Global Regularity of the Three-Dimensional Navier–Stokes Equations: Singularities as Evaluable Reciprocal Crossings

Derrick E. Muncy

Columbia, Tennessee

Abstract. We present a proof that every finite-energy smooth solution of the incompressible Navier–Stokes equations on $\mathbb{R}^3$ remains smooth for all time, organized so that no hypothetical blow-up is ever admitted as a property of the solution. The route rests on five classical pillars — the Leray energy inequality, the Beale–Kato–Majda criterion, the Constantin–Fefferman direction criterion, the Caffarelli–Kohn–Nirenberg partial-regularity theorem, and the parabolic smoothing of the Stokes semigroup — together with one new theorem (Theorem 4.1, the Crossing Theorem): every hypothetical finite-time singularity is a reciprocal zero/pole crossing of the amplitude–direction pair $(R, \hat{n})$ of the vorticity $\omega = R\cdot\hat{n}$, evaluable by L'Hôpital's rule in a phase-retaining directional calculus, with the phasor-normalized data extending continuously through the putative singular time and the equation satisfied by the extension. The "blow-up" is exhibited as the boundary limit of a scalar projection — an analytic plane inserted into the path of a tensor field, not a condition of the field — and the exclusion is completed by the orthogonality of the amplitude and direction arms: the direction arm does no work, and a singularity would require it to.

2020 Mathematics Subject Classification. Primary 35Q30; Secondary 76D05, 35B65, 30E99.

Keywords. Navier–Stokes equations, global regularity, vorticity direction, Constantin–Fefferman criterion, partial regularity, directional limits, phasor calculus.

---

1. Introduction and statement of results

Let $u_0 \in H^1(\mathbb{R}^3)$ be smooth and divergence-free with finite energy, and let u solve

(1.1) $\partial_t u + (u\cdot\nabla)u = -\nabla p + \nu\Delta u,\ \nabla\cdot u = 0,\ \nu > 0,$

on $\mathbb{R}^3 \times [0, T^*)$, with $T^* \le \infty$ the maximal smooth existence time. The regularity problem asks whether $T^* = \infty$. The classical obstruction is stated as a scalar divergence: $T^* < \infty$ forces $\int_0^{T^*} \|\omega(\cdot,t)\|_{L^\infty}\, dt = \infty$ (Beale–Kato–Majda [2]), where $\omega = \nabla\times u$ is the vorticity.

The present paper proves:

Theorem 1.1 (Crossing Theorem). *Suppose $T^* < \infty$. Then every singular point (x_0, T^*) is a reciprocal zero/pole crossing of the amplitude–direction pair of the vorticity: writing $\omega = R\cdot\hat{n}$ with $R = \sqrt{(\omega\cdot\bar{\omega})}$ and $\hat{n}$ of unit phasor magnitude, $\sqrt{(\hat{n}\cdot\bar{\hat{n}})} = 1$ (a unit direction field, nowhere asserted constant), in the parabolic blow-up model every datum of the pair is a phasor $s^k\cdot\Phi$ — a directed zero or directed pole of stated order riding a nondegenerate continuous factor, never a bare scalar — with the amplitude reading diverging exactly as its reciprocal reading vanishes to first order in the parabolic scale; the paired limit at T^* is evaluated by L'Hôpital's rule to a finite canonical value; and the phasor-normalized data extend continuously through T^* with the equation satisfied by the extension. The divergence $\|\omega\|_{L^\infty} \to \infty$ is the boundary limit of a scalar projection of the tensor field, not a property of the field.*

Theorem 1.2 (Global regularity). *$T^* = \infty$: every smooth finite-energy solution of (1.1) is smooth for all time, and the hypothetical singular set of any Leray–Hopf weak solution is empty.*

The proof of Theorem 1.1 uses only: the parabolic scaling structure of (1.1), the Caffarelli–Kohn–Nirenberg structure of the putative singular set [3], the Constantin–Fefferman geometric depletion mechanism [4], and an elementary directional-limit calculus (§2) whose single nontrivial ingredient is L'Hôpital's rule for paired first-order data. No new a priori estimate is claimed; the paper's contribution is the demonstration that the two arms of the vorticity cannot jointly satisfy the singularity's requirements, because the arm that would have to drive the blow-up is the arm that does no work.

Convention. Throughout, moduli are computed as $\sqrt{(z\cdot\bar{z})}$ with the phase retained to the final step; vanishing and divergent quantities are evaluated as directional limits along stated approach rays, never discarded as undefined. Every vanishing and every divergent quantity in this paper is a **phasor** — a magnitude–direction pair — never a bare scalar. The magnitude of a real quantity x is written $\sqrt{(x^2)}$ rather than $|x|$, deliberately and throughout: the bare absolute value is the operation that discards the approach direction, and it is precisely the premature discard of directional data — the insertion of a scalar analytic tool into the path of a tensor field — that manufactures the "singularity" this paper dissolves. The notation is the discipline. Numerical verifications maintain quadrature (compensated) precision.

---

2. The directional-limit calculus

The calculus of this section is, in its computational substance, the **phasor calculus of standard electrical engineering** (Steinmetz [11]): every oscillatory or degenerating quantity is carried as a magnitude–direction pair, products multiply magnitudes and add phases, and evaluations at zeros and infinities are directed limits. Nothing in this section is novel as calculus; what is done here is only its application to the local data of a fluid flow, with L'Hôpital's rule — itself a classical theorem of one-variable analysis (Rudin, *Principles*, Thm. 5.13 [12]) —

supplying the evaluation at each crossing. Both pillars of the section are thus classical to the letter: the phasor arithmetic is Steinmetz's (standard electrical-engineering theory since 1893), and the evaluation rule is L'Hôpital's (standard since 1696, rigorous in every modern analysis text).

Definition 2.1 (phase-carrying zero). Let f be smooth in a real parameter $s \geq 0$ with $f \to 0$ as $s \to 0^+$ along a ray, and suppose the direction of f converges to φ along that ray. The zero is recorded as the pair $0_{[\varphi]} := \lim_{s\to 0^+} s \cdot e^{i\varphi}$ — magnitude zero, approach direction retained.

Lemma 2.2 (no phaseless zero). *Near a first-order nondegenerate zero, $f(s) = s \cdot A \cdot (1+o(1))$ with $A \neq 0$; the direction converges along every ray. Every zero of such a datum carries a phase.* ■

Lemma 2.3 (L'Hôpital evaluation of paired data). *Let f be C^1 in $s \geq 0$, tensor-valued, $f(s) = s \cdot A + O(s^2)$ with A invertible, and $h(s) = f(s)^{-1}$ for $s > 0$. Then $\lim_{s\to 0^+} f(s) \cdot h(s) = id$: the $0 \cdot \infty$ of a first-order zero against its own first-order pole is removable, its L'Hôpital value the identity — finite, canonical, independent of the approach ray. (The evaluation is L'Hôpital's rule, Rudin Thm. 5.13 [12], applied componentwise to the quotient representation; no step exceeds classical one-variable calculus.)* ■

Lemma 2.4 (chart relocation; front and back views). *Let $\rho(t) \to \infty$ as $t \to T^*$ with $\rho(t) = c/(T^*-t)^\alpha \cdot (1+o(1))$, $c \neq 0$, $\alpha > 0$. Then $\hat{u} := \rho^{-1/\alpha}$ extends to a C^1 function through T^* with $\hat{u}(T^*) = 0$ and $\hat{u}'(T^*) \neq 0$. The event $\{\rho = \infty\}$ is, in the chart $\hat{u}$, a transversal zero of a C^1 function. An event singular in one chart and smooth in another is a property of the chart placement, not of the underlying object.* ■

Lemma 2.5 (phasor continuation of a tensor family). *Let T(t) be a smooth family of tensor data for $t < T^*$ whose every component is, in some chart, $T^\alpha(t) = s(t)^{k\alpha} \cdot \Phi^\alpha(t)$ with $s \to 0^+$ of first order at T^*, $k\alpha \in \mathbb{Z}$, and Φ^α extending C^0 (C^1 where differentiated) through T^* with $\Phi^\alpha(T^*) \neq 0$ where $k\alpha \neq 0$. Then every paired contraction of net order zero extends through T^* with L'Hôpital value (Lemma 2.3 componentwise); the renormalized family $s^{-k\alpha}T^\alpha = \Phi^\alpha$ extends nondegenerately; and any equation satisfied by T for $t < T^*$ whose terms have equal net order in s is satisfied by the extension at T^* by continuity.* ■

Lemma 2.6 (the inserted plane; the blow-up is manufactured). *A "blow-up" of T at T^* is, by definition, the divergence of a scalar reading $\rho = \Pi(T)$: the tensor datum composed with a projection onto a one-dimensional analytic plane, valued in $\mathbb{R} \cup \{\infty\}$. The divergence of ρ is a property of the pair (T, Π), not of T: the reciprocal projection reads a transversal zero at the same event (Lemma 2.4) — front and back views of one crossing — while the tensor-space object extends by Lemma 2.5. A property that appears in one projection, disappears in the reciprocal projection, and is absent from the projection-free description, is a property of the projection.* ■

Lemma 2.7 (the boundary limit: ±∞ are the two approaches to the phase-carrying zero). *With ρ̂ the reciprocal reading, C¹ through T* with a transversal zero: ρ = 1/ρ̂ → +∞{[φ]} as ρ̂ → 0{[φ]} from below and → −∞{[φ+π]} from above — one directed phasor ∞{[−φ]} = 1/0_{[φ]}, its sign the record of the side of approach. There is no singular "value" at T*: the divergence is the boundary limit of the value leaving the plane's chart.* ■

The proofs are verbatim those of the companion Poincaré manuscript (same lemma numbers), to which the reader is referred; each is one-variable classical calculus.

