

Predictive Fibres and Critical Oseen Transfer for the Three-Dimensional Navier-Stokes Equations

Luca Blanchi

June 2026

Abstract

This article gives a conditional proof strategy for local regularity of the three-dimensional incompressible Navier-Stokes equations. The strategy is presentation-theoretic. Around a putative singular point, one fixes the incoming and lateral information visible before the final core of a parabolic cylinder. Strong solutions generated from this information are the predictors. Suitable weak solutions compatible with the same observations form the fibre. A persistent singularity is then represented as a persistent failure of all admissible predictors to control the fibre across a nested sequence of critical cylinders.

The main theorem is a conditional reduction. Given a finite list of scale-invariant source theorems, it proves that no bad predictive chain can persist. The source theorems are stated explicitly in the body of the paper: they concern the production of the local predictor, the global representation of residual innovations in one Hilbert space, the packet-level rigidity of nearly invisible quadratic interactions, the localization of carrier escape, the synchronization of pressure gauges, and the final predictive Caffarelli-Kohn-Nirenberg gate.

Several parts of the reduction are proved here in full detail. These include the Serrin-critical local Oseen transfer estimate with cutoff and pressure-tail terms, the relaxed minimax-centre construction, the Hilbertian packing estimate for genuinely new normal innovations, the two-mode null classification of the Leray-projected quadratic symbol, the exact pressure absorption of Beltrami fields, the exact two-and-a-half-dimensional reduction, the relative Caffarelli-Kohn-Nirenberg gate once predictor smallness is available, the compactness of the final gate-excess defect, and the fixed-profile cutoff and fixed-drift sub-channels of the Oseen accounting. The source theorem interface then records the exact analytic inputs still needed to turn the conditional reduction into a nonconditional local regularity theorem.

1 Purpose

The incompressible Navier-Stokes equations in three space dimensions are

$$\partial_t u - \Delta u + (u \cdot \nabla)u + \nabla p = 0, \quad \nabla \cdot u = 0.$$

The global regularity problem remains open. Classical local theory studies suitable weak solutions through scale-invariant quantities, pressure decompositions, local energy inequalities, and epsilon-regularity criteria.

This paper proposes a complementary organization. Around a putative singular point, one asks whether the core of a parabolic cylinder is causally predictable from incoming and lateral information by a strong local model. The states compatible with the same incoming observables form a fibre. A singular chain is then interpreted as a sequence of scales at which no bounded predictor captures all states in the fibre.

The presentation-theoretic language is used through analytic statements. The paper separates three levels:

- (i) statements proved here;
- (ii) conditional propositions proved from explicitly stated analytic source theorems;
- (iii) analytic source theorems used as hypotheses in the conditional theorem.

The level of each statement is marked explicitly by its environment: lemma, proposition, theorem, or assumption.

The presentation layer fixes what is visible at a scale, what is invisible inside the fibre, what it means to update a predictor, and which errors are allowed to be counted at the next scale. Standard local regularity estimates enter as analytic engines inside that organization. The strategy turns persistent unpredictability into a finite alternative: either the fibre produces orthogonal innovations, which are square-summable; or it collapses onto a rigid branch; or it exits the current presentation through a registered critical defect. The conditional theorem below is this alternative written as a scale iteration.

2 Main Conditional Statement

The following theorem is the final form of the reduction proved in this paper. The detailed definitions are given in the subsequent sections, and the proof is repeated at the end after all ingredients have been stated.

Theorem 2.1 (Conditional predictive reduction). *Let (u, p) be a suitable weak solution in a neighbourhood of z_0 . Assume that the local predictive presentations centered at z_0 satisfy the source theorems stated below: predictor production and differentiability, global residual representation, packet rigidity, carrier localization, pressure and gauge synchronization, and final predictive epsilon regularity. Assume also that the explicit scale-selection errors are summable. Then there is no infinite nested sequence of cylinders at z_0 on which the predictive defect stays bounded below. Consequently, if failure of the predictive gate is the only remaining obstruction to the Caffarelli-Kohn-Nirenberg epsilon condition, then u is regular at z_0 .*

The theorem is conditional exactly in the sense expressed by its hypotheses. All use of an unproved analytic input occurs through one of those source theorems. The intervening logic—Oseen transfer, minimax selection, Hilbert packing, symbolic null classification, rigid branch reduction, compactness of the gate, and the final contradiction—is proved inside the paper.

Reading guide and dependency map

The proof has four layers.

- (i) The local analytic layer proves the perturbative estimates used inside one scale block: the Oseen transfer lemma, fixed-profile cutoff domination, fixed-drift stability, pressure-tail accounting, relative Caffarelli-Kohn-Nirenberg gating, and compactness of suitable fibres.
- (ii) The Hilbert layer turns genuinely new linear normal defects into a square-summable sequence, provided all scale-dependent defects admit a global residual representation with a Bessel bound.
- (iii) The geometric layer identifies two exact null mechanisms for the projected quadratic symbol–Beltrami and two-and-a-half-dimensional branches—and records the general coherent-branch alternative needed beyond pure modes.

- (iv) The source-theorem layer is the conditional interface. It specifies the PDE inputs that must be proved to construct predictors, synchronize pressures, transport defects, handle moving localization, produce branch predictors, and close the final bad-chain alternative.

Only the fourth layer is assumed. The first three layers are proved as standalone statements and then consumed by the modular regularity theorem.

3 Local Setting

3.1 Parabolic scaling

For $z_0 = (x_0, t_0)$, write

$$Q_R(z_0) = B_R(x_0) \times (t_0 - R^2, t_0).$$

The Navier-Stokes scaling is

$$u^{(R)}(x, t) = Ru(x_0 + Rx, t_0 + R^2t), \quad p^{(R)}(x, t) = R^2p(x_0 + Rx, t_0 + R^2t).$$

A quantity is critical if its normalized value is stable under this scaling.

3.2 Suitable weak solutions

A pair (u, p) is a suitable weak solution on Q if

$$u \in L_t^\infty L_x^2(Q) \cap L_t^2 \dot{H}_x^1(Q), \quad p \in L_{t,x}^{3/2}(Q),$$

the equations hold distributionally, and the local energy inequality holds. For every nonnegative $\phi \in C_c^\infty(Q)$,

$$\begin{aligned} & \int |u(x, t)|^2 \phi(x, t) dx + 2 \int_{-\infty}^t \int |\nabla u|^2 \phi dx ds \\ & \leq \int_{-\infty}^t \int |u|^2 (\partial_s \phi + \Delta \phi) + (|u|^2 + 2p) u \cdot \nabla \phi dx ds. \end{aligned}$$

3.3 Pressure cleaning

Let $B \subset \mathbb{R}^3$ be a smooth ball. For $1 < s < \infty$, the local Helmholtz projection $\mathbb{P}_B$ is

$$\mathbb{P}_B f = f - \nabla \pi_B,$$

where π_B solves the Neumann problem

$$\Delta \pi_B = \nabla \cdot f, \quad \partial_n \pi_B = f \cdot n,$$

with zero average normalization. Standard elliptic estimates give

$$\|\mathbb{P}_B f\|_{L^s(B)} + \|\nabla \pi_B\|_{L^s(B)} \leq C_{s,B} \|f\|_{L^s(B)}.$$

For a local argument, the pressure is decomposed schematically as

$$p = p_{\text{loc}} + p_{\text{harm}} + p_{\text{tail}}.$$

The local part is absorbed by $\mathbb{P}_B$. The harmonic and nonlocal pieces contribute a pressure-tail functional

$$\text{PressTail}(Q_R, Q_r).$$

In an Oseen estimate with residual w , pressure representative π , and cutoff ζ , this notation means a bound for the remaining pressure pairing

$$\left| \int_{Q_R} \pi_{\text{tail}} w \cdot \nabla(\zeta^2) dx dt \right|,$$

or for the corresponding supremum over the admissible energy tests in the same pressure gauge. Thus the local projection removes only the local gradient part; the nonlocal pressure contribution is paid explicitly.

