

The viscous scale as a test of blowup constructions

Justin Hill
EPAGOGUE LLC

8 September 2026

Status: scaling analysis, not a theorem. The criterion below is a corollary of two classical facts. Its application to each construction uses the standard linear estimate that a mode of wavenumber ξ dissipates at rate $\nu\|\xi\|^2$, and every constant is read from the published Lean sources. Nothing here verifies or refutes either proof. Prepared with Claude (Anthropic) under my direction; see Section 7.

Summary

Two blowup constructions for incompressible fluid equations were published within a day of each other, both formalized in Lean. Alpöge and Buckmaster prove finite-time blowup with smooth forcing for the 2D Boussinesq system and 3D Euler. OpenAI claims alternatives (C) and (D) of the Navier–Stokes Millennium problem. Tao’s exposition of the first notes that those methods have “high likelihood of also extending to Navier–Stokes,” which is stated as expectation rather than result.

Section 1 derives a necessary condition any viscous blowup must satisfy: its active wavenumber is pinned to $(\nu(T - t))^{-1/2}$, the viscous scale. Sections 2 and 3 test both constructions against it.

Construction	ratio to critical scaling
Alpöge–Buckmaster ladder, stage 1	e^{17551}
OpenAI <code>ChartScales</code> , $n = 10 \dots 10^6$	0.24 to 2.57
Certified Chen–Hou profile (Section 4)	$d \log k / d \log \nu = -0.037$

The first two are separated by twenty thousand orders of magnitude in the log. The third is an out-of-sample check on data neither construction has seen, where the criterion correctly reads the Euler scale rather than the viscous one. The criterion discriminates, which is the only reason it is worth stating. A test everything fails measures nothing.

The reading is that the frequency cascade route to Navier–Stokes is barred at its base, and that the construction claiming to have reached Navier–Stokes did not take that route. Those are the same statement seen twice.

1. The criterion

Leray's lower bound

Navier–Stokes is invariant under $u_\lambda(x, t) = \lambda u(\lambda x, \lambda^2 t)$ at fixed ν . Local existence continues smoothly from data of size M for a time at least $c\nu/M^2$, the only combination of ν and M carrying units of time. If u first loses smoothness at T , restarting at any $t < T$ must not reach past T :

$$t + \frac{c\nu}{\|u(t)\|_\infty^2} \leq T \quad \implies \quad \|u(t)\|_\infty \geq c \left(\frac{\nu}{T-t} \right)^{1/2}. \quad (1)$$

The rigorous route runs through subcritical L^p , $p > 3$; the scaling is the reason.

Scale selection

A structure carrying velocity $U \sim (\nu/(T-t))^{1/2}$ at length L has nonlinear rate U/L and viscous rate ν/L^2 . These coincide exactly at

$$L_\nu(t) = (\nu(T-t))^{1/2}, \quad \text{where both rates equal } \frac{1}{T-t}. \quad (2)$$

Above L_ν dissipation is negligible but the nonlinearity is too weak to finish in the time remaining. Below it, dissipation dominates. Blowup is not merely permitted at L_ν ; by (1) it is forced there.

The enstrophy budget

From the energy identity,

$$\nu \int_0^T \|\nabla u\|_{L^2}^2 dt \leq \frac{1}{2} \|u_0\|_{L^2}^2 + \int_0^T \int f \cdot u =: C_0 < \infty, \quad (3)$$

the work term being finite for any Fefferman-admissible f : condition (5) gives decay faster than any polynomial in t , and Grönwall bounds $\|u\|_{L^2}$. This is a hard budget and it survives the presence of forcing.

Criterion 1. *If blowup is carried by structures of amplitude a_n , wavenumber k_n , support volume V_n , held for durations τ_n , then $\|\nabla u\|^2 \sim a_n^2 k_n^2 V_n$ and*

$$\sum_n a_n^2 k_n^2 V_n \tau_n \leq \frac{C_0}{\nu}. \quad (4)$$

The sum must converge, so durations shrink at least as fast as k_n^{-2} :

$$\boxed{2 \log k_n \leq -\log \tau_n + O(\log a_n V_n)}. \quad (5)$$

Equivalently: the wavenumber tracks $(\nu(T-t))^{-1/2}$, and the local Reynolds number $\text{Re} = UL/\nu$ stays of order one.

