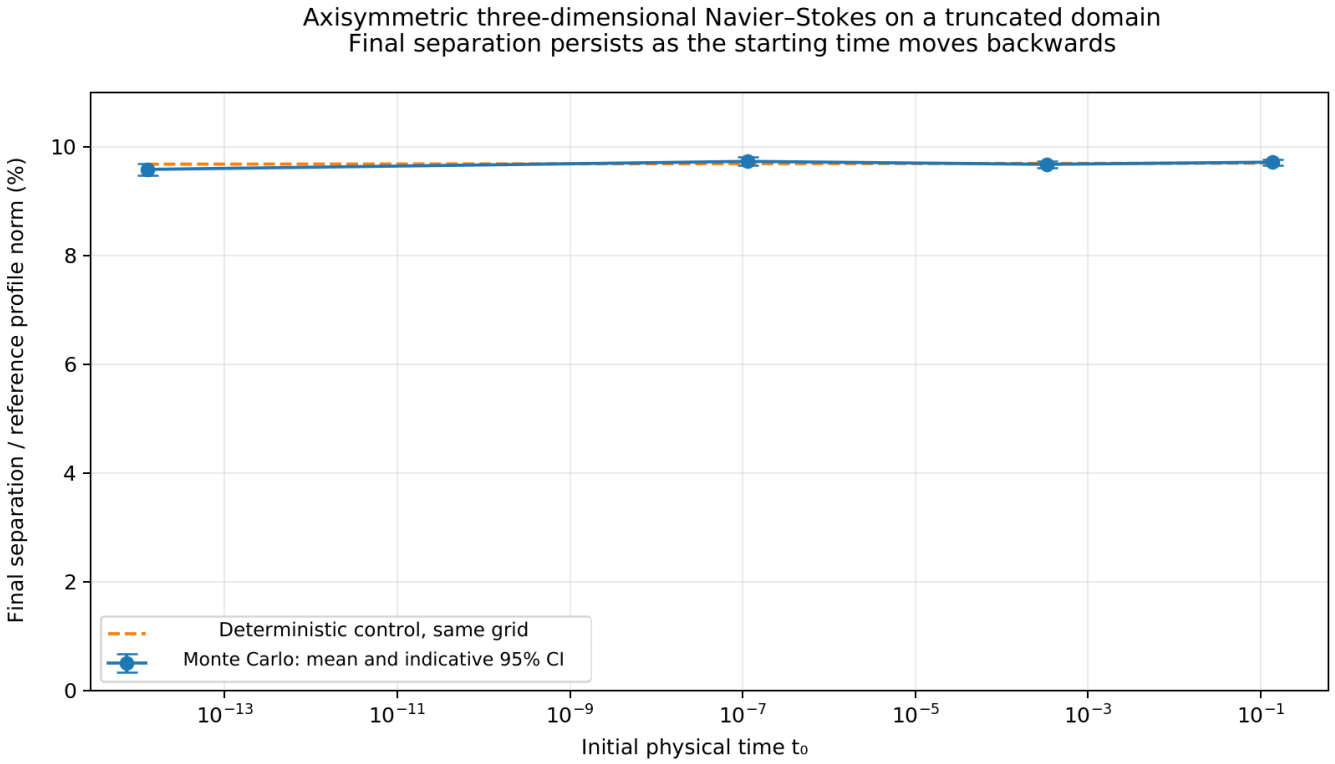
In physical variables,

$$t = e^{\tau}, \qquad u(t,x) = t^{-1/2} U\Big(\log t, \frac{x}{\sqrt{t}}\Big), \qquad \|u(t)\|_2 = t^{1/4} \|U(\log t)\|_2 .$$

Thus $\|u(t_0)\|_2 = e^{-T/4}\|U\|_2 \to 0$. The second initial datum is even closer to the first. The final time is $T = 1$ and $\|U\|_2 = 3041.531$.