4 Predictive Presentations

4.1 Incoming data

Fix $0 < \theta < 1$ and $0 < \lambda < 1$, with $\theta^2 < \lambda$. For $Q_R(z_0)$, define the incoming region

$$\mathcal{I}^-(Q_R) = B_R(x_0) \times (t_0 - R^2, t_0 - \lambda R^2) \cup (B_R(x_0) \setminus B_{\theta R}(x_0)) \times (t_0 - R^2, t_0).$$

It contains a past slab and a lateral collar, but excludes the final core.

4.2 Predictors and fibres

A causal strong predictor on Q_R is a strong solution g constructed from data measurable on $\mathcal{I}^-(Q_R)$, after the chosen pressure-cleaning procedure. The critical predictor budget is

$$\|g\|_{L_t^q L_x^p(Q_R)} \leq L, \quad \frac{2}{q} + \frac{3}{p} = 1, \quad p > 3.$$

Additional budget components may record pressure-tail size, observable cost, localization cost, and conditioning of the finite observational Gramian.

Let $\mathcal{O}$ be a finite family of incoming observables. The invisible fibre associated with (u, p) is the class

$$\mathcal{S}_{L, \mathcal{O}}^-(Q_R)$$

of suitable weak states (v, q) in Q_R with budget at most L and the same incoming observations as (u, p) :

$$\mathcal{O}(v, q) = \mathcal{O}(u, p).$$

Let $\mathcal{P}_{L, \mathcal{O}}^-(Q_R)$ denote the corresponding class of causal strong predictors.

Definition 4.1 (Local predictive defect). *Let d_{Q_R} be a pressure-cleaned distance on the core $Q_{\theta R}$. The minimax predictive defect is*

$$\mathfrak{D}_{L, \mathcal{O}}(Q_R) = \inf_{g \in \mathcal{P}_{L, \mathcal{O}}^-(Q_R)} \sup_{(v, q) \in \mathcal{S}_{L, \mathcal{O}}^-(Q_R)} d_{Q_R}((v, q), g).$$

When the infimum is taken in the closed convex hull of the predictor class, the distance from the relaxed centre to realizable predictors is recorded as a realization defect $\mathfrak{V}(Q_R)$.

A bad predictive chain at z_0 is a nested sequence $Q_{R_k}(z_0)$, $R_k \downarrow 0$, for which $\mathfrak{D}_{L, \mathcal{O}_k}(Q_{R_k})$ stays bounded below by a fixed positive number.

5 Critical Local Oseen Transfer

The following estimate is the local perturbative engine. It transports a residual around a Serrin-critical strong predictor from a cylinder to a smaller core, while recording the cutoff, drift, forcing, and pressure-tail costs that must be paid by the source theorem interface.

Lemma 5.1 (Critical local Oseen transfer). *Let $Q_\rho \Subset Q_R$. Let g be divergence-free and*

$$g \in L_t^q L_x^p(Q_R), \quad \frac{2}{q} + \frac{3}{p} = 1, \quad p > 3.$$

Let w be divergence-free and pressure-cleaned, satisfying

$$\partial_t w - \Delta w + (g \cdot \nabla)w + (w \cdot \nabla)g + \nabla \pi = F$$

in Q_R , distributionally. Choose $\zeta \in C_c^\infty(Q_R)$ with $\zeta = 1$ on Q_ρ . Assume that the nonlocal pressure contribution is bounded by PressTail_ζ , in the sense of the pressure tail pairing defined above, and define

$$\begin{aligned} \text{Coll}_{\zeta,g}(w)^2 &= \sup_{\text{initial slab}} \int |w|^2 \zeta^2 dx \\ &\quad + \int_{Q_R} |w|^2 \left(|\partial_t \zeta| + |\Delta \zeta| + |\nabla \zeta|^2 \right) dx dt \\ &\quad + \int_{Q_R} |g| |w|^2 |\nabla \zeta| dx dt. \end{aligned}$$

Then

$$\begin{aligned} &\|w\|_{L_t^\infty L_x^2(Q_\rho)} + \|\nabla w\|_{L_{t,x}^2(Q_\rho)} + \|w\|_{L_{t,x}^{10/3}(Q_\rho)} \\ &\leq C \left(\text{Coll}_{\zeta,g}(w) + \|F\|_{L_t^2 H_x^{-1}(Q_R)} + \text{PressTail}_\zeta \right), \end{aligned}$$

where C depends on $Q_\rho \Subset Q_R$, on p , and on $\|g\|_{L_t^q L_x^p(Q_R)}$.

Proof. Test the equation against $w\zeta^2$. The local pressure projection removes the local pressure gradient. The remaining pressure pairing is bounded by PressTail_ζ . We obtain

$$\frac{1}{2} \frac{d}{dt} \int |w|^2 \zeta^2 + \int |\nabla w|^2 \zeta^2 \leq I_{\text{cut}} + I_g + I_F + I_{\text{press}}.$$

The cutoff term is bounded by the collar part of $\text{Coll}_{\zeta,g}(w)^2$. The forcing term is estimated by duality:

$$|\langle F, w\zeta^2 \rangle| \leq \|F\|_{H^{-1}} \|w\zeta^2\|_{H^1} \leq \frac{1}{4} \|\zeta \nabla w\|_2^2 + C \|F\|_{H^{-1}}^2 + C \|w \nabla \zeta\|_2^2.$$

Since $\nabla \cdot g = 0$,

$$\int (g \cdot \nabla w) \cdot w \zeta^2 = -\frac{1}{2} \int |w|^2 g \cdot \nabla(\zeta^2),$$

which is included in the collar cost.

It remains to estimate the stretching term. Using $\nabla \cdot w = 0$,

$$\int w_i (\partial_i g_j) w_j \zeta^2 = - \int g_j w_i (\partial_i w_j) \zeta^2 - \int g_j w_i w_j \partial_i (\zeta^2).$$

The second term is again a collar term. For the first,

$$\left| \int g_j w_i (\partial_i w_j) \zeta^2 \right| \leq \|g\|_{L^p} \|w \zeta\|_{L^{2p/(p-2)}} \|\zeta \nabla w\|_{L^2}.$$

By Gagliardo-Nirenberg,

$$\|w\zeta\|_{L^{2p/(p-2)}} \leq C \|w\zeta\|_{L^2}^{1-3/p} \|\nabla(w\zeta)\|_{L^2}^{3/p}.$$

Writing

$$E(t) = \int |w|^2 \zeta^2, \quad D(t) = \int |\nabla w|^2 \zeta^2,$$

and absorbing the cutoff contribution into the collar cost, Young's inequality gives

$$|I_g^{\text{int}}| \leq \frac{1}{4} D(t) + C \|g(t)\|_{L^p}^q E(t), \quad q = \frac{2p}{p-3}.$$

The exponent is exactly the Serrin exponent because $2/q + 3/p = 1$.

Thus

$$\frac{d}{dt} E(t) + \frac{1}{2} D(t) \leq C \|g(t)\|_{L^p}^q E(t) + B(t),$$

where $B(t)$ contains only forcing, pressure-tail, and collar terms. Since $\|g(t)\|_{L^p}^q \in L_t^1$, Gronwall yields

$$\sup_t E(t) + \int D(t) dt \leq C \left(\text{Coll}_{\zeta,g}(w)^2 + \|F\|_{L_t^2 H_x^{-1}(Q_R)}^2 + \text{PressTail}_{\zeta}^2 \right).$$

Finally,

$$\|w\|_{L^{10/3}(Q_\rho)} \lesssim \|w\|_{L_t^\infty L_x^2(Q_\rho)}^{2/5} \|\nabla w\|_{L^2(Q_\rho)}^{3/5}.$$

This proves the claim. $\square$

5.1 Fixed-profile collars and the pressure interface

The Oseen estimate becomes a scale iteration only after the cutoff family has been fixed. The useful closed subchannel is the integrated derivative-weighted cutoff collar for a single anchored dyadic family.