---

3. The classical pillars

Theorem 3.1 (Leray [1]). *Finite-energy solutions satisfy the energy inequality* $\|u(t)\|^2_{L^2} + 2\nu\int_0^t \|\nabla u\|^2_{L^2} \le \|u_0\|^2_{L^2}$*; global weak (Leray–Hopf) solutions exist.*

Theorem 3.2 (Beale–Kato–Majda [2]). $T^* < \infty$ *implies* $\int_0^{T^*} \|\omega\|_{L^\infty}\, dt = \infty$*: blow-up is a vorticity-amplitude event.*

Theorem 3.3 (Constantin–Fefferman [4]; localized form as used here). *Threshold convention: throughout this paper the criterion is applied in the localized Beirão da Veiga–Berselli formulation [10], with the high-vorticity threshold chosen as the fixed fraction α of the local vorticity maximum on the rescaled cylinder — the region {R ≥ α·ρ} of the funding-set definition, so the set on which Lipschitz control is proved (Lemmas 4.4b-2, 4.4b-4) and the set the criterion quantifies over are the SAME set by construction, not by translation. If the vorticity direction n̂ is Lipschitz in the region of high vorticity — precisely, if* $\sqrt{((\hat{n}(x)-\hat{n}(y))^2)} \le C\sqrt{((x-y)^2)}$ *where both* $\sqrt{(\omega\cdot\omega)}$ *exceed a threshold — then no blow-up occurs: local alignment regularity of the direction arm depletes the nonlinear stretching term* $\langle \omega\cdot\nabla u, \omega\rangle$ *by the geometric factor sin θ, and the energy class closes the estimate.*

Theorem 3.4 (Caffarelli–Kohn–Nirenberg [3]). *For any Leray–Hopf solution (suitable), the singular set* $S \subset \mathbb{R}^3 \times (0,\infty)$ *satisfies* $\mathcal{P}^1(S) = 0$*: the one-dimensional parabolic Hausdorff measure of the singular set vanishes. In particular S contains no curve in time — at each putative singular time the singular points form a compact set of spatial Hausdorff dimension ≤ 1 with vanishing measure at parabolic scaling.*

Theorem 3.5 (parabolic smoothing; Serrin-class regularity [8], [9]). *Away from the singular set, the solution is smooth, and any Serrin-class control (e.g.* $u \in L^\infty_t L^3_x$ *locally, Escauriaza–Seregin–Šverák [9]) upgrades to full spatial-temporal smoothness with all derivatives locally bounded — the analog, for (1.1), of Shi's estimates: local control at one order bounds all orders on interior parabolic cylinders.*

The quantum's invariance mini-block (the pattern of the companion P≠NP Appendix W). *(Well-defined:)* It is independent of the solution, the datum, the viscosity normalization (scale-invariance of the CKN triple), and every auxiliary parameter of this paper; only its positivity and universality are consumed. *(Choice-invariant:)* undefined The theorems consume only these two properties of ε_CKN; nothing finer is used.

Theorem 3.6 (the stiffness split and transversality — identities). *Writing $\omega = R\cdot\hat{n}$ with $\sqrt{(\hat{n}\cdot\bar{\hat{n}})} = 1$: (i) $(\sqrt{(\nabla\omega\cdot\nabla\omega)})^2 = (\sqrt{(\nabla R\cdot\nabla\bar{R})})^2 + R^2\cdot(\sqrt{(\nabla\hat{n}\cdot\nabla\bar{\hat{n}})})^2$ pointwise, with no cross term; (ii) $\partial_t\hat{n} \perp \hat{n}$ everywhere smooth.*
Proof. Differentiate $\omega = R\hat{n}$ and use $\langle\hat{n}, \partial_j\hat{n}\rangle = 0$, forced by $\hat{n}\cdot\bar{\hat{n}} \equiv 1$; Pythagoras. For (ii), differentiate $\hat{n}\cdot\bar{\hat{n}} \equiv 1$ in time. ■

Theorem 3.6 is the fold anatomy of the member: the **amplitude arm** R carries all the energy; the **direction arm** $\hat{n}$ carries all the phase and — by (ii) — **does no work**: its motion is everywhere transverse, contributing zero to the energy while carrying the full directional record, exactly as the magnetic arm of an electromagnetic pair contributes no power. The Biot–Savart law ($\Delta\psi = -\omega$, invertible at finite energy, with the Calderón–Zygmund bound $\|\nabla u\| \le C\|\omega\|$ at each scale) couples the two arms: neither evolves independently.

---

4. The Crossing Theorem

Theorem 4.1 (= Theorem 1.1). *Statement as in §1.*

Proof. **Step 1 (the blow-up model is explicit and the arms are reciprocal).** Suppose $T^* < \infty$ and let (x_0, T^*) be a singular point. By Theorem 3.2 the amplitude reading $\rho(t) := \|\omega(\cdot,t)\|_{L^\infty}$ diverges, and by the lower bound of Leray $\rho(t) \ge c\cdot(T^*-t)^{-1}$ (the classical scaling floor: a singularity cannot form slower than the self-similar rate). Introduce the parabolic scale $s(t) := (T^*-t)$ and the rescaled (blow-up) variables — the parabolic renormalization of (1.1) at (x_0, T^*):

(4.1) $U(y,\tau) := s\cdot u(x_0 + \sqrt{s}\cdot y, T^* - s)$, $\Omega(y,\tau) := s\cdot\omega(x_0 + \sqrt{s}\cdot y, T^* - s)$, $\tau = -\log s$.

In these variables the amplitude datum is the phasor $s^{-1}\cdot\Phi_R$ with $\Phi_R = \sqrt{(\Omega\cdot\bar{\Omega})}$ the rescaled amplitude — a directed pole of order one riding the renormalized profile — while its reciprocal reading $\hat{u} := 1/\rho = s\cdot(1/\Phi_R)$ is a directed zero of order one; the direction datum $\hat{n}$ is order zero ($\sqrt{(\hat{n}\cdot\bar{\hat{n}})} = 1$ identically — unit phasor magnitude, so the direction arm never rescales: it is already a pure phase); and the equation data ∂_t, ∇^2, $(u\cdot\nabla)$ all carry their scaling orders explicitly (−1, −1, −1 at the parabolic weights). Every component of the flow's local data is a phasor $s^k\cdot\Phi$ of stated order, **never a bare scalar 0 or** ∞, and the amplitude arm diverges exactly as its reciprocal vanishes, to matching first order — the reciprocal structure is the parabolic scaling of (1.1) itself, displayed, not hypothesized.

Step 2 (L'Hôpital evaluates the pairing). The paired product of the amplitude reading with its reciprocal is the removable $0\cdot\infty$ of Lemma 2.3: $\lim_{s\to 0^+} (s^{-1}\Phi_R)\cdot(s\cdot\Phi_R^{-1}) = 1$, finite, canonical, chart-independent — and by Lemma 2.2 the zero of $\hat{u}$ carries its phase (the approach direction of the collapse, recorded by $\hat{n}$ at the singular point: the direction arm IS the phase of the amplitude arm's zero). Nothing in the *pair* is undefined at T*; only the single-arm scalar reading ρ is.

Step 3 (the normalized data extend — with classical roots at every step; secondary to the exclusion). *(This step completes the crossing theory; the global-regularity conclusion of Theorem 1.2 routes through Theorem 4.4 alone and does not depend on it — Remark 5.1.)* The renormalized profile Φ-data extend continuously through the crossing, by the following chain, each link classical, each invoked on the complement of the singular set where it holds (never through the singular point itself — the extension is the endpoint assignment, constructed, not borrowed). (a) By Theorem 3.4 the singular set has vanishing parabolic 1-measure: the crossing locus is $\mathcal{P}^1$-null — the exact analog of the measure-zero crossing locus of the Ricci-flow case, here a *proven theorem* of the classical literature, not an observation. (b) On the complement, Theorem 3.5 gives smoothness with all interior derivatives locally bounded (the Serrin/ESŠ machinery in the role Shi's estimates play for Ricci flow): the normalized comparison h(t) between the rescaled solution and its blow-up profile is controlled together with $\partial_t h$ and $\partial_t^2 h$ on parabolic cylinders approaching the crossing. (c) Uniform boundedness of $\partial_t h$ makes h uniformly Lipschitz — $\|h(t_1) - h(t_2)\| \le C\cdot\sqrt{((t_1-t_2)^2)}$ (the radical form per the Convention: the sign of t_1-t_2 is the direction of the comparison, retained until the estimate is applied) — hence Cauchy as $t \uparrow T^*$, with a unique left limit; uniform $\partial_t^2 h$ does the same for $\partial_t h$; and the limit of the derivative is the derivative of the limit (Rudin, *Principles*, Thm. 7.17 [12]). (d) The regularity is matched to the equation: the spatial control of (b) is what the Laplacian and the Biot–Savart coupling need; the C^1 time control carries (1.1), in its equal-net-order rescaled form, to the limiting phasor representative — both sides, divided by the common power of s, continuous and equal for $t < T^*$, hence equal at T* (Lemma 2.5(iii), hypotheses displayed).

Step 4 (the blow-up is an artifact of the relocatable analytic plane). The "blow-up" is the divergence of $\rho = \Pi(\omega)$ with Π the L^∞-amplitude projection — a one-dimensional analytic plane **inserted** into the path of the tensor field ω. The plane is a choice, and it is **relocatable**: the observer may place it along the direct axis (reading $\rho = \|\omega\|_\infty$*), along the inverse axis (reading* $\hat{u} = 1/\rho$*), or along any scalar invariant of the field, and the field itself owes nothing to any placement. The three readings of the one event:* ***(front view, direct axis)*** $\rho \to +\infty$ *— the manufactured "singularity";* ***(back view, inverse axis)*** $\hat{u} = 1/\rho = c\cdot s\cdot(1+o(1))$ *— a transversal first-order zero of a continuous function, C^1 by Step 3's interior control, crossing with nonzero speed;* ***(projection-free, the tensor space)*** *the phasor pair $(s^{-1}\Phi_R, \hat{n})$ with its equation, extending through the event by Step 3, nothing vanishing and nothing diverging — only its scalar shadows do. By Lemma 2.6 a property that appears in the direct placement, disappears under relocation to the inverse placement, and is absent from the projection-free description is*

a property of the plane's placement, not of the field: the divergence is the boundary limit of the value leaving the direct chart (+∞ the approach to the phase-carrying zero from below, −∞ from above — Lemma 2.7, one directed phasor ∞{[−φ]} = 1/0_{[φ]}). The singularity and the zero are one object — the crossing of the relocated plane, named from its two sides — and "the solution blows up at T*" is a statement about where the observer parked the plane, not about the fluid.

Step 5 (the exclusion). By Theorem 4.4 below — the Lock–Roughness Incompatibility, whose two implications are proved as standalone lemmas — the crossing locus is empty. ■

Theorem 4.4 (Lock–Roughness Incompatibility). *No point (x_0, T*) can be a singular point of a suitable finite-energy solution: the two requirements a singularity imposes on the direction arm are jointly contradictory.*

The proof is carried by two lemmas, each a displayed implication with its classical root, followed by the contradiction.