τ_0	t_0	Initial distance	Final MC distance	95% CI
-2	1.35×10^{-1}	3.32×10^{-1}	295.46	[293.93; 296.98]
-8	3.35×10^{-4}	4.50×10^{-10}	294.24	[292.50; 295.97]
-16	1.13×10^{-7}	6.75×10^{-22}	295.88	[293.57; 298.19]
-32	1.27×10^{-14}	1.52×10^{-45}	291.39	[288.08; 294.69]

Physical L^2 norms, using the weight $2\pi r$. Means of 12, 12, 8 and 6 repetitions respectively, with $m = 32$ Gaussian vectors per half-step. The Student intervals are indicative and concern only the mean across repetitions.



Observation. The final separation is about 9.6-9.7% of the reference profile norm; it does not collapse as the starting time moves backwards. At $\tau_0 = -32$, the initial profile still has norm 1.02032, or 0.03355% of its final norm: the data are not exactly zero at this finite time.

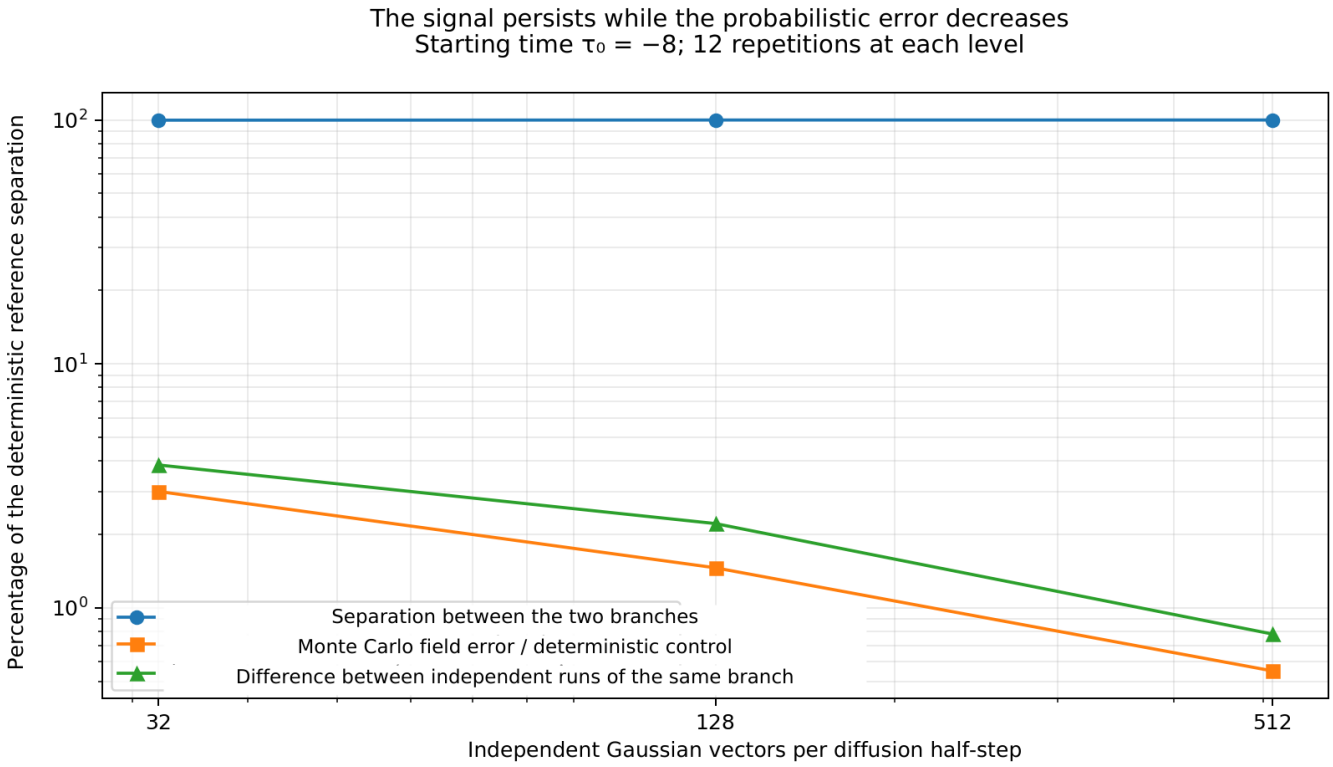
Positive initial times correspond to slightly different Cauchy problems. The purpose of the experiment is to examine their behaviour as $t_0 \to 0$, as in the self-similar mechanism of the theoretical manuscript. A finite collection of simulations does not prove this limit.