2. Construction I: Alpöge–Buckmaster

What was reproduced

The modulation system (`Reduced.lean`), state (C, Φ, Θ, V) , control $p(s)$, parameter ψ_0 :

$$\dot{C} = \Phi, \quad \dot{\Phi} = \frac{\sin(\psi_0 p)}{\psi_0}, \quad \dot{\Theta} = \frac{V}{\psi_0^2 + C^2}, \quad \dot{V} = \left[\cos(\psi_0 p) - \frac{C \sin(\psi_0 p)}{\psi_0} \right] \Theta. \quad (6)$$

(Θ, V) is a linear oscillator whose spring constant is that bracket, hyperbolic when positive and elliptic when negative, flipping exactly at $\text{overTilt} := C \tan(\psi_0 p)/\psi_0 = 1$. The control parks the system on that boundary with $V = 0$, over-tilts to rotate the Prüfer angle, and returns.

Gate	Check	Result
A	their <code>BridgeStatement</code> identity, 4000 random points	max dev 2.132×10^{-14}
B	Prüfer angle against $\arctan(V/\Theta)$, 2000 points	max dev 2.967×10^{-11}
C	full <code>LemmaP_Statement</code>	$D^* = 1.046671087216$

Gate A is load-bearing: it is *their* algebraic identity holding on an independent implementation, so the port is verified rather than plausible. Gate C reproduces every conclusion of the lemma, with `ParkedAt` exact to twelve digits. Their specification pins the selected maneuver to `Dsel` $\in [1.47, 1.60]$; seven crossing-data choices land D^* inside that interval, so the reduced system reproduces the *selected* maneuver.

The maneuver is a poor amplifier

Across the admissible box ($\psi_0 \in \{0.5, 0.1, 0.01, 0.001\}$, $\sigma_v \in \{0.5, 1, 3\}$, three crossing states, 36 maneuvers), the Θ gain has median 1.108 and the $\psi_0^2 + C^2$ throttle has median 9.161. The gain never once exceeds the throttle. Every maneuver drives C up and C^2 sits in the denominator of its own growth term. The $\psi_0 = 0.01$ and 0.001 rows agree to five digits, so this is uniform structure, not a small-parameter artifact.

`Main.lean` requires $A(n+1) \geq q^n A(n)$ for *every* q , which an 11 percent per-stage gain cannot supply. Their own definition locates the growth: $\text{Agrad} = |\Theta| \cdot \|\xi\|$ (`S6/State.lean:84`). It is in the frequency.

The reduction under dissipation

The ansatz is WKB, with $\xi = N \cdot \xi_N$, $\xi_N = \nabla X(0)$ (`S6/State.lean:72`), and their own bound (`S4/Objects.lean:518`) gives $2^{-1/2} \leq \|\xi_N\| \leq 6.5/\psi_0$, so the carrier satisfies $\|\xi\| \geq 0.707 N$. This settles by derivation a question reading alone leaves open: $\ell = N^{-2\beta}$ is an envelope scale, not the carrier.

On a plane wave $\nu\Delta$ returns $-\nu\|\xi\|^2$ times the same wave. It is diagonal on the ansatz, generating no structure the ansatz cannot carry, so the reduction closes and dissipation enters as exactly one term, $\dot{V} = -(bC/\text{otDen})\Theta - \delta V$ with $\delta = \nu\|\xi\|^2$.

Inside one maneuver this is nearly harmless. With the initial layer included,

$$\text{gain} = 1 + \frac{0.207886}{\delta}, \quad 0.207886 = \underbrace{\frac{V_0}{\psi_0^2 + C_0^2}}_{0.399840} - \underbrace{\int \frac{bC}{\text{otDen} \cdot \text{feedDen}} ds}_{0.191954}, \quad (7)$$

and Radau integration at $\delta = 10^6$ gives $\delta(\text{gain} - 1) = 0.20788661$, six digits. The gain approaches 1 from above and never reverses. Viscosity does not destroy the maneuver; it neutralizes it. That is also why the reduced system alone cannot settle the question: $\dot{\xi} = -\Phi^\top \xi$ is pure advection, so dissipation never touches the wavenumber, and Φ is the lowest-frequency object present.