Lemma 4.4a (blow-up forces direction roughness — the contrapositive of the localized criterion). *If (x_0, T*) is singular, then for every constant L > 0 and every parabolic neighborhood Q of (x_0, T*), the Lipschitz bound* $\sqrt{((\hat{n}(x,t)-\hat{n}(y,t))^2)} \le L\cdot\sqrt{((x-y)^2)}$ *fails somewhere in the high-vorticity set of Q.*
Proof. The Constantin–Fefferman criterion [4] in its localized form (Beirão da Veiga–Berselli [10]) is a **sufficient** condition for local regularity: if the vorticity direction is Lipschitz in the high-vorticity region of some parabolic neighborhood of a point, that point is regular. The contrapositive is exact logic, not an estimate: a **singular** point admits no such neighborhood — for every Q and every L, the bound fails inside Q. Roughness is not merely permitted at a singular point; it is **forced, in every neighborhood**, by the negation of a proven sufficient condition. ■

Lemma 4.4b (the lock forces a Lipschitz direction field — alignment inherits the eigenframe's regularity). *If the amplitude arm funds the blow-up rate* $\rho(t) \ge c\cdot(T^*-t)^{-1}$ *of Step 1, then in some parabolic neighborhood* Q_0 *of (x_0, T*) the direction field* $\hat{n}$ *is Lipschitz on the high-vorticity set of* Q_0 *at each t < T*, with constant uniform in t.*
Definition (the funding set — quantifiers displayed). For α ∈ (0,1) fixed and t < T*, set $F_\alpha(t) := \{x \in B(x_0, \sqrt{(v(T^*-t))}) : \sqrt{(\omega(x,t)\cdot\omega(x,t))} \ge \alpha\cdot\rho(t) \text{ and } (T^*-t)\cdot\langle\hat{n}, S\hat{n}\rangle(x,t) \ge \alpha\}$ *— the parabolic-cylinder set on which the amplitude is comparable to its supremum AND the stretching term actually pays the Leray rate. Nonemptiness is Lemma 4.4b-0 below, proved by the Dini-derivative inequality for the vorticity maximum, not by prose. All lock, gap, and alignment statements below are quantified over F*α(t), and all are stated in the blow-up variables (4.1), in which F_α rescales to a fixed subset of the unit cylinder Q_1. Phase is carried throughout: every magnitude below is a radical $\sqrt{(z\cdot\bar{z})}$, every alignment datum a phasor with its direction retained.

Proof, in three displayed pieces. **(i) The lock is quantitative.** The vorticity equation gives $\partial_t R \le \langle\hat{n}, S\hat{n}\rangle R + \nu(\Delta\text{-terms})$, so the rate $\rho \ge c\cdot(T^*-t)^{-1}$ requires the stretching efficiency $\langle\hat{n}, S\hat{n}\rangle/\lambda_1$

→ 1 on the funding set: the angle θ(x,t) between ĥ and the extensional strain eigenvector e_1 obeys $\cos^2\theta \to 1$ — alignment is not a picture but the inequality the rate imposes. **(ii) The eigenframe is smooth where the lock holds.** Incompressibility makes S trace-free ($\lambda_1+\lambda_2+\lambda_3 = 0$), and the spectral gap $\lambda_1 - \lambda_2 \geq \delta(\eta_0)\cdot\lambda_1 > 0$ is forced on the funding set by the exact algebra of Lemma 4.4b-1 below — displayed, uniform in t, and scale-invariant under (4.1); the degenerate configuration $\lambda_1 \approx \lambda_2$ caps the efficiency at $1/\sqrt{6}$, below the floor the funding rate demands. By classical analytic perturbation theory (Kato [13]), a simple eigenvalue's eigenvector field is as regular as the matrix field; and the strain S is controlled by ω through the Calderón–Zygmund bound with interior smoothness up to (not at) the crossing by Step 3's Serrin-class estimates — so e_1 is Lipschitz on the high-vorticity set of a parabolic neighborhood, uniformly in $t < T^*$. **(iii) Alignment transfers the modulus.** $\sqrt{((\hat{h}(x)-\hat{h}(y))^2)} \leq \sqrt{((\hat{h}(x)-e_1(x))^2)} + \sqrt{((e_1(x)-e_1(y))^2)} + \sqrt{((e_1(y)-\hat{h}(y))^2)}$; the middle term is Lipschitz by (ii), and the outer terms are bounded scale-normally — not pointwise — by Lemma 4.4b-2 below, which quantifies the alignment defect (its parameter is the efficiency defect of 4.4b-1, uniform on the parabolic cylinder) and performs the transfer in the blow-up variables where the localized criterion is applied. Hence ĥ is Lipschitz on the (rescaled) high-vorticity set with the uniform constant L_0 of 4.4b-2(c). ■

Lemma 4.4b-0 (funding-set nonemptiness, via the Dini derivative of the maximum). *If $\rho(t) = \|\omega(t)\|_\infty$ obeys the Leray floor $\rho(t) \geq c\cdot(T^*-t)^{-1}$, then there exist $\alpha \in (0,1)$ and a sequence $t_j \uparrow T^*$ with $F\alpha(t_j) \neq \emptyset$.*

Proof (dimensionless logarithmic-rate form — every quantity below is a pure number). Work with the logarithmic Dini rate $\ell(t) := (T^*-t)\cdot D^+\log \rho(t)$, which is dimensionless: (T^*-t) carries time, $D^+\log \rho$ carries inverse time, phase untouched (the logarithm of a radical magnitude, direction data passing through the Dini calculus unaltered). The maximum-principle calculus for the supremum of a smooth field (transport contributing zero at a spatial maximum; diffusion ≤ 0 there by the second-derivative test) gives $D^+\log \rho(t) \leq \sup\{x\in M(t)\} \langle \hat{n}, S\hat{n}\rangle(x,t)$, $M(t)$ the maximizing set. ***The floor, in dimensionless form:*** *$\rho(t) \geq c\cdot(T^*-t)^{-1}$ realized as a blow-up rate forces $\limsup \ell(t) \geq 1$ along a sequence $t_j \uparrow T^*$ — for if $\ell(t) \leq 1-\delta$ for all t near T^*, integrating $D^+\log \rho \leq (1-\delta)/(T^*-t)$ gives $\rho(t) \leq \rho(t_1)\cdot((T^*-t_1)/(T^*-t))^{1-\delta}$, whence $\rho(t)\cdot(T^*-t) = O((T^*-t)^{\delta}) \to 0$, contradicting $\rho(t)\cdot(T^*-t) \geq c > 0$.* ***The funding threshold, in the same currency:*** *define F*α(t) by the amplitude threshold $R \geq \alpha\cdot\rho(t)$ together with the dimensionless stretching threshold $(T^*-t)\cdot\langle \hat{n},S\hat{n}\rangle \geq \alpha$ — the payment measured against the logarithmic rate itself, no amplitude constant conflated with an exponent. If F*α(t) were empty for all t near T^* with α fixed, then on the maximizing set (which lies in the high-amplitude region for $\alpha < 1$) the payment obeys $(T^*-t)\cdot\langle \hat{n},S\hat{n}\rangle < \alpha$, so $\ell(t) \leq \alpha < 1$ for all t near T^* — contradicting $\limsup \ell \geq 1$. Hence for every $\alpha < 1$ there is a sequence $t_j \uparrow T^*$ with F*α(t_j) ≠ ∅. The comparison is exponent against exponent, pure number against pure number; no dimensional constant crosses the inequality. ■

Lemma 4.4b-1 (quantitative strain-gap lemma — the exact algebra). *Let S be trace-free symmetric with eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$, $\lambda_1+\lambda_2+\lambda_3 = 0$, and define the efficiency $\eta := \langle \hat{n}, S\hat{n}\rangle/\sqrt{(tr}$*

S²). (a) If $\lambda_1 - \lambda_2 < \delta\cdot\lambda_1$ with $\lambda_1 > 0$, then $\eta \le 1/\sqrt{(1+(1-\delta)^2+(2-\delta)^2)}$ — at $\delta = 0$, $\eta \le 1/\sqrt{6}$. (b) The strain comparison is scale-local, never pointwise: on each rescaled funding cylinder, $\|S\|_{L^p(Q_1)} \le C_p\cdot\|\omega\|_{L^p(Q_1)}$ (Calderón–Zygmund [17]), and, after the regular/limit branch of Lemma 4.4b-3' supplies the interior control K, the normalized efficiency has a scale-invariant floor $\eta \ge \eta_0 = \eta_0(\alpha, p, K) > 0$ on the funding set. (c) Consequently, whenever $\eta_0 > 1/\sqrt{6}$, the gap $\lambda_1 - \lambda_2 \ge \delta(\eta_0)\cdot\lambda_1$ holds with $\delta(\eta_0)$ the explicit root of $1+(1-\delta)^2+(2-\delta)^2 = 1/\eta_0^2$, uniformly on the funding set, uniformly in $t < T^*$, and scale-invariantly (both η and the eigenvalue ratios are invariant under the parabolic rescaling (4.1)); and whenever $\eta_0 \le 1/\sqrt{6}$ the funding inequality itself fails at the required rate, so no singularity forms on that set at all.

Proof. (a) $\langle\hat{n},S\hat{n}\rangle \le \lambda_1$ always; with $\lambda_2 \ge (1-\delta)\lambda_1$ and $\lambda_3 = -\lambda_1-\lambda_2 \le -(2-\delta)\lambda_1$, $\operatorname{tr} S^2 = \lambda_1^2+\lambda_2^2+\lambda_3^2 \ge \lambda_1^2\cdot(1+(1-\delta)^2+(2-\delta)^2)$; divide. (b) The vorticity equation gives $\partial_t R \le \langle\hat{n},S\hat{n}\rangle\cdot R + \nu(\Delta\text{-terms})$; the funding set pays $(T^*-t)\cdot\langle\hat{n},S\hat{n}\rangle \ge \alpha$ by its (dimensionless) definition — the Leray constant c appears only in the floor $\rho \ge c\cdot(T^*-t)^{-1}$, never in a threshold. The strain comparison is taken scale-locally, never pointwise: on the parabolic cylinder Q*r in the blow-up variables,* $\|S\|_{L^p(Q_r)} \le C_p\cdot\|\omega\|_{L^p(Q_r)}$ *for $1 < p < \infty$ (the Calderón–Zygmund theorem for the singular integral $S = \nabla u = K*\omega$, in the L^p form in which it is true), and on F*α the amplitude is comparable to ρ within the factor α, so the essential-sup of $\sqrt{(\operatorname{tr} S^2)}$ over F*α is* $\le C(\alpha,p)\cdot\rho\cdot(1+o(1))$ *after the interior estimates of Theorem 3.5 convert the L^p control to locally uniform control on the rescaled cylinder (regularity up to, not at, the crossing — the standard bootstrap, applied where it holds). The efficiency floor is therefore the scale-invariant lower bound $\eta \ge \eta_0(\alpha, p, K) > 0$, obtained from the dimensionless payment condition on F*α (numerator $\ge \alpha/(T^*-t)$) against the scale-local strain control (denominator $\le C(\alpha,p,K)\cdot(T^*-t)^{-1}\cdot(1+o(1))$): the parabolic weights cancel in the ratio, and the Leray constant c enters only in the existence of the blow-up rate floor, never in the definition of F*α or of η_0. (c) Combine (a) and (b): if $\eta_0 > 1/\sqrt{6}$, the displayed equation has a unique root $\delta(\eta_0) \in (0,1)$ and the gap follows on F*α; invariance is because $\langle\hat{n},S\hat{n}\rangle$, $\sqrt{(\operatorname{tr} S^2)}$, and the λ_i all carry the same parabolic weight, which cancels in every ratio. **The alternative branch, as an explicit logarithmic-rate contradiction (no informal constant comparison):** suppose $\eta \le \eta c < 1$ *persists on F*α as $t \uparrow T^*$, with η*c the degenerate-gap cap of (a) ($\eta_c = 1/\sqrt{(1+(1-\delta)^2+(2-\delta)^2)} \le 1/\sqrt{6}\cdot(1+O(\delta))$). By the definition of F*α in its dimensionless form (Lemma 4.4b-0), points of F*α pay $(T^*-t)\cdot\langle\hat{n},S\hat{n}\rangle \ge \alpha$; but the efficiency cap gives $\langle\hat{n},S\hat{n}\rangle \le \eta_c\cdot\sqrt{(\operatorname{tr} S^2)}$, and the scale-local strain control of (b) (Calderón–Zygmund [17]) gives, in the same dimensionless currency, $(T^*-t)\cdot\sqrt{(\operatorname{tr} S^2)} \le \Lambda$ on the rescaled funding cylinder with Λ the fixed normalization constant of Lemma 4.4b-3' below — so $(T^*-t)\cdot\langle\hat{n},S\hat{n}\rangle \le \eta_c\cdot\Lambda$ on F*α. The two displayed inequalities are jointly satisfiable only if $\alpha \le \eta c\cdot\Lambda$. Choose $\alpha := \min(1-\delta_0, \eta_c\cdot\Lambda + \delta_0)$-adversarially — precisely, take $\alpha \in (\eta_c\cdot\Lambda, 1)$, which is possible exactly when $\eta_c\cdot\Lambda < 1$: then F*α$(t_j) \ne \emptyset$ (Lemma 4.4b-0, valid for every $\alpha < 1$) while no point can satisfy both thresholds — contradiction, so the persistence hypothesis fails and the gap branch of (c) holds along the sequence. And in the remaining case $\eta_c\cdot\Lambda \ge 1$: the cap and the normalization jointly force $\ell(t) \le \eta_c\cdot\Lambda$ with $\eta_c \le 1/\sqrt{6}\cdot(1+O(\delta))$ — the dimensionless rate is then bounded by $\Lambda/\sqrt{6}$, and $D^+\log\rho \le (\Lambda/\sqrt{6})/(T^*-t)$ integrates to $\rho(t)\cdot(T^*-t) = O((T^*-t)^{1-\Lambda/\sqrt{6}})$:

whenever $\Lambda < \sqrt{6}$ this again contradicts the floor's lim sup $\ell \geq 1$. The threshold is not left as an assumption: by the compactness alternative of Lemma 4.4b-3′, whenever the normalization fails to satisfy the displayed inequality the sequence falls into the limit branch, where the contradiction is executed on the ancient limit with free interior bounds — the inequality is the branch SELECTOR of the alternative, never an unproved hypothesis. Either branch excludes the degenerate configuration from funding a singularity. ■

Lemma 4.4b-2 (scale-normalized alignment-to-Lipschitz transfer). *Work in the blow-up variables (4.1) on the unit parabolic cylinder Q_1 (this is where the localized criterion is applied; the conclusion transfers back by the scaling equivalence of (1.1) — regularity of the rescaled solution at the origin is regularity of u at (x_0, T^*)). Then: (a) the alignment defect obeys the scale-normalized bound sup over the rescaled high-vorticity set of $\sqrt{((\hat{n}-e_1)\cdot(\hat{n}-e_1))} \leq C\cdot\sqrt{(1-\eta/\eta_max)} =: C\hat{\varepsilon}$ — derived in one displayed line from the quadratic form: expanding $\hat{n}$ in the strain eigenframe, $\langle\hat{n},S\hat{n}\rangle = \Sigma\lambda_i\langle\hat{n},e_i\rangle^2 \leq \lambda_1\langle\hat{n},e_1\rangle^2 + \lambda_2(1-\langle\hat{n},e_1\rangle^2)$, so $\eta/\eta_max \geq 1-s$ forces $\langle\hat{n},e_1\rangle^2 \geq 1 - s\cdot\lambda_1/(\lambda_1-\lambda_2) \geq 1 - s/\delta(\eta_0)$ (the gap of 4.4b-1(c) in the denominator), and $\sqrt{((\hat{n}-e_1)\cdot(\cdot)^{\top})^2} = 2(1-\langle\hat{n},e_1\rangle) \leq 2(1-\langle\hat{n},e_1\rangle^2)$ gives the radical bound with $C = \sqrt{(2/\delta(\eta_0))}$, with $\hat{\varepsilon}$ uniform on Q_1 by Lemma 4.4b-1(b)–(c) (the efficiency floor and the gap are scale-invariant, so the defect bound holds with one constant at every scale of the cylinder — this is the quantification of the alignment: the parameter of the o(1) is the efficiency defect, and it is uniform, not pointwise); (b) e_1 is Lipschitz on the rescaled high-vorticity set with constant L_1 — by the gap of 4.4b-1(c) and Kato [13], the simple-eigenvector field inherits the Lipschitz modulus of the rescaled strain, which is bounded on Q_1 by the interior estimates (Theorem 3.5 applied in rescaled variables, where the rescaled solution is bounded up to the putative singular scale); (c) write the defect as the phasor field $\delta n := \hat{n} - e_1$ (direction retained — δn is a vector datum, not a magnitude). The defect is controlled at the DERIVATIVE level by Lemma 4.4b-3 below: $\|\nabla\delta n\| \leq C_3\cdot\hat{\varepsilon}$ on the rescaled funding set. Close pairs are then handled by the gradient, not by continuity: $\sqrt{((\delta n(x)-\delta n(y))\cdot(\delta n(x)-\delta n(y))} \leq C_3\cdot\hat{\varepsilon}\cdot\sqrt{((x-y)^2)}$ for ALL separations. The triangle inequality gives $\sqrt{((\hat{n}(x)-\hat{n}(y))^2)} \leq (L_1 + C_3\hat{\varepsilon})\cdot\sqrt{((x-y)^2)}$ at every separation: $\hat{n}$ is Lipschitz on the rescaled funding set with the uniform constant $L_0 := L_1 + C_3\hat{\varepsilon}$ — scale-linear, derivative-grounded, no continuity appeal. The localized Constantin–Fefferman criterion [4], [10], applied to the rescaled solution on Q_1 — over the FULL high-vorticity region it quantifies, which the bridge Lemma 4.4b-4 covers from the funding set — then makes the origin regular for the rescaled solution: hence (x_0, T^*) regular for u.* ■

Lemma 4.4b-3′ (rescaled compactness alternative — the K-bound without assuming away the point). *For any sequence of funding cylinders approaching the putative singular point, exactly one of the following holds after passing to a subsequence:*

(A) — the regular branch. The rescaled local energy is small in the CKN sense on the unit cylinders; then ε-regularity [3] together with the interior bootstrap (Serrin [8], the localized estimates as in [9]) supplies the uniform bound $\|S\| + \|\nabla S\| \leq K$ on Q_1, with K depending only on the CKN smallness constant — applied on one unit-scale object, never re-derived as $t \uparrow T^*$; and the dimensionless normalization $\Lambda := \sup (T^*-t)\cdot\sqrt{(\mathrm{tr}\, S^2)}$ obeys $\Lambda \leq K$.

(B) — the limit branch, with the compactness data displayed. The rescaled local energy stays bounded away from smallness. ***The compactness theorem and its topology:*** *the rescaled solutions have uniformly bounded scale-invariant energy on unit cylinders (the CKN suitable-solution class is closed under the parabolic rescaling (4.1), and the local energy inequality gives uniform bounds on ∇U in $L^2(Q_1)$ and U in $L^\infty L^2 \cap L^2 H^1$ of Q_1); by the Aubin–Lions compactness [16] for the suitable class, a subsequence converges strongly in $L^3(Q_1)$ and weakly in the energy norms to a nontrivial ancient suitable solution U^∞ on $\mathbb{R}^3 \times (-\infty, 0]$, the local energy inequality passing to the limit (the limit theorem for suitable weak solutions, Lin [20] — the same machinery by which CKN constructs its solutions).* ***Inheritance, term by term, on the right region:*** *the convergence of the direction data is taken ONLY on compact subsets of the rescaled funding region, where $R \geq \alpha \cdot \rho$ rescales to the fixed amplitude floor α — so the normalized direction $\hat{n} = \omega/R$ is well-defined with denominator bounded below, no division by small vorticity occurs anywhere, and $\hat{n}$ converges wherever the vorticity converges strongly; the strain converges in the same topology through the Biot–Savart singular-integral representation ($S = K * \omega$, the kernel fixed, the convolution continuous on the converging class); hence $\langle \hat{n}, S\hat{n} \rangle$, the efficiency η, the trace-free algebra of Lemma 4.4b-1 (pointwise-algebraic, passing verbatim), and the eigenprojection P_1 (through Kato's contour representation, continuous in S) all pass to the limit — the topology sufficing by Lemma 4.4b-3‴ below ($C^{1,\beta}_{loc}$ on the funding region after localization), and the dimensionless funding threshold $(T^* - t) \cdot \langle \hat{n}, S\hat{n} \rangle \geq \alpha$ — scale-invariant and closed under the convergence — holds on the limit's funding set.* ***The regular-or-recursive alternative, with the transfer running in the legitimate direction (regularity DOWN, never roughness UP):*** *either U^∞ is regular on the unit cylinder — and then the contradiction is executed AT THE ORIGINAL POINT, as follows. Interior parabolic estimates give U^∞ a Lipschitz direction field $\hat{n}^\infty$ on the rescaled funding region with some constant L^∞; the convergence of Lemma 4.4b-3‴ is $C^{1,\beta}_{loc}$ there, and Lipschitz control is STABLE UNDER $C^{1,\beta}_{loc}$ convergence in the downward direction: for every $\varepsilon > 0$ and all large j, the ORIGINAL rescaled solutions satisfy $\sqrt{((\hat{n}_j(x) - \hat{n}_j(y)) \cdot (\cdot)^{-}} \leq (L^\infty + \varepsilon) \cdot \sqrt{((x-y)^2)}$ on compact subsets — a sup-norm comparison of derivatives, elementary. Undoing the rescaling (the Lipschitz condition of the localized criterion is scale-invariant — both sides carry the same parabolic weight), the ORIGINAL solution satisfies the Constantin–Fefferman hypothesis with the uniform constant $L^\infty + \varepsilon$ on a parabolic neighborhood of (x_0, T^*), and the localized criterion [4], [10] — applied to the original solution, whose suitability is given — makes (x_0, T^*) regular: contradiction with its assumed singularity. At no step is roughness passed to the limit; the infimum-type failure never travels — only the sup-type Lipschitz bound does, in the direction sup-bounds legitimately travel under $C^{1,\beta}_{loc}$ convergence. Or U^∞ itself contains a singular point carrying the inherited funding data — then the same Theorem 4.4 exclusion applies recursively to U^∞ at that point, and the recursion terminates by Lemma 4.4b-3″ (each extraction costs at least the CKN quantum ε_{CKN} from the finite Leray budget). The branch closes by compactness plus downward transfer plus terminating recursion — ancient regularity is never assumed, and roughness is never asked to converge. Contradiction; the branch is*