Proposition 5.2 (Fixed-profile dyadic cutoff domination). *Fix $0 < a < b < 1$, $0 < \lambda < 1$, a centre $z_* = (x_*, t_*)$, and $r_k = 2^{-k}R$. Let*

$$\zeta_k(x, t) = \phi\left(\frac{x - x_*}{r_k}\right) \psi\left(\frac{t_* - t}{r_k^2}\right),$$

where $\phi \in C_c^\infty(B_b(0))$, $\phi = 1$ on $B_a(0)$, and ψ is a fixed smooth past-time cutoff. Put

$$W_k = |\partial_t \zeta_k| + |\Delta \zeta_k| + |\nabla \zeta_k|^2.$$

Then, inside the fixed outer cylinder,

$$\sum_{k \geq K} W_k(x, t) \leq C \left(|x - x_*|^{-2} + R^{-2} \right)$$

for a constant depending only on the profiles and on a, b, λ . Hence, for every local energy field w ,

$$\sum_{k \geq K} \int |w|^2 W_k \leq C \left(\int |\nabla w|^2 + R^{-2} \int |w|^2 \right)$$

on the outer cylinder, after the standard local Hardy inequality.

Proof. The derivative bounds give

$$W_k \leq Cr_k^{-2} \left(\mathbf{1}_{A_k^{\text{lat}}} + \mathbf{1}_{A_k^{\text{time}}} \right),$$

where A_k^{lat} is the spatial annulus $ar_k \leq |x - x_*| \leq br_k$ during the active time interval, and A_k^{time} is the past time floor on which the time cutoff varies. On A_k^{lat} , r_k is comparable to $|x - x_*|$, and a fixed point belongs to only boundedly many dyadic annuli. Thus the sum of the lateral weights is bounded by $C|x - x_*|^{-2}$.

On A_k^{time} , the condition $t_* - t \simeq r_k^2$ allows only boundedly many dyadic indices for each fixed time. For those indices, the spatial support lies in $|x - x_*| \leq br_k$; the same dyadic comparison gives the pointwise bound away from $x = x_*$, which is sufficient for the integrated estimate. A possible outer cutoff contributes only CR^{-2} . Integrating against $|w|^2$ and applying the local Hardy inequality gives the displayed bound. $\square$

Proposition 5.3 (Fixed-drift collar stability). *Let g be fixed on a block of nested cylinders and suppose*

$$\|g\|_{L_t^q L_x^p(Q_R)} \leq L, \quad \frac{2}{q} + \frac{3}{p} = 1, \quad p > 3.$$

For every perturbation w satisfying the local Oseen equation on the block, the drift-dependent part of the collar is controlled by the same annular perturbation norm used in the Oseen transfer:

$$\int_{Q_R} |g| |w|^2 |\nabla \zeta_R| \leq C(L, \theta, p) \|w\|_{L_t^\infty L_x^2(Q_R \setminus Q_{\theta R})}^{2(1-3/p)} \|\nabla w\|_{L^2(Q_R \setminus Q_{\theta R})}^{6/p}.$$

Thus a fixed drift does not introduce a new scale owner: it is absorbed by the same annular energy and dissipation already paid in the Oseen estimate.

Proof. Use Hölder exactly as in the preceding proposition and interpolate

$$\|w\|_{L_x^{2p/(p-2)}} \leq C \|w\|_{L_x^2}^{1-3/p} \|\nabla w\|_{L_x^2}^{3/p}$$

on each time slice. Integrating in time with the Serrin conjugate exponent gives the claimed bound. Since g and the cutoff profile are fixed before the block is selected, the term is a predetermined Oseen collar contribution. $\square$

Proposition 5.4 (Pressure-pair source interface). *For a fixed dyadic Oseen chart, the pressure pairing*

$$\text{PressPair}_k = \left| \int \pi_k^0 w_k \cdot \nabla(\zeta_k^2) \right|$$

reduces, after parabolic normalization and Young's inequality, to critical pressure and velocity collar densities:

$$\text{PressPair}_k \leq \varepsilon V_k + C_\varepsilon P_k,$$

where

$$V_k = r_k^{-2} \int_{A_k} |w_k|^3, \quad P_k = r_k^{-2} \int_{A_k} |\pi_k^0|^{3/2}.$$

The fixed-profile Hardy estimate above does not by itself imply summability of $\sum_k P_k$ or $\sum_k V_k$. Those two normalized collar sums are part of the pressure/localization source theorem interface.

Proof. The first displayed inequality is the usual local pressure estimate on the normalized chart, written back in scale-invariant variables. It separates the pressure pairing into a velocity collar and a pressure collar.

The obstruction is purely scaling. Let A_k be disjoint dyadic annuli and choose a nonnegative density m with

$$\int_{A_k} m = r_k^2/k.$$

Then $\sum_k \int_{A_k} m < \infty$, while

$$\sum_k r_k^{-2} \int_{A_k} m = \sum_k \frac{1}{k} = \infty.$$

Taking $m = |\pi|^{3/2}$, or $m = |w|^3$, shows that finite local critical integrability does not produce normalized dyadic collar summability. Therefore the pressure-pair term cannot be closed by the fixed cutoff Hardy argument alone; it needs an independent pressure/localization theorem, a gate, or a registered source alternative. $\square$

6 Minimax Centres

The minimax construction used by the predictive framework is a Hilbert-space statement. The PDE content lies in verifying that the chosen predictor classes satisfy the hypotheses.

Lemma 6.1 (Relaxed minimax centres). *Let $\mathcal{H}$ be a Hilbert space. Let $P \subset \mathcal{H}$ be nonempty, bounded, closed, and convex, and let $S \subset \mathcal{H}$ be bounded. Define*

$$R(g) = \sup_{v \in S} \|v - g\|_{\mathcal{H}}, \quad g \in P.$$

Then R attains its minimum on P .

Proof. The map $g \mapsto \|v - g\|_{\mathcal{H}}$ is convex and weakly lower semicontinuous for each v . Hence R , as a supremum of weakly lower semicontinuous functions, is weakly lower semicontinuous. Since Hilbert spaces are reflexive, a bounded closed convex subset is weakly compact. Therefore R attains its minimum on P . $\square$

Remark 6.2. *In the Navier-Stokes application, P is the weakly closed convex hull of admissible strong predictors at a fixed finite observation budget. The gap between a relaxed minimax centre and an actual strong predictor is exactly the realization defect $\mathfrak{V}$. This defect must be paid for in the scale-transfer inequality.*

7 Innovation Packing

7.1 Abstract packing

Let $\mathfrak{H}$ be a Hilbert space of normalized adjoint observables. Suppose that a sequence of normal defects $\mathcal{N}_k \in \mathfrak{H}^*$ has normalized innovative witnesses $e_k \in \mathfrak{H}$. The filtration condition is

$$e_k \perp \mathcal{P}_{k-1}, \quad \mathcal{P}_k = \overline{\mathcal{P}_{k-1} + \text{span}\{e_k\}},$$

so (e_k) is orthonormal on innovative scales.