Failure against Criterion 1

Stages do not chain the same oscillator. Each builds a child whose V is read off the parent's vorticity (`S6/RunClass.lean:105`), and vorticity is exactly what $\nu\Delta$ attacks. Because $q_\beta(m_0 + \beta) = \beta$ identically, combining `Esc(a)` with `MoveOut.Λge` gives

$$\ln N_{n+1} = \frac{0.634[(5n+1)\ln \mathfrak{a}_n + 104]}{\beta_n}, \quad 1/\beta_n \geq 80, \quad (8)$$

so $1/\beta$ multiplies the entire bracket and every halving doubles it. The maneuver duration is $\varsigma_n = 1.1588 \mathfrak{a}_{n-1}^{-1/2}$, giving $-\log \tau_n = \frac{1}{2} \log \mathfrak{a}_{n-1}$. Criterion 1 demands

$$\underbrace{2 \log N_n}_{17,557} \leq \underbrace{\frac{1}{2} \log \mathfrak{a}_{n-1}}_{5.76} \quad \text{at } n = 1, \quad (9)$$

a ratio of 3050, exceeded by e^{17551} . Volume factors are powers of length scales and cannot close a gap of that size in the log.

Both branches of the halving dichotomy lose. If halving fires, `Esc` forces ρ up and N explodes. If it does not, β stays at $1/80$ where $m_0 = 0$ and $\log_{10} N = 2.2 \times 10^6$ immediately. Equivalently, $\nu_{\text{crit}} \approx 10^{-4.4 \times 10^6}$ at stage 1.

3. Construction II: OpenAI

Normalization

The construction is built at $\nu = 1$ and the general statement follows by scaling, in `ComparatorBridge.lean` at lines 64 and 217:

```
def rescaledForce (nu) (f) : VelocityField :=
  rescale (nu ^ 2) nu f
theorem normalized_solution ... :
  GlobalSolutionOne f (rescale nu-inv nu-inv ...)
```

This is standard and correct, and it is why the construction files contain no viscosity: in normalized variables there is none to contain. The rescaling preserves Fefferman's condition (5), polynomial decay being scale invariant as a class.

The wavenumber schedule

From `ChartScales.lean` and `SlotColoring.lean`, with $T_g = 4 + \sqrt{2}$, $\Lambda = 4 - \sqrt{2}$, $Q_n = 2^{-n}$, $\kappa = 10^{-5}$, $\rho = \log \Lambda / \log T_g$ and $m = \lfloor (n(1+h) \log 2 - 2 \log n) / \log T_g \rfloor$:

$$\text{radialExponent}(h) = 2((1+h)\rho - h\kappa), \quad (10)$$

$$\text{timeCoefficient} = T_g^m Q_n^{1+h} \quad (\text{reciprocal time, } \text{slotLength} = 2r_0/\text{timeCoeff}), \quad (11)$$

$$\text{radialCoefficient} = \Lambda^m Q_n^{\text{radialExponent}/2} \quad (\text{reciprocal length, the wavenumber}). \quad (12)$$

Criterion 1 demands $\text{radialCoeff} = \text{timeCoeff}^{1/2}$ up to order one. Note $\rho = 0.5624713724$, which is *not* $\frac{1}{2}$, so at first sight the scaling looks wrong. It is not, because the ρ dependence cancels identically:

$$\begin{aligned} \log \frac{\text{radialCoeff}}{\text{timeCoeff}^{1/2}} &= \underbrace{m \log \frac{\Lambda}{\sqrt{T_g}}}_{[n(1+h) \log 2 - 2 \log n](\rho - \frac{1}{2})} + \underbrace{\log Q_n[(1+h)(\rho - \frac{1}{2}) - h\kappa]}_{-n \log 2 [(1+h)(\rho - \frac{1}{2}) - h\kappa]} \\ &= nh\kappa \log 2 - 2(\rho - \frac{1}{2}) \log n. \end{aligned} \quad (13)$$

The cancellation holds only for their specific `radialExponent`, and it is insensitive to Λ : mutation testing (Section 5) replaces $\Lambda = 4 - \sqrt{2}$ by 3 and the residual is unchanged, because ρ is *defined* as $\log \Lambda / \log T_g$. Λ is a free parameter of the scaling; the exponent form is what carries it. Two terms survive: a harmless polynomial correction $n^{-2(\rho-1/2)} = n^{-0.125}$ from the $-2 \log n$ inside `nativeArgument`, and $+nh\kappa \log 2$, a deliberate small *upward* detuning off the critical scale.

h	n	$\log(\text{ratio})$	predicted	ratio
0.10	10	-0.371	-0.288	0.690
0.10	10^4	-1.175	-1.144	0.309
0.25	10^6	-0.058	+0.007	0.943
0.40	10^2	-0.575	-0.575	0.563
0.40	10^6	+0.944	+1.046	2.570

Across $n = 10$ to 10^6 and $h \in \{0.1, 0.25, 0.4\}$ the ratio stays in $[0.24, 2.57]$, matching the derived formula to within 0.10 in the log, the residual being floor jitter. The construction sits on the critical viscous scaling and leans off it by $\kappa = 10^{-5}$: enough excess to blow up, small enough that dissipation cannot arrest it. The last row shows that detuning becoming visible at large n . The width of the needle's eye is written down as a constant.