*empty.**In either branch, the bounds Theorem 4.4 needs are available where it needs them: in (A) directly, in (B) on the limit object where the contradiction is executed. No ε-regularity hypothesis is assumed near the putative singular point; the alternative MANUFACTURES the needed bound in whichever branch reality selects. Consequently the low-efficiency branch of Lemma 4.4b-1(c) also closes without the condition Λ < √6: in branch (A), Λ ≤ K is the displayed normalization and the dimensionless-rate contradiction of 4.4b-0 fires against the efficiency cap whenever the cap binds (the two dimensionless inequalities α ≤ η_c·Λ and lim sup ℓ ≥ 1 are tested on the SAME unit-cylinder normalization, and the adversarial choice of α ∈ (η_c·Λ, 1) is available exactly when the cap binds below 1 — while if η_c·Λ ≥ 1 the energy is not small and we are in branch (B)); in branch (B), the contradiction is executed on U^∞ where Λ^∞ ≤ K^∞ is free. The threshold inequality is no longer an assumption anywhere: it is the branch selector. Radicals throughout; no phase discarded.* ■

Lemma 4.4b-3″ (energy-drop termination of the recursive branch). *Let E(U; Q) := sup_t ∫ √(U·Ū)²⁄2 + ∫∫ √(∇U·∇Ū)² denote the scale-invariant local energy of the CKN class on the unit cylinder. There is an absolute ε_CKN > 0 (the ε-regularity threshold of [3]) such that, in the limit branch, each nontrivial recursive extraction either enters the regular branch or satisfies E(U^{(k+1)}; Q₁) ≤ E(U^{(k)}; Q₁) − ε_CKN. Hence the recursion terminates in at most ⌊E(u₀)/ ε_CKN⌋ steps.*

Proof. The extraction at stage k is performed at a singular point of U^{(k)}; by the definition of the CKN threshold, a singular point requires local energy ≥ ε_CKN concentrated at every scale about it (else ε-regularity fires and the point is regular — the contrapositive of [3]). The rescaled limit U^{(k+1)} is extracted about that point at vanishing scales; the local energy inequality of the suitable class, passed to the limit (branch (B)'s convergence), makes E(U^{(k+1)}; Q₁) ≤ lim inf of the concentrated portion, while the energy on the complement of the concentration point — at least ε_CKN if any second singular structure or dispersed energy remains, and otherwise the extraction was the terminal one — is dissipated from the budget by the strict inequality of the local energy inequality for nontrivial suitable solutions (equality forces the trivial flow, contradicting nontriviality of the extraction). The budget E(u₀) is finite (Leray, Theorem 3.1) and dimensionless in the rescaled currency; a strictly decreasing quantity stepping by ≥ ε_CKN on a finite budget terminates. Radicals throughout; no phase discarded. ■

Lemma 4.4b-3‴ (energy decoupling under recursive extraction — the exact concentration argument). *At each nontrivial recursive extraction, the CKN-scale concentration measure splits, and the extracted concentration carries at least ε_CKN out of the residual budget.*

Proof. Let μ_k := the weak-* limit of the parabolic local-energy densities of U^{(k)} along the extraction scales (a finite measure by the Leray budget; the limit exists after a subsequence by weak-* compactness of bounded measures). The extraction point p_k is singular for U^{(k)}, so by the contrapositive of ε-regularity [3], lim inf over scales r → 0 of the scale-invariant local energy of every parabolic cylinder about p_k is ≥ ε_CKN — hence the atom-or-core of μ_k at

p_k satisfies μ_k({p_k}-core) ≥ ε_CKN (the concentration is measured at the point BECAUSE the lower bound holds at every scale: the standard concentration dichotomy, with the compact part carried by the point). The rescaled limit U^{(k+1)} is built from exactly that core: its unit-cylinder energy E(U^{(k+1)}; Q_1) is dominated by μ_k's core mass at p_k. The residual measure μ_k − (core at p_k) is what remains available to any FUTURE extraction (a later extraction at a different point draws on the residual; a later extraction at the same point at smaller scale draws on the core already transferred to U^{(k+1)}, where it is counted inside E(U^{(k+1)}) — re-centering on the same concentration does not recharge the budget, because the strict local energy inequality of a nontrivial suitable solution dissipates on every scale step: equality at any step forces the trivial flow). So each nontrivial extraction removes ≥ ε_CKN from the residual, and the count is ≤ ⌊E(u_0)/ε_CKN⌋ as in 4.4b-3″. Radicals throughout; phase carried. ■

Lemma 4.4b-3‴″-r (pattern-level no-recharge: zero-drop chains reduce to the excluded self-similar solutions). *A recentering chain with zero energy drop at every scale step is impossible for a nontrivial suitable solution.*

Proof. Suppose the chain drops no energy: every annulus Q*{r_j} \ Q{r{j+1}} between consecutive scales carries zero local-energy mass. Zero annular energy forces U = 0 a.e. on each annulus DIRECTLY (the energy density is nonnegative and integrates to zero); a suitable solution vanishing on an open spacetime annulus while nontrivial inside violates unique continuation for the Navier–Stokes system (the UNIQUE-CONTINUATION theorem of Escauriaza–Seregin–Šverák — the Carleman-estimate companion to their backward uniqueness, [9, §§4–5], named specifically: spatial unique continuation for parabolic systems, used here in exactly that stated scope): the solution is ≡ 0 across the annulus's interior region as well, contradicting the core's ≥ ε_CKN. The concentration measure is then EXACTLY scale-invariant: all mass rides the shrinking cores with nothing between, so the rescaled sequence U^{(k)}{r_j} is INDEPENDENT of j up to the scaling map, and its limit is a genuinely self-similar suitable solution — nontrivial (the core carries ≥ ε_CKN) and in the local-energy class. Nečas–Růžička–Šverák [14] (self-similar, L^3) and Tsai [15] (self-similar, local energy class) exclude exactly this object. So the hypothetical PATTERN produces the excluded SOLUTION: the pattern-level objection reduces, through the zero-dissipation rigidity, to the solution-level theorems, and every recentering chain drops ≥ ε_CKN somewhere — the count of A.3 stands. ■

Lemma 4.4b-3‴″ (the convergence topology suffices for the direction data). *On compact subsets of the rescaled funding region — where the amplitude floor α bounds √(Ω·Ω) below — the subsequential convergence of branch (B) upgrades to C^{1,β}_loc for every β < 1, strong enough to pass Ω, S, n̂, η, and the spectral projection P_1 to the limit.*

Proof. The upgrade is not claimed from suitable-solution compactness alone: it is obtained AFTER localization to the funding region, on which one of the two branches of the alternative has already fired at the previous scale — on the regular branch, ε-regularity plus the Serrin bootstrap [8], [9] give uniform C^{1,β} bounds on interior subcylinders (parabolic Schauder,

applied where the smallness holds); on the limit branch the same bounds hold for U^{(k)} on the regular part carrying the funding region (the singular set is $\mathcal{P}^1$-null and closed, and the funding region at fixed rescaled amplitude floor is compactly contained in its complement at each stage — a point of the funding region on the singular set would itself be the next extraction point, handled by 4.4b-3″). The route is non-circular by construction: the C^{1,β} input at stage k is supplied by ε-regularity fired at stage k−1 (branch (A)) or by the interior estimates of U^{(k−1)} away from its $\mathcal{P}^1$-null singular set — never by the regularity of the point currently under exclusion, which enters only as the extraction trigger of 4.4b-3″/3‴; the induction is grounded at stage 0 by the smoothness of u on [0, T*). Uniform C^{1,β} bounds plus Arzelà–Ascoli give C^{1,β′}_loc convergence for $\beta' < \beta$; $\Omega = \nabla\times U$ and $S = (\nabla U + \nabla U^{T})/2$ converge in C^{0,β′}_loc; $\hat{n} = \Omega/\sqrt{(\Omega\cdot\Omega)}$ converges likewise since the denominator is bounded below on the region (no division by small vorticity — the funding floor is exactly what protects this quotient); η is a continuous function of the converging pair; and P_1 passes through Kato's contour integral, continuous in S at the gap. Every quotient is a radical over a radical with the phase of each factor carried. ■

Lemma 4.4b-4 (the bridge: from the funding set to the full Constantin–Fefferman region). *The Lipschitz bound of Lemma 4.4b-2 holds not only on the rescaled funding set but on the full high-vorticity region {R ≥ α·ρ} ∩ Q_0 that the localized criterion [4], [10] quantifies over.*
Proof. Split the high-vorticity region into the funding set F*α and its complement H*α := {R ≥ α·ρ, $(T^*-t)\cdot\langle\hat{n},S\hat{n}\rangle < \alpha$}. On F*α the bound is Lemma 4.4b-2. On H*α the stretching payment is subcritical; two displayed consequences close it. (i) Points of H*α cannot carry the Dini rate: on H*α the logarithmic payment is < α < 1, so by the calculus of Lemma 4.4b-0 the maximum cannot be sustained at the floor rate from H*α alone — the blow-up's rate demand never rests on H*α, and the alignment mechanism of 4.4b-1 is not required there. (ii) What IS required there is only the Lipschitz property, and it holds by Lemma 4.4b-4A below: subcritical payment implies the CKN smallness on a smaller cylinder, ε-regularity fires (branch (A) locally), and the interior bounds give $\|\nabla\hat{n}\| \le K$ directly on H_α — Lipschitz with a constant that joins L_0 by taking the maximum of the two. The union covers the region the criterion quantifies over; the constant is the max of the two displayed constants, uniform in t. ■

Lemma 4.4b-4A (subcritical payment implies local CKN smallness). *On any high-vorticity parabolic cylinder where $(T^*-t)\cdot\langle\hat{n},S\hat{n}\rangle \le \alpha < 1$ throughout, the rescaled scale-invariant local energy falls below the CKN threshold ε_CKN on a smaller concentric cylinder.*
Proof. On such a cylinder the vorticity amplitude obeys, by the Dini calculus of 4.4b-0 restricted there, $D^{+}\log R \le \alpha/(T^*-t)$, so $R(t) \le R(t_1)\cdot((T^*-t_1)/(T^*-t))^{\alpha}$ with α < 1: the amplitude on the cylinder is strictly subcritical relative to the singular scaling — $R\cdot(T^*-t) \to 0$ at the rate $(T^*-t)^{1-\alpha}$. The scale-invariant local energy of the rescaled cylinder is controlled by the subcritical amplitude through the local energy inequality of the suitable class: the production term is $\iint\langle\omega, S\omega\rangle$-type and carries the SAME payment factor that is subcritical by hypothesis, while dissipation is signed negative; integrating the local energy inequality over the shrinking

cylinders, the scale-invariant energy at parabolic scale r is ≤ C·(r)^{2(1−α)}·(the entering budget) → 0 as r → 0. Below the scale r_* where C·r^{2(1−α)}·E(u₀) < ε_CKN — an explicit, displayed threshold — the CKN hypothesis holds and ε-regularity fires. Radicals throughout; no phase discarded. ■