Lemma 7.1 (Hilbertian innovation packing). *Assume that there exist $\mathcal{R} \in \mathfrak{H}^*$, errors ρ_k , and $c > 0$ such that*

$$\langle \mathcal{N}_k, e_k \rangle = \langle \mathcal{R}, e_k \rangle + \rho_k, \quad \sum_k |\rho_k|^2 < \infty,$$

and

$$|\langle \mathcal{N}_k, e_k \rangle| \geq c \|\mathcal{N}_k\|_{\mathfrak{H}^*}.$$

Then

$$\sum_k \|\mathcal{N}_k\|_{\mathfrak{H}^*}^2 \leq C_c \left(\|\mathcal{R}\|_{\mathfrak{H}^*}^2 + \sum_k |\rho_k|^2 \right).$$

Proof. By Bessel's inequality,

$$\sum_k |\langle \mathcal{R}, e_k \rangle|^2 \leq \|\mathcal{R}\|_{\mathfrak{H}^*}^2.$$

Using the coefficient representation,

$$\sum_k |\langle \mathcal{N}_k, e_k \rangle|^2 \leq 2 \sum_k |\langle \mathcal{R}, e_k \rangle|^2 + 2 \sum_k |\rho_k|^2.$$

The witness lower bound then gives the stated estimate. $\square$

Proposition 7.2 (Bessel representation with moving charts). *Let H_{def} be a Hilbert space of registered defect vectors. For each scale k , let $d_k \in H_{\text{def}}$, let*

$$i_k = d_k - P_{\text{span}\{d_j : j < k\}} d_k,$$

and, when $i_k \neq 0$, set $e_k = i_k / \|i_k\|$. Suppose that there are a Hilbert space $\mathcal{K}$, a residual vector $\mathcal{R} \in \mathcal{K}$, bounded operators $A_k : \mathcal{K} \rightarrow H_{\text{def}}$, and errors ρ_k such that

$$d_k = A_k \mathcal{R} + \rho_k, \quad \sum_k \|\rho_k\|_{H_{\text{def}}}^2 < \infty,$$

and such that the pulled-back innovation tests satisfy the Bessel bound

$$\sum_k |\langle f, A_k^* e_k \rangle_{\mathcal{K}}|^2 \leq B \|f\|_{\mathcal{K}}^2 \quad \text{for every } f \in \mathcal{K}.$$

Then

$$\sum_k \|i_k\|_{H_{\text{def}}}^2 \leq 2B \|\mathcal{R}\|_{\mathcal{K}}^2 + 2 \sum_k \|\rho_k\|_{H_{\text{def}}}^2.$$

Proof. For nonzero i_k , the vector e_k is orthogonal to all previous d_j , so

$$\|i_k\| = \langle d_k, e_k \rangle = \langle \mathcal{R}, A_k^* e_k \rangle + \langle \rho_k, e_k \rangle.$$

After squaring, summing, and using $\|e_k\| = 1$, the Bessel bound gives the displayed estimate. Zero-innovation indices contribute nothing. $\square$

7.2 PDE meaning

The abstract packing lemma applies when the transported defect vectors already live in one fixed Hilbert space. The moving-chart proposition is the form needed for Navier-Stokes: scale changes, pressure gauges, carrier frames, and cutoffs may vary with k , and the maps A_k must satisfy a genuine Bessel or finite-overlap estimate. The PDE source theorem supplies exactly this representation and the square-summability of the error vectors ρ_k .

8 The Bilinear Symbol and Null Fibres

The next lemma is a finite-dimensional calculation for the Leray-projected quadratic symbol. It explains why large invisible fibres with very small quadratic leakage should be close to special geometric branches, but by itself it treats only pure modes.

Lemma 8.1 (Two-mode null classification). *Let*

$$u = ae^{i\xi \cdot x}, \quad v = be^{i\eta \cdot x},$$

with

$$a \cdot \xi = 0, \quad b \cdot \eta = 0.$$

Set $k = \xi + \eta$, $k \neq 0$, and consider the symmetrized projected symbol

$$\mathfrak{b}((\xi, a), (\eta, b)) = \mathbb{P}_k[(a \cdot \eta)b + (b \cdot \xi)a].$$

Write

$$a = \alpha k + A, \quad b = \beta k + B, \quad A \perp k, \quad B \perp k.$$

Then $\mathfrak{b} = 0$ if and only if

$$\alpha B + \beta A = 0.$$

The nonzero solutions fall into the following branches:

- (i) the shear-null branch $\alpha = \beta = 0$;
- (ii) the balanced equal-radius branch $\alpha\beta \neq 0$, which forces $|\xi| = |\eta|$.

Proof. Since $a \cdot \xi = 0$ and $b \cdot \eta = 0$,

$$a \cdot \eta = a \cdot (\xi + \eta) = a \cdot k, \quad b \cdot \xi = b \cdot (\xi + \eta) = b \cdot k.$$

Therefore

$$(a \cdot \eta)b + (b \cdot \xi)a = (a \cdot k)b + (b \cdot k)a = |k|^2(\alpha b + \beta a).$$

Substituting the decompositions gives

$$\alpha b + \beta a = 2\alpha\beta k + \alpha B + \beta A.$$

The projection $\mathbb{P}_k$ kills the component parallel to k , so the projected symbol vanishes exactly when $\alpha B + \beta A = 0$.

If $\alpha = \beta = 0$, both polarizations are perpendicular to k , giving the shear-null branch. If exactly one of α, β is zero, then the relation forces one polarization to vanish, so there is no nonzero two-mode interaction. If $\alpha\beta \neq 0$, then $B = -(\beta/\alpha)A$. The incompressibility condition for a gives

$$A \cdot \xi = -\alpha k \cdot \xi.$$

Since $A \perp k$, we have $A \cdot \eta = -A \cdot \xi = \alpha k \cdot \xi$. The incompressibility condition for b gives

$$0 = b \cdot \eta = \beta k \cdot \eta + B \cdot \eta = \beta k \cdot \eta - \frac{\beta}{\alpha} A \cdot \eta = \beta(k \cdot \eta - k \cdot \xi).$$

Thus $k \cdot \eta = k \cdot \xi$, which is

$$|\eta|^2 + \xi \cdot \eta = |\xi|^2 + \xi \cdot \eta.$$

Hence $|\xi| = |\eta|$. □

Remark 8.2. *The next analytic step is a quantitative packet-level null-graph rigidity theorem: it should promote this exact two-mode calculation to localized parabolic wave packets and identify when a large fibre with small leakage is close to a Beltrami, shear, or carrier-escape branch.*

9 Rigid Branches

Proposition 9.1 (Beltrami absorption). *Suppose U is divergence-free and*

$$\nabla \times U = \lambda U.$$

Then

$$(U \cdot \nabla)U = \nabla \left(\frac{|U|^2}{2} \right).$$

Thus the quadratic term is pure pressure and is eliminated by Leray projection or by local pressure-cleaning.

Proof. Use the identity

$$(U \cdot \nabla)U = \nabla \left(\frac{|U|^2}{2} \right) - U \times (\nabla \times U).$$

If $\nabla \times U = \lambda U$, then $U \times (\lambda U) = 0$. □

Proposition 9.2 (Two-and-a-half-dimensional reduction). *Let*

$$U(x, t) = V(y, t) + W(y, t)n,$$

where y belongs to a fixed plane Π , $V(y, t) \in \Pi$, $n \perp \Pi$, and $\nabla_\Pi \cdot V = 0$. Then the three-dimensional Navier-Stokes system reduces to

$$\partial_t V + V \cdot \nabla_\Pi V - \Delta_\Pi V + \nabla_\Pi P = 0,$$

$$\partial_t W + V \cdot \nabla_\Pi W - \Delta_\Pi W = 0.$$

Thus the planar part is a two-dimensional Navier-Stokes solution and the transverse component is a passive transported-diffused scalar.