4. A third test, on data neither construction has seen

Sections 2 and 3 test the criterion against constructions whose scaling I had already worked through by hand, so neither could surprise it. This section applies it to a third object, using run data produced weeks before either formalization was public.

Setup

The runs are a certified Chen–Hou self-similar profile integrated at three viscosities, $\nu \in \{10^{-4}, 3 \times 10^{-6}, 10^{-6}\}$, with a measured blowup time $T^* = 0.0036185$ and reference scale $\Omega^* = 1618.388$. All three share initial data to six digits.

The buoyancy θ is materially transported at bounded amplitude, so $|\nabla \theta|$ is proportional to the wavenumber: $k(t) \propto \sup |\nabla \theta|$. This is the same structure as $\text{Agrad} = |\Theta| \cdot \|\xi\|$ with $|\Theta|$ bounded, which is what makes the comparison meaningful.

The discrimination

Criterion 1 requires $k \sim (\nu(T - t))^{-1/2}$, that is

$$\frac{d \log k}{d \log \nu} = -\frac{1}{2} \quad (\text{viscous scale}), \quad \text{against} \quad \frac{d \log k}{d \log \nu} = 0 \quad (\text{Euler scale}). \quad (14)$$

t	$T^* - t$	$k(10^{-4})$	$k(3 \times 10^{-6})$	$k(10^{-6})$	$d \log k / d \log \nu$
0.00200	0.001619	5.55×10^5	5.87×10^5	5.91×10^5	-0.0145
0.00240	0.001219	7.65×10^5	8.30×10^5	8.42×10^5	-0.0213
0.00270	0.000919	1.15×10^6	1.28×10^6	1.30×10^6	-0.0280
0.00290	0.000719	1.67×10^6	2.01×10^6	2.07×10^6	-0.0476
0.00300	0.000619	2.09×10^6	2.79×10^6	2.87×10^6	-0.0720

The mean is -0.0367 , against -0.5 for the viscous scale and 0 for the Euler scale: **a factor 12.6 closer to Euler**. The sharpest single reading is $k(3 \times 10^{-6})/k(10^{-6}) = 0.9718$ at $t = 0.0030$, where viscous scaling would demand $\sqrt{3} = 1.732$. Tripling the viscosity moves the wavenumber by under three percent.

The same statement as a collapse test: under the Leray rescaling $R = k(\nu(T^* - t))^{1/2}$, a viscous-scale flow collapses to spread ≈ 1 . Measured spread is 7.3 to 9.4 , against the $\nu^{1/2} = 10.0$ predicted for a wavenumber set by the inviscid dynamics. It does not collapse; it separates by almost exactly the predicted factor.

Why this is the load-bearing test

The criterion classifies a construction it had no knowledge of, correctly, by more than an order of magnitude. A criterion that measured nothing could not produce that separation, and unlike Sections 2 and 3 this outcome was not known in advance.

It also relates to a result measured on these runs months earlier, and the direction of that relation matters. The implications run one way:

$$\begin{aligned}
 k \text{ } \nu\text{-independent} &\implies \nu k^2 \text{ regular} \implies \text{linear response,} \\
 k \text{ } \nu\text{-independent} &\implies \text{not at the viscous scale.}
 \end{aligned} \tag{15}$$

They do not reverse. A regular perturbation requires only $k = o(\nu^{-1/2})$, so a wavenumber going like $\nu^{-1/4}$ would give linear response while being far from ν -independent. The earlier paper measured linear response, which is the *consequence*; the measurement above is the *premise* that implies it, taken from the same runs through a different diagnostic.

What is fair to say is that two results obtained months apart by different methods converge on one fact about this profile, and that the earlier one reached the same content empirically rather than deductively: narrower in scope, one object rather than a general condition, and carrying quantities the deductive route does not produce.