5. The signal is distinct from Monte Carlo error

For $\tau_0 = -8$, increase the number m of independent Gaussian vectors per half-step. Each row uses 12 realizations.

m	Mean final norm	Field RMSE	Same-branch difference
32	294.236	3.000%	3.851%
128	294.675	1.458%	2.214%
512	294.721	0.551%	0.780%

Errors relative to the deterministic perturbation field on the same grid with the same time step. The same-branch difference uses six disjoint pairs of independent realizations.



With m = 512, the mean result and its indicative interval are

$$\widehat{D} = 294.7207, \qquad I_{95\%} = [294.3236;\ 295.1179].$$

The deterministic control is 294.7555.

The field RMSE falls from 3.000% to 0.551%. The difference between two calculations of the same branch falls from 3.851% to 0.780% of the reference separation, while the distance between branches remains close to 100% of that reference.

Caution. A small Monte Carlo error does not remove spatial or domain errors. These intervals are not total-error bounds or certificates of non-uniqueness. The deterministic reference serves as a numerical control, not as an exact truth.

6. Additional checks and a discarded trial

Time refinement

Step Δτ	Final norm, T = 8	Field difference / step 0.001
0.004	294.755542	0.05440%
0.002	294.877923	0.01090%
0.001	294.908597	0.00000%

On the main grid, the field difference between time steps 0.004 and 0.001 is 0.0544%. It is small compared with the observed separation, without providing a rigorous bound on time-discretization bias.

The nonlinear terms genuinely matter

Comparing two evolutions with time step 0.001, one linearized and one complete, gives

$$\frac{\|v_{\mathrm{nonlin}}(0) - v_{\mathrm{lin}}(0)\|_2}{\|v_{\mathrm{nonlin}}(0)\|_2} = 10.908\%.$$

The difference of their norms alone is smaller; it does not describe the full change in the field. The result is therefore not a prescribed exponential multiplied by an eigenfunction.

Zero and geometric controls

Four Monte Carlo evolutions initialized with q = 0 remain exactly zero. The perturbation formulation preserves the stationary profile by construction. The measured discrete divergence has RMS of order 2×10^{-16}. This precision follows in particular from the stream function; it does not certify the overall accuracy of the Navier-Stokes solution.

An initial choice of mode was unsuitable

A complex mode with real part near 1.675 was used initially. The separation did not stabilize as the starting time moved backwards: the calculation reached about 3538 at T = 16. A wider search then found the faster real mode, near 3.153.