Proof. All fields are constant in the n -direction. The in-plane component of the equation is exactly the two-dimensional incompressible Navier-Stokes equation for V . The n -component contains no pressure gradient and gives the passive scalar equation for W . □

10 Relative Caffarelli-Kohn-Nirenberg Gates

The Caffarelli-Kohn-Nirenberg criterion is used here in a relative form. The predictor is allowed to carry the large, but controlled, Serrin component; the unknown part is the perturbation around that predictor.

For a cylinder $Q_R(z_0)$, write

$$C_R(u, p; z_0) = R^{-2} \int_{Q_R(z_0)} |u|^3 dx dt + R^{-2} \int_{Q_R(z_0)} |p - (p)_{B_R}(t)|^{3/2} dx dt.$$

Proposition 10.1 (Relative Caffarelli-Kohn-Nirenberg criterion). *Let g be a strong divergence-free solution on Q_1 with pressure π_g , and suppose*

$$\|g\|_{L_t^q L_x^p(Q_1)} \leq L, \quad \frac{2}{q} + \frac{3}{p} = 1, \quad p > 3.$$

Let (u, p) be suitable on Q_1 , set $w = u - g$, and write

$$p - \pi_g = \pi_w + \pi_{\text{tail}}$$

with the local pressure part normalized by zero spatial mean on balls. For every $\varepsilon_{\text{CKN}} > 0$ there exist $\theta = \theta(L, p, \varepsilon_{\text{CKN}}) \in (0, 1)$ and $\delta = \delta(L, p, \varepsilon_{\text{CKN}}) > 0$ such that, if

$$\|w\|_{L^{10/3}(Q_1)} + \|\pi_w\|_{L^{3/2}(Q_1)} + \|\pi_{\text{tail}}\|_{L^{3/2}(Q_1)} \leq \delta,$$

then

$$C_\theta(u, p; 0) \leq \varepsilon_{\text{CKN}}.$$

In particular, choosing ε_{CKN} below the classical Caffarelli-Kohn-Nirenberg threshold gives regularity in a smaller cylinder.

Proof. By Hölder and the Serrin bound, $g \in L^3(Q_\theta)$ with

$$\theta^{-2} \int_{Q_\theta} |g|^3 \leq C(L, p) \theta^\alpha$$

for some $\alpha > 0$, because $p > 3$ gives subcritical local integrability on shrinking cylinders after the Serrin normalization. The associated pressure π_g obeys the same shrinking estimate after subtracting its spatial mean, by local Calderon-Zygmund estimates and the harmonic-pressure decomposition. Choose θ so these two predictor contributions are at most $\varepsilon_{\text{CKN}}/4$.

For the perturbation, interpolation gives

$$\theta^{-2} \int_{Q_\theta} |w|^3 \leq C \|w\|_{L^{10/3}(Q_1)}^3 \theta^{\beta-2}$$

with a fixed positive exponent β after the parabolic volume factor is accounted for. The pressure perturbation contributes

$$\theta^{-2} \int_{Q_\theta} |\pi_w|^{3/2} \leq \theta^{-2} \|\pi_w\|_{L^{3/2}(Q_1)}^{3/2},$$

and the same estimate applies to π_{tail} . With θ fixed, choose δ so the three perturbative contributions are together at most $\varepsilon_{\text{CKN}}/2$. Summing predictor and perturbation pieces gives the displayed bound for $C_\theta(u, p; 0)$. $\square$

10.1 Final gate-excess defect

The presentation uses a final gate to decide whether a scale is already within the relative Caffarelli-Kohn-Nirenberg regime. The following compactness fact is the precise topological statement needed in the iteration.

Definition 10.2 (Gate-excess defect). *Fix a bounded presentation chart, a predictor class P , a suitable weak fibre S , a core cylinder Q_θ , and a gate threshold $\varepsilon_{\text{gate}} > 0$. For $g \in P$ and $(v, q) \in S$, define*

$$\Gamma((v, q), g) = [C_\theta(v, q \mid g) - \varepsilon_{\text{gate}}]_+,$$

where $C_\theta(v, q \mid g)$ denotes the relative Caffarelli-Kohn-Nirenberg quantity of $v - g$ with the corresponding pressure difference and pressure-tail allowance. The gate-excess defect of the chart is

$$\mathfrak{G} = \inf_{g \in P} \sup_{(v, q) \in S} \Gamma((v, q), g).$$

Proposition 10.3 (Gate-excess dichotomy). *Assume P is compact in the strong topology used for predictors, S is compact in the suitable-weak topology, and the relative Caffarelli-Kohn-Nirenberg functional is lower semicontinuous on $S \times P$. Then:*

- (i) $\mathfrak{G} = 0$ if and only if every state in the fibre is captured by some predictor at the final gate;
- (ii) if the final gate is closed and no predictor captures the whole fibre, then $\mathfrak{G} > 0$;
- (iii) for each $\delta > 0$, the condition $\mathfrak{G} \geq \delta$ is stable under compact convergence of charts.

Proof. The map

$$g \mapsto \sup_{(v, q) \in S} \Gamma((v, q), g)$$

is lower semicontinuous as a supremum of lower semicontinuous functions. Since P is compact, the infimum is attained. If the attained value is zero, the chosen predictor has zero gate excess against every fibre state, which is exactly capture by the final gate. Conversely, a capturing predictor gives zero excess. If no such predictor exists, the attained value cannot be zero; hence it is strictly positive. Finally, lower semicontinuity under compact convergence gives closedness of the superlevel sets $\{\mathfrak{G} \geq \delta\}$. $\square$

11 Compactness and Zero-Cost Branches

The conditional proof repeatedly passes to limits of bounded scale-normalized objects. The following propositions isolate what is used and what is not.

Proposition 11.1 (Suitable limit passage). *Let (u_n, p_n) be suitable weak solutions on Q_1 with uniform bounds*

$$\sup_n \left(\|u_n\|_{L_t^\infty L_x^2(Q_1)} + \|\nabla u_n\|_{L^2(Q_1)} + \|p_n\|_{L^{3/2}(Q_1)} \right) < \infty.$$

Then a subsequence converges to a suitable weak solution (u, p) on compact subcylinders of Q_1 , with strong convergence of u_n in L_{loc}^2 and lower semicontinuity of the local energy inequality.

Proof. The equation gives a uniform bound for $\partial_t u_n$ in a negative Sobolev space on compact subcylinders. Aubin-Lions compactness gives strong L_{loc}^2 convergence of u_n , while the pressure bound gives weak $L_{\text{loc}}^{3/2}$ convergence of p_n . The nonlinear term passes to the limit because $u_n \rightarrow u$ strongly in L_{loc}^2 and $\nabla u_n \rightharpoonup \nabla u$ weakly in L_{loc}^2 . The local energy inequality passes by weak lower semicontinuity of the dissipation and strong convergence of the lower-order terms. $\square$

Proposition 11.2 (Normal-defect limit passage). *Suppose a sequence of normalized presentations has uniformly bounded predictors, uniformly bounded suitable fibres, convergent pressure gauges, and normal defect functionals $\mathcal{N}_n$ converging weakly to $\mathcal{N}$ in a common dual space. If the corresponding witness spaces converge by compact embeddings and the witness lower bounds are uniform, then every nonzero limiting normal defect is detected by a limiting witness. In particular, a vanishing witness pairing for all limiting witnesses forces $\mathcal{N} = 0$.*

Proof. Let e be a unit witness for the limit space. By the compact convergence of witness spaces, choose witnesses e_n in the n -th space with $e_n \rightarrow e$ strongly. Weak convergence of $\mathcal{N}_n$ and strong convergence of witnesses give

$$\langle \mathcal{N}_n, e_n \rangle \rightarrow \langle \mathcal{N}, e \rangle.$$

If every limiting pairing is zero, then $\mathcal{N}$ annihilates the closure of the witness span. The uniform witness lower bound says this span is norming for normal defects, hence $\mathcal{N} = 0$. $\square$

Proposition 11.3 (Euler-coherent pressure branch). *Let U be a smooth divergence-free field on a cylinder and assume its quadratic nonlinearity is a pure pressure gradient:*

$$\mathbb{P}((U \cdot \nabla)U) = 0.$$

Then the field U is invisible to the Leray-projected quadratic normal defect. Consequently, a zero-cost branch in the predictive presentation must include not only Beltrami and two-and-a-half-dimensional examples, but every coherent branch for which the projected quadratic term vanishes after the selected pressure gauge.