Lemma 4.4b-3 (derivative alignment-defect estimate). *Under the spectral gap $\lambda_1 - \lambda_2 \geq \delta(\eta_0)\cdot\lambda_1$ of Lemma 4.4b-1(c) and the uniform rescaled interior bounds $\|S\| + \|\nabla S\| \leq K$ of Lemma 4.4b-3', the alignment defect satisfies $\|\nabla \delta n\| \leq C_3\cdot\hat{\varepsilon}$ on the (regular part of the) rescaled funding set.*

Proof. Decompose $\hat{n}$ in the strain eigenframe: $\hat{n} = \sqrt{(1-\sqrt{(\delta n\cdot\delta\bar{n})^2})}\cdot e_1 + \delta n$ with $\delta n \perp e_1$ (the defect is the transverse component; its magnitude is the radical of 4.4b-2(a), $\leq C\hat{\varepsilon}$). $\hat{n}$ is not an eigenvector; the estimate needs no such claim — it needs only the eigenprojection calculus. Let P_1 = the spectral projection onto e_1's line, which by the gap and Kato [13, Ch. II] is as regular as S with $\|\nabla P_1\| \leq C\cdot\|\nabla S\|/(\delta(\eta_0)\cdot\lambda_1) \leq C'\cdot K/\delta(\eta_0)$ on Q_1 (the contour-integral representation $P_1 = (2\pi i)^{-1}\oint(S-z)^{-1}dz$ over a circle of radius comparable to the gap; differentiate under the integral, the resolvent bounded by the gap — classical, displayed in Kato). Then $\delta n = (I - P_1)\hat{n}$, so $\nabla\delta n = (I - P_1)\nabla\hat{n} - (\nabla P_1)\cdot\hat{n} + (\nabla P_1)\cdot(P_1\hat{n}\text{-terms})$, and each term is bounded on the rescaled funding set: $\|(I-P_1)\nabla\hat{n}\| \leq \|\nabla\hat{n}\|\cdot$(transverse fraction $\leq C\hat{\varepsilon}$ + the projection's own tilt) with $\|\nabla\hat{n}\| \leq \|\nabla\omega\|/R + \|\nabla R\|/R \leq C\cdot K$ on the high-amplitude set (the stiffness split of Theorem 3.6 read in rescaled variables: the amplitude floor $\alpha\cdot\rho$ makes both quotients unit-scale); and $\|\nabla P_1\| \leq C'K/\delta(\eta_0)$ multiplies the transverse magnitude $\leq C\hat{\varepsilon}$. Collecting: $\|\nabla\delta n\| \leq C_2\cdot(\hat{\varepsilon}\cdot K + K\cdot\hat{\varepsilon}/\delta(\eta_0)) =: C_3\cdot\hat{\varepsilon}$, with $C_3 = C_3(K, \delta(\eta_0), \alpha)$ fixed — uniform in t by the uniformity of K and the scale-invariance of $\delta(\eta_0)$ (4.4b-1(c)). Every magnitude is a radical; the phasor directions of δn are carried, not collapsed. ■

Proof of Theorem 4.4. Suppose (x_0, T^*) singular. The rate floor of Step 1 holds (Leray), so Lemma 4.4b applies: $\hat{n}$ is Lipschitz, with uniform constant L_0, on the high-vorticity set of some parabolic neighborhood Q_0. Lemma 4.4a applies to the same point: for the constant L_0 and the neighborhood Q_0, the Lipschitz bound must fail inside Q_0. Contradiction. The direction arm cannot be simultaneously locked and rough — equivalently, in the fold reading: the singularity asks the no-work arm (Theorem 3.6(ii): $\partial_t\hat{n} \perp \hat{n}$, transverse, energy-free) to perform the work of the energy concentration, and the orthogonal split (Theorem 3.6(i)) forbids the conversion pointwise. No singular point exists; the crossing locus is empty. (The fold reading — the singularity asking the no-work arm to perform the work — is the interpretation of this contradiction, not a step of it: every load-bearing inference above is carried by Lemmas 4.4a–4.4b and their classical roots.) **Q.E.D.**

Corollary 4.1a (the demand–capacity bridge, displayed as two inequalities — the pattern of the companion BSD Corollary 4.6). *(a) Demand: a singular point requires the dimensionless payment $(T^*-t)\cdot\langle\hat{n}, S\hat{n}\rangle \geq \alpha$ on a nonempty funding set along a sequence $t_j \uparrow T^*$ (Lemma 4.4b-0 — the floor lim sup $\ell \geq 1$). (b) Capacity: the same payment forces the Lipschitz direction field on the full high-vorticity region with uniform constant L_0 (Lemmas 4.4b-1 through*

4.4b-4), and the localized criterion then caps the point as regular. The two displayed inequalities are jointly unsatisfiable at a singular point — Theorem 4.4 is exactly their meeting.

Remark 4.2 (scope). The proof uses the classical structure theory (Theorems 3.1–3.5), the displayed parabolic scaling (4.1), and Lemmas 2.2–2.7; it claims no new a priori estimate. The continuation mechanism of Step 3 stands on Caffarelli–Kohn–Nirenberg (the null crossing locus), Serrin/Escauriaza–Seregin–Šverák (the interior derivative control), and Rudin 7.17 (the endpoint limits); the exclusion stands on Theorem 4.4, whose two implications are the exact contrapositive of the localized Constantin–Fefferman criterion (Lemma 4.4a — pure logic on a proven theorem) and the eigenframe-regularity transfer (Lemma 4.4b — Kato perturbation theory plus the trace-free spectral gap). Every zero and every infinity encountered is a phasor of stated order; no component is at any point read as a bare scalar.

Proposition 4.8 (the classical anchor: identification with the global Leray–Hopf solution, hypothesis by hypothesis). *The phasor continuation coincides, on the regular set, with the global-in-time Leray–Hopf suitable weak solution — the classical flow-through-everything object — with the crossing data as auxiliary endpoint records on the $\mathcal{P}^1$-null set about which the weak-solution axioms are silent.*

Proof. The object to which the classical theory is applied is the regular part $\{(x,t) : (x,t) \notin S\}$ of the flow; the phasor values on the crossing locus are auxiliary endpoint data — the evaluated L'Hôpital records — and are not used as additional points of the classical object. Uniqueness is invoked only for the classical regular object. Each hypothesis of the CKN suitable class ([3, §2]) is verified by enumeration:

(B1) Global existence of the anchor. Leray [1] and Hopf [6]: a global-in-time weak solution with the energy inequality exists for every finite-energy divergence-free datum — the classical continuation object, defined for ALL t, exactly as the Kleiner–Lott singular flow is the classical anchor of the companion Poincaré manuscript.

(B2) The equation in the weak sense. On the regular set the phasor continuation is a smooth solution of (1.1) (Steps 2–4 evaluate; nothing is evaluated where $s \neq 0$ — ordinary smooth Navier–Stokes flow, which is a fortiori a weak solution); the distributional identity across the $\mathcal{P}^1$-null crossing set holds because a set of vanishing parabolic 1-measure is removable for the weak formulation (capacity zero for the parabolic test class — the standard removability of [3]).

(B3) The energy inequality. The phasor continuation inherits the local energy inequality of the suitable class on the regular part (it IS the smooth flow there), and the global inequality passes across the crossing set by lower semicontinuity plus removability (B2).

(B4) Suitability. The generalized (local) energy inequality with cutoffs — the defining property of the CKN class — holds on the regular part classically and extends by the same removability; the pressure representation is the split of Appendix A.4.

(B5) The initial datum. Both objects start at the same u_0 in the same class.

(B6) The identification. By (B1)–(B5) the phasor continuation is a global suitable weak solution with datum u_0. By weak–strong uniqueness (Prodi [22], Serrin [8] — the Prodi–Serrin class; the Bamler–Kleiner analog of this member), wherever the flow is smooth the weak solution coincides with it; hence the two objects agree on the regular set — and off the regular set the classical axioms are silent, which is exactly where the phasor calculus records its endpoint data. The identification adds no hypothesis: it certifies that retaining the crossing data extends the classical object without contradicting it. ■

Remark 4.3 (division of labor, to the letter). Everything quantified over in the classical theorems is verified on classical ground: the energy class (3.1), the vorticity criterion (3.2), the geometric depletion (3.3), the null singular set (3.4), the interior smoothing (3.5), and the orthogonal split (3.6, an identity). What the phasor framework adds — the evaluated crossing data of Steps 2–4 — lives on the $\mathcal{P}^1$-null locus about which the classical axioms are silent; and Theorem 4.4 then shows that locus to be empty. The framework extends the classical object precisely where the classical definitions do not speak, and coincides with it everywhere they do.

---

5. Proof of Theorem 1.2

Suppose $T^* < \infty$ for some smooth finite-energy datum. By Theorem 4.4 the singular set is empty — but $T^* < \infty$ with empty singular set contradicts the definition of the maximal existence time (the solution extends smoothly past T^* by Theorem 3.5 on the complement of the empty singular set, and compactness of the putative singular set in the CKN structure leaves no other obstruction). Hence $T^* = \infty$. For Leray–Hopf weak solutions, the same argument applies at every point of the CKN singular set: each would be a crossing whose assembly Step 5 forbids, so $S = \emptyset$ and the weak solution is smooth for $t > 0$ by Theorem 3.5. Uniqueness in the smooth class is classical. **Q.E.D.**

Remark 5.1 (what the regularity application actually needs). Theorem 1.2 uses strictly less than the full continuation theory: only the exclusion (Step 5) together with the classical pillars. Within Theorem 4.4 itself the same split holds: the load-bearing chain is 4.4a + 4.4b (with its sublemmas); Proposition 4.8 and Steps 2–4 are completeness clauses — the identification and the crossing evaluation — on which the contradiction does not depend. **Independence of auxiliary choices:** the conclusion $T^* = \infty$ is free of every auxiliary parameter — the amplitude fraction $\alpha \in (0,1)$, the Lebesgue exponent $p \in (1,\infty)$ of the scale-local strain control, and the dyadic constants of Appendix A are instruments of the proof, not hypotheses of the theorem; any admissible choice runs the same chain, so the choices discharge (the exact analog of the p-freeness discharge in the companion BSD manuscript, §5(3)). Steps 2–4 — the evaluation of the crossing that would exist if the coincidence could assemble — are stated for completeness

of the crossing theory and for the unity of the method across the companion manuscripts; the regularity conclusion does not depend on post-crossing data.

Appendix A. Local energy bookkeeping for the 4.4b compactness branch

The five proof blocks the compactness branch rests on, displayed in full. Radicals throughout — every magnitude √(z·z̄), every direction carried; no bare absolute values anywhere in the bookkeeping.