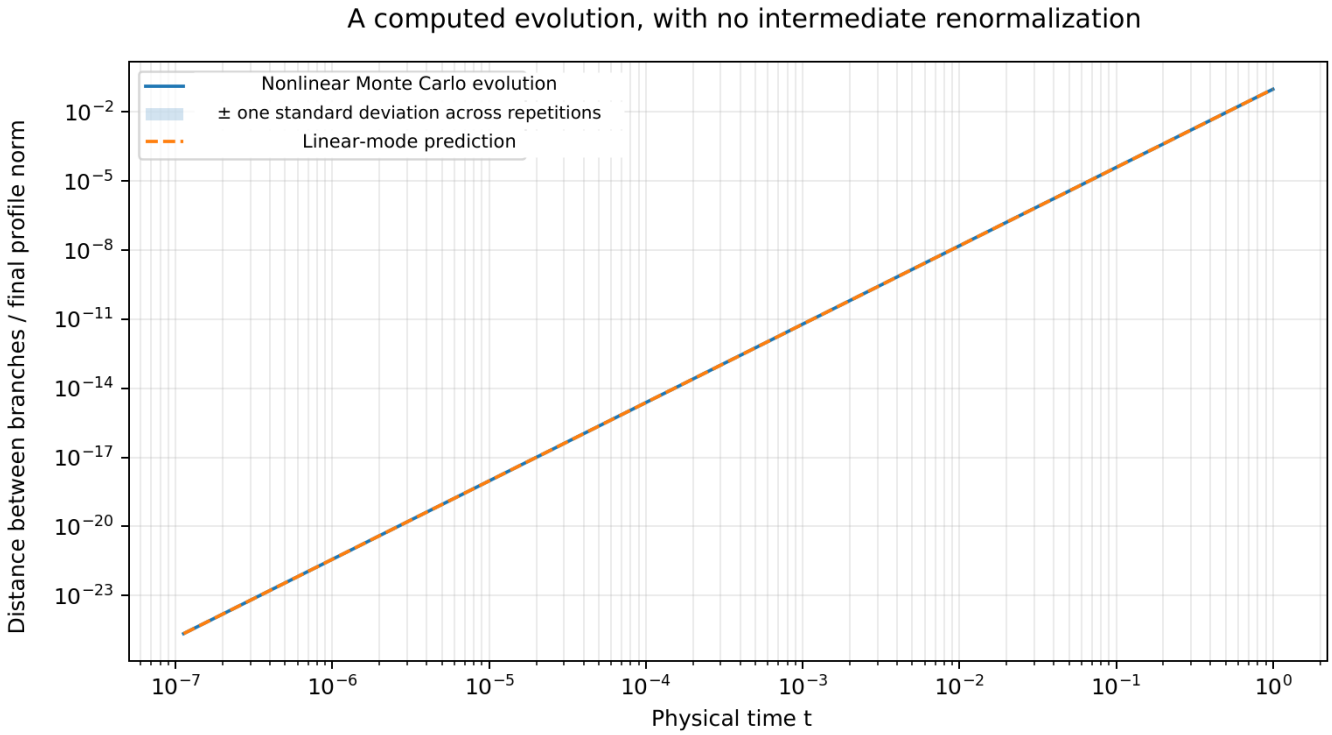
Initialization with a slower mode is sensitive to contamination by faster modes, including contamination from numerical errors. The campaign was therefore repeated with the real mode. The outputs of the inconclusive trial were retained in the original campaign directory `diagnostic_nonleading_mode`. They are not used as evidence for non-uniqueness.

Summary of the checks. On the main grid, the observed field differences are 2.67% relative to the fine grid and 1.27% relative to the largest domain on the common region. They support qualitative robustness of the signal. They should not be added together and presented as a certified bound on the total error.

7. What this experiment supports

A defensible statement: the computation gives a probabilistic numerical illustration of the self-similar non-uniqueness mechanism in a discretization of the forced three-dimensional equations. A branch selected by an unstable mode remains separated from the reference profile at a fixed time, while the initial physical data approach zero. The separation persists under refinement and substantially exceeds the measured statistical fluctuations.

What is not established: two distinct Leray solutions on all of $\mathbb{R}^3$, rigorous validation of the continuous spectrum for this particular profile, convergence of all the numerical limits, or the energy inequality for the limits. The profile used is not Vishik's exact construction. Initialization selects the branches; the noise does not discover them spontaneously.



Evolution of the physical distance in a nonlinear calculation from $\tau_0 = -16$. The linear law is superimposed as a control; it is not used to replace the evolution.

Reproduction and traceability

The original report describes a campaign directory containing `ring_solver.py`, the campaigns, seeds, JSON and NPZ outputs, plots, an integrity manifest and a detailed README. Python ≥ 3.10, NumPy, SciPy and Matplotlib suffice for those calculations. Its main commands are `python campaign.py det` and `python campaign.py mc`; further controls are listed in that README. The present website distributes its runnable implementation as `navier_stokes.py`.

Sources and contributions of this campaign

[1] T. Giarrizzi and Y. Wang, *Non-uniqueness of Turbulent Solutions of the Navier-Stokes Equation in Dimension N=3*, 15 February 2025, supplied theoretical manuscript, Sections 4-7.

[2] D. Albritton, E. Brué and M. Colombo, *Non-uniqueness of Leray solutions of the forced Navier-Stokes equations*, Annals of Mathematics 196 (2022), 415-455. DOI: 10.4007/annals.2022.196.1.3.

The theoretical mechanism comes from these sources. The explicit profile, computational method and tabulated results belong to this experiment. The signs are those of the self-similar equation, checked against the original article. No theoretical novelty is claimed.