Proof. The Leray projection annihilates gradients. If $(U \cdot \nabla)U = \nabla P_U$ in the chosen local gauge, then the projected quadratic contribution is zero. Therefore no normal witness built only from the projected quadratic leakage can distinguish such a branch. This proves the claim and explains why the branch alternative in the conditional theorem is formulated as a general coherent zero-cost branch rather than only the two explicit model branches proved above. $\square$

12 Branch Gates and Nonlinear Witnesses

The branch alternative becomes useful only when it feeds the final gate rather than merely explaining a loss of quadratic leakage. The next statement records the exact implication needed from a branch construction.

Proposition 12.1 (Branch gate criterion). *Let a coherent branch predictor G be produced on a cylinder Q_R , with branch pressure P_G , and suppose:*

- (i) *the predictor has a Serrin bound $\|G\|_{L_t^q L_x^p(Q_R)} \leq L$, $2/q + 3/p = 1$, $p > 3$;*
- (ii) *the relative perturbation obeys*

$$\|u - G\|_{L^{10/3}(Q_R)} + \|p - P_G\|_{L^{3/2}(Q_R)} + \text{PressTail}(Q_R, Q_{\theta R}) \leq \delta;$$

- (iii) *all pressure gauges used to define P_G and p are synchronized on the smaller cylinder.*

Then, for δ small depending only on L, p, θ , the final relative Caffarelli-Kohn-Nirenberg gate opens on $Q_{\theta R}$.

Proof. After parabolic normalization, this is exactly the relative Caffarelli-Kohn-Nirenberg criterion. The pressure synchronization hypothesis is needed only to ensure that the pressure difference entering the relative quantity is the same pressure difference controlled in the branch construction. With the gauges aligned, the smallness assumption gives the perturbative quantity required by the criterion. $\square$

Proposition 12.2 (Nonlinear Caffarelli-Kohn-Nirenberg witnesses are not Hilbert innovations). *The raw Caffarelli-Kohn-Nirenberg quantity*

$$C_R(u, p; z_0) = R^{-2} \int_{Q_R} |u|^3 + R^{-2} \int_{Q_R} |p - (p)_{B_R}(t)|^{3/2}$$

cannot, by itself, be inserted as a linear Hilbert innovation in the Bessel packing argument. Any use of this quantity in the predictive proof must be routed through one of the following mathematical alternatives: a relative gate, a linearized source functional with a norming witness, a compactness-rigidity branch, or a separate monotone/Carleson estimate.

Proof. The Hilbert packing lemma applies to linear functionals paired with orthonormal witnesses. The map $u \mapsto \int |u|^3$ is nonlinear and has a large kernel under linearization at $u = 0$; the pressure part is also nonlinear through the quadratic pressure relation. Therefore raw Caffarelli-Kohn-Nirenberg mass is not a Hilbert coefficient. To use it in a square-summability argument one must first produce a genuine linear defect functional, or else use the quantity as a gate whose success exits the iteration. The listed alternatives are exactly those two possibilities and their compactness variants. $\square$

Proposition 12.3 (Gated witnesses exit the source count). *If a Caffarelli-Kohn-Nirenberg witness opens the final relative gate at a scale, then that scale is terminal for the bad-chain argument and its witness mass is not counted in the Hilbert source budget. If the gate does not open, the witness may be used only after it has been converted into a source functional, a rigidity alternative, or an escape alternative.*

Proof. Opening the final gate gives regularity in a smaller cylinder by the relative criterion and the classical epsilon theorem. A bad chain cannot continue through that cylinder, so the scale is no longer part of the infinite source count. If the gate remains closed, no regularity conclusion has been obtained; counting the raw nonlinear witness as a Hilbert coefficient would be exactly the invalid step excluded by the preceding proposition. Thus the two cases are disjoint. $\square$

13 Source Theorem Interface

We now state the analytic inputs consumed by the reduction. Each source theorem has a specific output: an object, a summability statement, a branch, a gate, or a registered defect. The modular proof never uses an unnamed analytic input.

Proposition 13.1 (Finite presentation normal form). *Assume that along a tail of a predictive chain the presentation alphabet is fixed, the admissible transition rule between consecutive scale states is total, and every accepted positive component is assigned to exactly one source ledger. Then every transition on that tail has one of the following statuses:*

- (i) *accepted source term;*
- (ii) *finite update;*
- (iii) *relative gate;*
- (iv) *registered defect;*
- (v) *coherent branch;*
- (vi) *carrier recentering;*
- (vii) *impossible branch.*

Consequently, on a no-gate, no-defect tail with only finitely many updates, the modular proof may consume only the named source ledgers and the branch or carrier alternatives stated below.

Proof. This is a finite normal-form argument. The fixed alphabet prevents new source names or new gauges from appearing after the tail has been selected. Totality of the transition rule classifies each adjacent scale transition. The owner map assigns every accepted positive component to one source ledger, so no accepted cost can be used twice or disappear. After finitely many finite updates are discarded, the listed statuses exhaust all transitions on the tail. $\square$

Assumption 13.2 (Predictive epsilon regularity source theorem). *For every L and every admissible pressure-cleaning scheme, there exist $\varepsilon(L) > 0$ and $0 < \theta < 1$ such that the following holds. If (u, p) is suitable in Q_1 , and if there is a strong predictor g with*

$$\|g\|_{L_t^q L_x^p(Q_1)} \leq L, \quad \frac{2}{q} + \frac{3}{p} = 1, \quad p > 3,$$

such that

$$\|u - g\|_{X(Q_\theta)} + \text{PressTail}(Q_1, Q_\theta) \leq \varepsilon(L),$$

then u is regular in a smaller cylinder $Q_{\theta/2}$. Here X is the pressure-cleaned local energy-critical control norm used in the predictive distance, and the pressure-tail term is measured in the same gauge used for the relative Caffarelli-Kohn-Nirenberg gate.

Assumption 13.3 (Predictor production and differentiability source theorem). *For every bounded incoming presentation there is a nonempty class of admissible strong predictors. The causal strong solution map Φ , from incoming data to core predictors, is twice differentiable in the selected critical topology on each bounded predictor class. Quantitatively, for incoming data a and perturbations h ,*

$$\|\Phi(a + h) - \Phi(a) - D\Phi(a)h\|_X \leq C_L \|h\|_{\mathcal{A}}^2,$$

with constants stable under parabolic rescaling and compatible with the selected pressure-cleaning procedure. The theorem also supplies realizable predictors within the stated realization error of the relaxed minimax centre.

Assumption 13.4 (Global residual representation source theorem). *Along every bad predictive chain, the registered local defects determine vectors d_k in a fixed Hilbert defect space H_{def} . There exist a Hilbert space $\mathcal{K}$, a residual vector $\mathcal{R} \in \mathcal{K}$, bounded maps $A_k : \mathcal{K} \rightarrow H_{\text{def}}$, and square-summable errors ρ_k such that*

$$d_k = A_k \mathcal{R} + \rho_k.$$

Moreover, the pulled-back innovation tests satisfy the Bessel bound required in the moving-chart packing proposition. The errors include, with no double counting, pressure tails, cutoff losses, gauge changes, finite-window transport errors, carrier recentering errors, and predictor-realization errors.