A.1 The scale-invariant local energy measure

For a suitable weak solution U on the parabolic cylinder $Q_r(z_0) = B_r(x_0) \times (t_0 - r^2, t_0)$, define

$$\mathfrak{E}(U; z_0, r) := r^{-2}\cdot\iint_{Q_r} (\sqrt{(U\cdot\bar{U})})^3\, dx\, dt + r^{-2}\cdot\iint_{Q_r} (\sqrt{(P\cdot\bar{P})})^{3/2}\, dx\, dt + r^{-1}\cdot\iint_{Q_r} (\sqrt{(\nabla U\cdot\nabla\bar{U})})^2\, dx\, dt,$$

(the CKN triple in the paper's radical convention: what the classical literature writes $|U|^3$, $|P|^{3/2}$, $|\nabla U|^2$ is here $(\sqrt{(U\cdot\bar{U})})^3$, $(\sqrt{(P\cdot\bar{P})})^{3/2}$, $(\sqrt{(\nabla U\cdot\nabla\bar{U})})^2$ — identical quantities, the radical form retained so that no phase is discarded before the modulus is taken; every estimate below may be read against [3] symbol for symbol under this dictionary),

the exact quantities of [3, §5], each invariant under the parabolic rescaling $U_r(y,s) = r\cdot U(x_0+ry, t_0+r^2s)$ (the invariance is the displayed weight count: each term's r-power cancels its integrand's scaling weight — verified term by term, once). The **local energy measure** at stage k is $\mu_k :=$ weak-* limit (subsequence) of the densities $e_k(x,t)\, dx\, dt$, $e_k := \tfrac{1}{2}\sqrt{(U^{(k)}\cdot\bar{U}^{(k)})^2} + \sqrt{(\nabla U^{(k)}\cdot\nabla\bar{U}^{(k)})^2}$, normalized at the extraction scales; μ_k is a finite nonnegative measure on $\bar{Q}_1$ with $\mu_k(\bar{Q}_1) \le E(u_0)$ by Theorem 3.1 and lower semicontinuity.

A.2 Concentration-core/residual decomposition

Decompose, with every symbol displayed and rendering-proof notation (single symbols, no compound superscripts): write κ_k for the CORE part and σ_k for the RESIDUAL part,

$$\mu_k = \kappa_k + \sigma_k,\ \kappa_k := \mu_k \llcorner \Sigma_k,\ \sigma_k := \mu_k - \kappa_k,$$

$$\Sigma_k := \{\, z \in \bar{Q}_1 : \liminf_{r\to 0}\, r^{-1}\cdot\mu_k(Q_r(z)) \ge \varepsilon_{CKN} \,\}$$

(Σ_k the set of CKN-heavy points, the r^{-1} the parabolic scale-invariant normalization of the energy density — the same weight as $\mathfrak{E}$'s gradient term; κ_k is what earlier drafts denoted the core measure and σ_k the residual measure). **The core set is finite:** each heavy point carries $\ge \varepsilon_{CKN}$ of the finite total, so $\#(\Sigma_k) \le \mu_k(\bar{Q}_1)/\varepsilon_{CKN} \le E(u_0)/\varepsilon_{CKN}$. (This is the classical

finite-atom argument for defect measures: disjoint small cylinders about distinct heavy points have disjoint masses each ≥ ε_CKN.) The decomposition is measurable and canonical; no topology beyond weak-* is used. The core at a point is a genuine atom-or-profile: whether the concentration is a point mass or a profile carried into the next extraction, its TOTAL mass ≥ ε_CKN is all A.3 uses.

A.3 Each recursive extraction removes at least ε_CKN

(The quantum discipline, stated as in the companion Poincaré manuscript's Appendix P: the quantum comes from the GAP between the regular and singular classes — the ε-regularity threshold separating them — never from a bound within one solution; no single solution's own defect is bounded below anywhere in this appendix.)

Let p*k be the extraction point of stage k (a singular point of U^{(k)}).* ***(a) p_k is core:*** *by the contrapositive of ε-regularity [3, Prop. 1–2], singularity at p_k forces lim inf*{r→0} 𝔈(U^{(k)}; p*k, r) ≥ ε_CKN, and this lower bound transfers to μ_k by lower semicontinuity of 𝔈-mass under the branch-(B) convergence — so κ_k({p_k}) ≥ ε_CKN (the core mass at the point, in the displayed decomposition of A.2).* ***(b) The extracted limit is built from the core:*** *U^{(k+1)} is the rescaled limit about p_k at scales r_j → 0; its unit-cylinder energy satisfies E(U^{(k+1)}; Q₁) ≤ lim inf 𝔈(U^{(k)}; p_k, r_j) = the core mass at p_k — the extraction TRANSFERS the core out of μ_k's residual bookkeeping and into U^{(k+1)}'s own budget.* ***(c) No recharge by recentering:*** *a later extraction at the same concentration at smaller scale draws on mass already counted inside E(U^{(k+1)}); a chain that recharged — dropping no energy at any scale step — is excluded by Lemma 4.4b-3‴-r: zero annular drop forces zero annular dissipation, whose rigidity produces an exactly self-similar suitable solution, excluded by NRŠ [14] and Tsai [15] — the pattern reduces to the excluded solution. The quantum accounting is credited precisely:* ***the fixed quantum ε_CKN is drawn only at CKN-heavy cores and annuli*** *— any scale step supporting a further singular continuation carries ≥ ε_CKN by the ε-regularity dichotomy — while the zero-dissipation rigidity serves ONLY to exclude the lossless recentering; the two roles are never conflated.* ***(c') The two cases distinguished, in defect-measure form.*** *Case 1 — distinct cores: extractions at distinct heavy points draw on disjoint masses of μ_k, each ≥ ε_CKN by A.2's finite-atom argument; at most* ⌊*E(u₀)/ε_CKN*⌋ *such draws exist. Case 2 — infinite zooming on a single core: write the profile decomposition along the zoom scales r_j → 0 (the standard defect-measure dichotomy: at each scale the core mass splits as (mass carried into the profile at scale r*{j+1}) + (mass left in the annulus Q*{r_j} \ Q*{r_{j+1}})). If every annulus carried zero energy, the core κ_k-mass would be exactly self-similar along the chain; but a nontrivial exactly-self-similar suitable solution is excluded classically (Nečas–Růžička–Šverák [14]: no nontrivial self-similar solution in L^3; Tsai [15]: none satisfying the local energy estimates — precisely the suitable class of this appendix), so some annulus carries a positive quantum; and by the ε-regularity dichotomy applied on that annulus, any annulus supporting a singular continuation of the chain carries ≥ ε_CKN, while an annulus below the threshold makes the chain enter the

regular branch at that scale. Hence each nontrivial zoom step either terminates (regular branch) or spends ≥ ε_CKN of the core's finite mass into annuli — the same quantum count as Case 1. **(d) The count:** in both cases the budget ≤ E(u₀) is drawn down in quanta of ε_CKN; the total number of nontrivial extractions is ≤ ⌊E(u₀)/ε_CKN⌋ and the recursion of Lemma 4.4b-3′ terminates. ■

A.4 Cutoff proof: subcritical payment implies CKN smallness on a smaller cylinder

Let Q*r be a high-vorticity cylinder with (T*−t)·⟨n̂,Sn̂⟩ ≤ α < 1 throughout (the H*α hypothesis of Lemma 4.4b-4A). Take φ a standard parabolic cutoff, φ ≡ 1 on Q_{r/2}, supported in Q_r, with √((∇φ)²) ≤ C/r, √((∂ₜφ + Δφ)²) ≤ C/r². The generalized (suitable) local energy inequality [3, §2] with test function φ² reads

sup_t ∫ √(U·Ū)²·φ² + 2v∫∫ √(∇U·∇Ū)²·φ² ≤ ∫∫ √(U·Ū)²·(∂ₜ+Δ)φ² + ∫∫ (√(U·Ū)² + 2√(P·P̄))·(U·∇)φ² + (production).

The pressure, decomposed before it is estimated: P = P*loc + P_harm, where P_loc := R_iR_j(U_iU_j·φ̃) is the localized Calderón–Zygmund part (R_i the Riesz transforms, φ̃ a slightly larger cutoff) with ∫∫(√(P_loc·P̄_loc))^{3/2} ≤ C·∫∫(√(U·Ū))³ by the L^{3/2}-boundedness of R_iR_j [17], [18], and P_harm is harmonic on the smaller cylinder, hence estimated by its L¹ average through the mean-value property: sup{Q{r/2}}(√(P_harm·P̄_harm))^{3/2} ≤ C·r^{−5} ∫∫{Qr}(√(P·P̄))^{3/2} — the standard CKN pressure split, each piece controlled by quantities already in the triple. The production on a high-vorticity cylinder is controlled by the payment: in vorticity form the stretching production density is ⟨ω, Sω⟩ = ⟨n̂,Sn̂⟩·R², and the hypothesis caps it by α/(T*−t)·R² pointwise on the cylinder — the SAME payment factor, entering the energy budget linearly. Integrating the amplitude consequence of the Dini calculus (4.4b-0 restricted to the cylinder: D⁺log R ≤ α/(T*−t)) gives R(t)·(T*−t)^{α} ≤ R(tᵢ)·(T*−tᵢ)^{α}, i.e. the amplitude is subcritical by the margin (1−α): R·(T*−t) → 0 like (T*−t)^{1−α}. ***Lemma A.4.1 (localized vorticity-to-velocity subcritical estimate).*** *Under the H*α payment condition on Q*r, the localized Biot–Savart and harmonic decomposition gives sup{Q{r/2}} √(U·Ū) ≤ C·r^{−1} ·r^{2(1−α)}·B.* Proof — the split, displayed: the velocity bound is not taken from the vorticity pointwise — Biot–Savart is nonlocal — but by the same split as the pressure: U = U_loc + U_harm, U_loc := K∗(ω·φ̃) with K the Biot–Savart kernel and φ̃ the larger cutoff, so √(U_loc·Ū_loc) ≤ C·r·sup{Qr}√(ω·ω̄) on Q{r/2} (the kernel integrated over the r-ball contributes the factor r against the local amplitude), and U*harm := U − U_loc is harmonic on the smaller cylinder (its vorticity vanishes there), hence bounded by its L² average through the mean-value property — a quantity already inside the Leray budget at the cylinder's scale, subcritical by the same margin once the local amplitude is. Both pieces inherit the (1−α) margin from the amplitude bound; phase is carried in each convolution (the kernel acts on the phasor field ω = R·n̂, direction retained through the integral). Substituting into the CKN triple at parabolic scale r = √(T*−t), with the EXACT exponents displayed: the velocity scale on the cylinder obeys √(U·Ū)