Limits

The runs stop at $t/T^* = 0.83$, so this is not asymptotic. The slope drifts from -0.0145 to -0.0720 across the window, so dissipation does matter more as scales refine; it would have to steepen roughly sevenfold to reach $-\frac{1}{2}$, and these runs cannot exclude that. And $k \propto \sup |\nabla \theta|$ assumes $\sup |\theta|$ bounded, the same assumption the criterion makes elsewhere.

5. Adversarial checks

Recorded because they are the reasons to doubt the above. Two of them narrowed it and one nearly produced a false claim.

Can the forcing absorb $-\nu\Delta u$ in Construction I? Partly. Layer velocity gives $\nu\Delta u \sim \nu N^{1-m_0}$, so at stage 3 and beyond ($m_0 = 10, 70, \dots$) the quantity is negligible and an admissible forcing absorbs it. The objection wins in the tail. It fails at the base: `HalvingRule` fixes $\beta(1) = \beta(2) = 1/80$ where $m_0 = 0$ exactly, the quantity is νN^{+1} with $\log_{10} N = 2.2 \times 10^6$, and `SmoothForce.bounded` requires a single finite M uniform over $[0, T)$. This is why the claim is stated at the base only.

Does Construction II ignore viscosity? No, and asserting so would have been a false claim. Only 6 of 643 construction files mention `nu`, and four of those use it as a generic smallness threshold. The apparent absence is the normalization of Section 3, detected by noticing that `ComparatorBridge.lean` contains no `nu` while being called with one, which led to the Unicode ν and the rescaling lemmas.

Is $\|\xi_N\|$ order one? Bounded below by $2^{-1/2}$ in their own file. The upper bound only increases dissipation.

Is the duration right? ϖ_S uses `Agrad` at `manStart`, not at t_k , so the duration is an overestimate by at most $\rho^{1/6} \approx 46$, i.e. 3.8 in the log against 17,557.

Does the criterion measure anything, or pass everything? Mutation testing on Construction II: perturbing κ from 10^{-5} to 10^{-1} drives the residual to +17,327; halving the leading 2 in `radialExponent` gives +243,671; replacing ρ by the naive $\frac{1}{2}$ gives +54,127. All three are caught. A control mutation outside the ratio is correctly not flagged. The criterion discriminates rather than condemning.

Does a PASS certify the roles? No, and this is the sharpest limitation. The test relates `radialCoefficient` to `timeCoefficient`; that they are a wavenumber and a reciprocal time is read from naming and from `slotLength = 2r0/timeCoefficient`. It is not derived. In a $\nu = 1$ normalized construction length and time are not independent, so dimensional homogeneity cannot separate them either. A PASS is conditional on that reading.

What would refute this. A demonstration that the effective dissipation is not $\nu\|\xi\|^2$ at the carrier scale, for instance if an envelope or stage cutoff shields the carrier. That was looked for and not found, over one working session.

6. Status and limitations

Criterion 1 is a corollary of (1) and the energy identity. Applying it to a specific construction uses the linear heat kernel estimate on the stage timescale, and that is the weakest link.

Nothing here verifies either proof. Construction II is 616K lines of Lean and was not compiled for this note; what was checked is that its wavenumber schedule has the character Criterion 1 requires. Nothing here contradicts Construction I's published theorem, which concerns the inviscid forced system and is untouched.

The contribution is a number where there was an expectation, and a test that separates rather than condemns. Any viscous extension of Construction I must clear the bar set by (8) with $1/\beta_n \geq 80$.

7. Method and reproduction

All numbers were produced with `scipy` on a laptop in seconds. Gate A is the transcription check; run it before trusting anything downstream. Two naming traps in Construction I: `nu` in the source is $N\sqrt{\kappa/A\text{flat}}$ and is not viscosity, and κ is the buoyancy constant from the theorem statement and is also not viscosity.

```
git clone https://github.com/tristanbuckmaster/fluid_lean
git clone https://github.com/openai/NavierStokesAndEuler
python3 rs_gate.py      # gates A, B, C and the 36-maneuver sweep
python3 rs_viscous.py   # viscous reduction, regression at delta = 0
python3 rs_stiff.py     # large-delta law, gain = 1 + 0.207886/delta
python3 seam3.py        # frequency ladder and nu_crit table
python3 adversarial.py  # Dsel range, forcing-absorption test
python3 chartscales.py  # Construction II wavenumber ratio
```

This work was done with substantial assistance from Claude (Anthropic), under my direction and subject to the verification gates described above. The reproduction gates exist so that no claim here rests on trust, including trust in me.