Assumption 13.5 (Packet rigidity source theorem). *If a localized parabolic wave-packet family has quadratically small normal leakage on a quantitatively connected interaction graph, then, after removing pressure and gauge directions, it is close in the critical norm to a coherent zero-cost branch. The branch class contains the explicit Beltrami and two-and-a-half-dimensional branches and every smooth Euler-coherent branch whose projected quadratic nonlinearity is pure pressure in the selected gauge. If the packet family is not close to such a branch, the theorem produces a registered normal defect detected by the Hilbert witness system.*

Assumption 13.6 (Carrier localization source theorem). *High-frequency carrier escape has one of three outcomes. Either the carrier can be recentered so that the dominant frequency returns to order one in a new admissible presentation; or the escaped component is locally subcritical in every carrier cell and therefore cannot maintain the bad-chain defect; or the escape produces a registered normal, pressure, gauge, localization, or transport defect which enters the global residual representation with square-summable error.*

Assumption 13.7 (Pressure and gauge synchronization source theorem). *For every scale block used in the iteration, pressure representatives, harmonic tails, gauge normalizations, and mean-zero conventions can be chosen before the block is evaluated. Under those fixed choices, all pressure-force pairs used by the Oseen estimate, the branch gate, and the relative Caffarelli-Kohn-Nirenberg functional refer to the same pressure difference on the core cylinder. Any failure of synchronization is recorded as a pressure or gauge defect in the global residual representation.*

Assumption 13.8 (Localization and critical collar source theorem). *For every controlled Oseen block, the fixed-profile derivative collars are handled by the dyadic Hardy estimate proved above. All remaining localization terms—moving cutoffs, carrier collars, pressure tails created by localization, and the critical collar sums*

$$\sum_k r_k^{-2} \int_{A_k} |w_k|^3, \quad \sum_k r_k^{-2} \int_{A_k} |\pi_k^0|^{3/2}$$

—either have a scale-invariant source bound, open a relative gate, produce a registered pressure/localization/carrier defect, or enter a coherent branch covered by the branch source theorem.

Assumption 13.9 (Branch production and size source theorem). *Whenever the packet rigidity theorem selects a coherent zero-cost branch, the presentation produces a strong branch predictor on a smaller cylinder together with its synchronized pressure. The predictor has a Serrin bound, its pressure tail is controlled in the selected gauge, and its size data imply the relative Caffarelli-Kohn-Nirenberg smallness needed by the branch gate, unless a registered branch defect, pressure defect, force defect, or carrier escape is produced.*

Assumption 13.10 (Finite presentation and update source theorem). *The iteration uses a finite presentation alphabet at each bounded complexity level. Predictor charts, pressure gauges, finite observational Gramians, branch states, and carrier charts either stabilize along a tail of the bad chain, are updated only finitely many times inside a fixed complexity level, or produce a registered update/escape defect. Thus an infinite bad chain cannot hide an unpaid cost in infinitely many silent changes of presentation.*

Assumption 13.11 (Quantitative closure alternative source theorem). *For every L and $\delta > 0$, there are $m = m(L, \delta)$ and $c = c(L, \delta) > 0$ such that the following holds along any m -step controlled block of a predictive chain. If $\mathfrak{D}_k \geq \delta$ throughout the block and all registered normal, holonomy, pressure, carrier, update, and realization defects in the block are at most c , then either the defect decreases below $\delta/2$ at the next controlled scale, or the chain enters a coherent branch or a carrier branch covered by the previous source theorems.*

13.1 Analytic work packages

The source theorems above are the complete conditional interface used by the proof. Expanded into concrete analytic obligations, they require the following statements.

- (i) *Critical predictor construction.* The incoming and lateral observations must generate a nonempty family of strong predictors with Serrin control, stable local pressure gauges, and quantitative dependence on the incoming data. The differentiability estimate must be proved in the same critical topology used by the predictive defect, not in a stronger topology that loses scale invariance.
- (ii) *Realization of minimax centres.* The relaxed minimax centre exists by Hilbert compactness, but the PDE iteration needs a realizable strong predictor near that relaxed centre. The realization theorem must either produce such a predictor with the stated error or register the error as a summable realization defect.
- (iii) *Global residual representation.* Normal defects produced at different scales must be represented as vectors d_k in a fixed defect Hilbert space and must satisfy $d_k = A_k \mathcal{R} + \rho_k$, with a Bessel bound for the pulled-back innovation tests $A_k^* e_k$. The error vectors ρ_k must control connection changes, cutoff changes, pressure gauges, carrier recentering, finite-window truncation, and chart changes. This is the statement that makes the moving-chart packing proposition applicable to Navier-Stokes.
- (iv) *Packet pairing and row summability.* The localized wave-packet decomposition must supply actual pairings between physical residuals and adjoint witnesses. Repeated packet rows must obey either a Bessel/Carleson bound or a direct operator bound in the critical topology. A row that cannot be made summable must exit by a gate, a coherent branch, carrier escape, or a registered defect.
- (v) *Packet rigidity.* Small projected quadratic leakage must force the localized packet family toward a coherent zero-cost branch after pressure and gauge directions are removed. If this rigidity conclusion fails, the failure must produce a normal defect detected by the witness system. The theorem must cover localized packets, not only exact Fourier modes.
- (vi) *Pressure and gauge source closure.* Pressure representatives, harmonic tails, mean-zero conventions, gradient gauges, and pressure-force pairs must be chosen in a block-independent way. Variations of representatives, pressure tails, gauge shrink losses, cutoff pressure losses, and pressure-force couplings must be assigned once and only once to source terms, gates, or escape alternatives.

- (vii) *Native force and old-data closure.* Forcing terms created by predictor updates, old visible data, finite-memory truncations, and transport of past information must either be absorbed by the predictor construction or registered as force defects with a summable source bound. The proof cannot use the same force contribution simultaneously as a predictor input and as a separate residual source.
- (viii) *Cutoff and localization closure.* The fixed-profile cutoff and fixed-drift subchannels are proved above. The remaining localization theorem must handle moving cutoffs, annular critical mass, carrier collars, pressure tails created by localization, and compatibility between localization cells and the Oseen transfer. Any failure must be routed to a localization defect, carrier escape, pressure defect, or gate.
- (ix) *Oseen source closure beyond fixed subchannels.* The local Oseen energy transfer is proved, and fixed cutoff/drift costs are controlled. What remains is the source theorem for moving drifts, pressure-force pairs, native forcing, localization-carrier-cutoff compatibility, and explicit Oseen error terms. These contributions must be separated from the packet, pressure, force, and carrier source counts so that no cost is counted twice.
- (x) *Carrier source closure.* A high-frequency carrier must either be recentred into a new admissible chart, become locally subcritical in each carrier cell, or produce a summable registered defect. The theorem must also control carrier tails, containment of carrier cells in the physical cylinder, and compatibility with the pressure gauge used by the final gate.
- (xi) *Coherent branch production.* Once packet rigidity selects a zero-cost branch, the proof needs an actual branch predictor with synchronized pressure. The branch theorem must produce outer branch data, interior size data, derivative data when needed, pressure synchronization, and a route from branch size to relative Caffarelli-Kohn-Nirenberg smallness. If any of these data fail, the failure must become a registered branch, pressure, force, or carrier defect.
- (xii) *Branch radius and margin compatibility.* The final gate is opened on a smaller cylinder, so the branch radius, analytic radius, cutoff radius, pressure tail allowance, and gate margin must be synchronized. The theorem must show that downstream consumption of these margins leaves a positive final gate allowance, or else identifies a definite registered defect.
- (xiii) *Nonlinear Caffarelli-Kohn-Nirenberg witnesses.* Raw Caffarelli-Kohn-Nirenberg mass is not a Hilbert coefficient. Every use of such mass must therefore be routed through the relative gate, converted to a genuine linear source functional, controlled by an independent monotone/Carleson estimate, or absorbed by compactness-rigidity. Gated rows exit the iteration and are not counted as source terms.
- (xiv) *Finite presentation and update closure.* The finite observational atlas, predictor charts, pressure gauges, memory windows, coefficient ranks, branch states, and carrier states must stabilize on controlled tails or pay a finite update/escape cost. This prevents an infinite bad chain from changing presentations indefinitely without producing a summable source term or an exit.
- (xv) *Final floor and compactness closure.* The no-gate condition, the positive gate-excess defect, the compactness of bounded charts, and the lower-semicontinuity of the final floor must be compatible with scale iteration. This is the compactness statement that lets a persistent bad block survive in a limit only if it retains a positive defect, and it is the point where the quantitative closure alternative is consumed.