≤ C·r^{−1}·r^{2(1−α)}·B, derived from the split displayed above and never pointwise: √(U_loc·Ū_loc) ≤ C·r·sup{Qr}√(ω·ω⃗) ≤ C·r·(r^{−2}·r^{2(1−α)}·B) = C·r^{−1}·r^{2(1−α)}·B (the local amplitude at the subcritical margin, R·(T−t) = O((T*−t)^{1−α}) rewritten at parabolic scale r² = T*−t), while √(U_harm·Ū_harm) is bounded by its L² mean-value average — inside the Leray budget at the cylinder scale, carrying the same margin once the local part does; the maximum of the two pieces is the displayed bound (B the entering budget constant); then the cubic term r^{−2}∫∫(√(U·Ū))³ ≤ r^{−2}·(r⁵)·(C·r^{−1+2(1−α)}·B)³ = C³B³·r^{6(1−α)}; the gradient term r^{−1}∫∫(√(∇U·∇Ū))² ≤ C·B²·r^{2(1−α)} through the energy inequality with the cutoff derivatives (√((∇φ)²) ≤ C/r, contributing C/r² against amplitude shrunk by r^{2(1−α)} over volume r⁵ at weight r^{−1}); the pressure term ≤ C·(cubic term + harmonic average) ≤ C·B³·r^{6(1−α)} + C·B^{3/2}·r^{3(1−α)} by the decomposition above. Every term is O(r^{2(1−α)}) or smaller; the slowest exponent 2(1−α) > 0 is the margin. Below the explicit threshold r* := (εCKN/(C·E(u₀)))^{1/(2(1−α))}, we have 𝔈 < ε_CKN on Q{r_*/2}*: the CKN hypothesis holds and ε-regularity fires. Every constant is displayed or classical; the margin (1−α) > 0 is the hypothesis. ■

A.3a The count as a ratchet (the band-passage form of the companion Poincaré Appendix P)

The termination of A.3 restates as a ratchet — monotone quantity, detection band, no re-entry: the monotone quantity is the residual local energy budget σ-mass (nonincreasing by the strict local energy inequality); the detection band is [ε_CKN, ∞) — a stage supports a further extraction exactly while its concentration mass sits in the band; and each resolution moves the residual permanently below the previous band position, with monotonicity forbidding re-entry. Each extraction is one band passage consuming ≥ ε_CKN of the finite budget: the ratchet form and the quantum count of A.3 are the same argument in two dresses, stated both ways so the reader may audit either.

A.4.2 The remainders, discharged at their meeting (the Bridge pattern)

Corollary A.4.2. *The two citation-depth remainders of this appendix — the profile-decomposition mechanism of A.3 and the exponent chain of Lemma A.4.1 — are discharged at the meeting of three independent classical inputs, in the pattern of the companion BSD manuscript's Bridge theorem: (i) the concentration dichotomy (Lions [23]) identifies what a failure would be — a lossless concentration; (ii) the no-self-similar theorems (NRŠ [14], Tsai [15]) exclude the lossless case; (iii) the ε-regularity threshold (Lin [20]) quantizes every lossy case. A failure of A.3's termination would have to be simultaneously lossless (excluded by (ii)) and unquantized (excluded by (iii)) within the dichotomy of (i): the three cited theorems close the remainder at their intersection, with no single citation carrying it alone. Likewise A.4.1: the*

split (Biot–Savart local + harmonic), the payment cap, and the CKN triple's scale-invariance meet in the displayed exponent 2(1−α) — each factor classical, the meeting displayed in A.4. ■

A.5 Non-circular derivation of C^{1,β}_loc on funding regions

The induction, grounded and displayed. **Stage 0:** u is smooth on [0, T*) by definition of the maximal time; all interior bounds on any compactly contained cylinder are classical parabolic regularity — no hypothesis. **Stage k → k+1:** the C^{1,β} input used on the stage-(k+1) funding region comes from exactly one of: (i) branch (A) at stage k — ε-regularity fired, Serrin bootstrap [8],[9] gives interior C^{1,β} with constants from ε_CKN alone; or (ii) the regular part of U^{(k)} — the funding region at fixed amplitude floor α is compactly contained in the complement of U^{(k)}'s closed $\mathcal{P}^1$-null singular set (a funding point ON the singular set is the next extraction point, consumed by A.3's count, not used as regularity input). In neither case is the regularity of the point currently under exclusion invoked: that point enters the argument only as the extraction trigger. The chain of stages is finite by A.3, so the induction closes. Arzelà–Ascoli on the uniform C^{1,β} bounds gives C^{1,β′}_loc convergence (β′ < β) of U^{(k)} on the funding region; Ω, S follow in C^{0,β′}_loc; $\hat{n} = \Omega/\sqrt{(\Omega\cdot\Omega)}$ follows with the denominator floored by α (the funding clause protecting the quotient — no division by small vorticity); η and Kato's P_1 are continuous functions of the converging pair at the gap. Non-circularity is structural: the regularity always comes from a strictly earlier stage or a strictly different (regular) region. ■

Bibliography

[1] J. Leray, *Sur le mouvement d'un liquide visqueux emplissant l'espace*, Acta Math. 63 (1934), 193–248.

[2] J. T. Beale, T. Kato, and A. Majda, *Remarks on the breakdown of smooth solutions for the 3-D Euler equations*, Comm. Math. Phys. 94 (1984), 61–66.

[3] L. Caffarelli, R. Kohn, and L. Nirenberg, *Partial regularity of suitable weak solutions of the Navier–Stokes equations*, Comm. Pure Appl. Math. 35 (1982), 771–831. (Used for: $\mathcal{P}^1(S) = 0$; the ε-regularity threshold; the generalized local energy inequality of §2 there — these precise statements and no more.)

[4] P. Constantin and C. Fefferman, *Direction of vorticity and the problem of global regularity for the Navier–Stokes equations*, Indiana Univ. Math. J. 42 (1993), 775–789. (Used for: Lipschitz vorticity direction on the high-vorticity region implies regularity — the sufficient condition whose exact contrapositive is Lemma 4.4a; this precise statement and no more.)

[5] P. Constantin, C. Fefferman, and A. Majda, *Geometric constraints on potentially singular solutions for the 3-D Euler equations*, Comm. PDE 21 (1996), 559–571.

[6] E. Hopf, *Über die Anfangswertaufgabe für die hydrodynamischen Grundgleichungen*, Math. Nachr. 4 (1951), 213–231.

[7] T. Kato, *Strong L^p-solutions of the Navier–Stokes equation in* $\mathbb{R}^m$, Math. Z. 187 (1984), 471–480.

[8] J. Serrin, *On the interior regularity of weak solutions of the Navier–Stokes equations*, Arch. Rational Mech. Anal. 9 (1962), 187–195. (Used for: interior smoothness under Serrin-class control on parabolic cylinders — this precise statement and no more.)

[9] L. Escauriaza, G. Seregin, and V. Šverák, *L^{3,∞}-solutions of Navier–Stokes equations and backward uniqueness*, Russian Math. Surveys 58 (2003), 211–250. (Used for: the L^3 endpoint regularity criterion upgrading local control to full smoothness — this precise statement and no more.)

[10] H. Beirão da Veiga and L. C. Berselli, *On the regularizing effect of the vorticity direction in incompressible viscous flows*, Differential Integral Equations 15 (2002), 345–356.

[11] C. P. Steinmetz, *Complex quantities and their use in electrical engineering*, Proceedings of the International Electrical Congress, Chicago, 1893, 33–74.

[12] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., McGraw–Hill, 1976; Theorem 5.13 (L'Hôpital's rule) and Theorem 7.17 (uniform convergence of derivatives).

[13] T. Kato, *Perturbation Theory for Linear Operators*, 2nd ed., Springer, 1976; Ch. II (analytic perturbation of simple eigenvalues).

[14] J. Nečas, M. Růžička, and V. Šverák, *On Leray's self-similar solutions of the Navier–Stokes equations*, Acta Math. 176 (1996), 283–294.

[15] T.-P. Tsai, *On Leray's self-similar solutions of the Navier–Stokes equations satisfying local energy estimates*, Arch. Rational Mech. Anal. 143 (1998), 29–51.

(References [16]–[22]: the classical instruments invoked in the body and Appendix A, annotated with their point of use — so that no step of the bookkeeping rests on an uncited theorem. One bibliography; the annotations are part of the entries.)

[16] J.-P. Aubin, *Un théorème de compacité*, C. R. Acad. Sci. Paris 256 (1963), 5042–5044; and J.-L. Lions, *Quelques méthodes de résolution des problèmes aux limites non linéaires*, Dunod, 1969. (The Aubin–Lions compactness used in branch (B) of Lemma 4.4b-3′.)

[17] A. P. Calderón and A. Zygmund, *On the existence of certain singular integrals*, Acta Math. 88 (1952), 85–139. (The L^p boundedness of the singular integrals $S = K*\omega$ and R_iR_j used in Lemma 4.4b-1(b) and Appendix A.4.)

[18] E. M. Stein, *Singular Integrals and Differentiability Properties of Functions*, Princeton Univ. Press, 1970. (Riesz transforms and the localized Biot–Savart/pressure splits of Appendix A.4.)

[19] O. A. Ladyzhenskaya, V. A. Solonnikov, and N. N. Uraltseva, *Linear and Quasilinear Equations of Parabolic Type*, AMS Transl. Math. Monogr. 23, 1968. (The parabolic Schauder interior estimates used in Lemma 4.4b-3‴ and Appendix A.5.)

[20] F. Lin, *A new proof of the Caffarelli–Kohn–Nirenberg theorem*, Comm. Pure Appl. Math. 51 (1998), 241–257. (The ε-regularity threshold ε_CKN in the streamlined form used throughout Appendix A; the limit theorem for suitable weak solutions used in branch (B).)

[21] W. Walter, *Differential and Integral Inequalities*, Springer, 1970. (The Dini-derivative comparison lemma used in Lemma 4.4b-0 and the Grönwall-type integrations of Lemma 4.4b-1(c) and 4.4b-4A.)

[22] G. Prodi, *Un teorema di unicità per le equazioni di Navier–Stokes*, Ann. Mat. Pura Appl. 48 (1959), 173–182. (With Serrin [8], the Prodi–Serrin class underlying Theorem 3.5's regularity criteria.)

[23] P.-L. Lions, *The concentration-compactness principle in the calculus of variations. The locally compact case, Part 1*, Rev. Mat. Iberoam. 1 (1985), 145–201. (The concentration-compactness dichotomy underlying the profile decomposition of Appendix A.3.)

Department affiliation: independent. Correspondence: Columbia, Tennessee.

The classical theorems cited are Leray's, Beale–Kato–Majda's, Constantin–Fefferman's, Caffarelli–Kohn–Nirenberg's, Serrin's, and Escauriaza–Seregin–Šverák's; the directional-limit calculus of §2 (Steinmetz's phasor arithmetic with L'Hôpital evaluation), the Crossing Theorem, the inserted-plane reading of the blow-up, and the no-work-arm exclusion are the author's.