14 Conditional Newton Transfer

Under the predictor production and differentiability source theorem, the Newton bridge becomes a precise conditional proposition.

Proposition 14.1 (Conditional saddle Newton bridge). *Assume critical predictor-map differentiability and the local Oseen transfer estimate. Let g_k be a relaxed minimax predictor at scale k , and write $u = g_k + w_k$. Then, after passing to the next core scale,*

$$\mathfrak{D}_{k+1} \leq C_L \mathfrak{D}_k^2 + C_L \left(\mathfrak{R}_k^\perp + \mathfrak{H}_k^\perp + \mathfrak{F}_k^\perp + \mathfrak{V}_k \right) + \eta_k.$$

Here $\mathfrak{R}_k^\perp$ is the non-absorbed normal quadratic defect, $\mathfrak{H}_k^\perp$ the holonomy or gauge defect, $\mathfrak{F}_k^\perp$ the pressure/localization defect, $\mathfrak{V}_k$ the realization defect of the relaxed centre, and η_k the chosen minimization error.

Proof. The predictable linear part of w_k is absorbed by the first variation $D\Phi(a)h$. The tangential quadratic part is absorbed by the second-order correction supplied by the differentiability source theorem. The non-absorbed quadratic part is exactly $\mathfrak{R}_k^\perp$; pressure, cutoff, and connection errors give $\mathfrak{F}_k^\perp$ and $\mathfrak{H}_k^\perp$; and non-realisability of the relaxed centre gives $\mathfrak{V}_k$. Comparing the suitable weak solution with the updated strong predictor, the local Oseen estimate bounds the remainder on the smaller core. Since the updated predictor is an admissible competitor for the minimax problem at the next scale, the displayed inequality follows. $\square$

15 Modular Regularity Theorem

Theorem 15.1 (Conditional predictive regularity). *Let (u, p) be a suitable weak solution near z_0 . Suppose that the predictive presentation system at z_0 satisfies the analytic source theorems stated above, and that the explicit scale-selection errors η_k are summable. Then there is no bad predictive chain at z_0 . Consequently, if the predictive defect is the only obstruction to the predictive epsilon-regularity hypothesis, then u is regular at z_0 .*

Proof. Assume that a bad predictive chain exists. Then for some $\delta_0 > 0$,

$$\mathfrak{D}_k \geq \delta_0$$

on all sufficiently small scales.

By the global residual representation source theorem and the Hilbertian packing lemma, the innovative normal defects are square-summable. Hence, after discarding finitely many initial scales, one can find arbitrarily long controlled blocks on which the average registered normal defect is as small as required. The pressure and gauge synchronization theorem ensures that the pressure quantities appearing in those blocks are the same quantities used by the Oseen estimate and by the relative Caffarelli-Kohn-Nirenberg gate.

Apply the quantitative closure alternative on such a block. Since $\mathfrak{D}_k \geq \delta_0$ throughout the bad chain, the first outcome-defect decay below $\delta_0/2$ —is impossible on sufficiently late blocks. Therefore the chain must enter either a coherent branch or a carrier branch, unless a registered defect has been produced. A registered defect is also impossible infinitely often with fixed positive size, because the global residual representation and the packing lemma make the registered innovative cost square-summable.

If the chain enters a Beltrami branch, the quadratic term is pressure-cleaned and the Newton transfer gives decay of the predictive defect. If it enters a two-and-a-half-dimensional branch, the reduced two-dimensional system and passive scalar equation provide the strong predictor and again force decay, up to a registered anisotropic defect. If it enters a more general Euler-coherent zero-cost branch, the branch production and size source theorem supplies the

branch predictor and pressure data needed for the branch gate, unless a registered branch, force, pressure, or carrier defect is produced. In the gate case, the relative Caffarelli-Kohn-Nirenberg criterion opens the final gate and the bad chain terminates. In the defect case, the defect has already been shown to be square-summable and therefore cannot maintain the fixed lower bound δ_0 indefinitely.

If the chain enters carrier escape, the carrier localization source theorem either recenters the scale, makes the escaped contribution locally subcritical, or produces one of the registered defects already controlled by packing. Recentring gives a new admissible presentation and returns to the same alternative; subcriticality contradicts persistence of the defect; and a registered defect is square-summable. Every possible outcome contradicts the existence of a bad chain with $\mathfrak{D}_k \geq \delta_0$ on all sufficiently small scales.

Therefore no bad predictive chain exists. The final implication is exactly the predictive epsilon regularity source theorem. $\square$

16 Verified Components and Source Inputs

The verified pieces are:

- (i) the Serrin-critical Oseen transfer estimate, with collar and pressure-tail terms explicitly paid for;
- (ii) fixed-profile dyadic cutoff-collar control, fixed-drift collar stability, and the pressure-pair source interface;
- (iii) the Hilbert-space existence of relaxed minimax centres;
- (iv) the abstract Bessel packing mechanism for innovative normal witnesses;
- (v) the exact two-mode null classification of the projected bilinear symbol;
- (vi) exact absorption of Beltrami fields and exact reduction of two-and-a-half-dimensional fields;
- (vii) the relative Caffarelli-Kohn-Nirenberg gate, conditional only on the stated predictor and pressure smallness;
- (viii) compactness of the gate-excess defect under the selected chart topology;
- (ix) suitable limit passage and normal-defect limit passage under the stated compactness hypotheses;
- (x) the separation between nonlinear Caffarelli-Kohn-Nirenberg witnesses and Hilbert innovations.

The conditional theorem depends on the following scale-invariant PDE source theorems:

- (i) production, differentiability, and realization of causal strong predictors;
- (ii) construction of a global residual representation across a bad chain, with a Bessel bound and square-summable pressure, cutoff, gauge, carrier, and transport errors;
- (iii) packet-level rigidity for localized nearly invisible quadratic interactions;
- (iv) carrier localization and recentring with pressure-tail control;
- (v) localization and critical velocity/pressure collar source bounds beyond the fixed-profile Hardy subchannel;

- (vi) pressure and gauge synchronization across Oseen, branch, and final-gate quantities;
- (vii) production of branch predictors, branch pressure data, and branch size smallness;
- (viii) finite presentation/update stabilization;
- (ix) predictive epsilon regularity and the quantitative closure alternative.

The useful conclusion of the paper is therefore a dependency graph with a proved core. Presentation theory supplies the language of predictors, observables, fibres, costs, and legal transfer; Oseen estimates supply the perturbative transport; Hilbert packing supplies a square-summability mechanism for genuinely new normal defects; relative Caffarelli-Kohn-Nirenberg gates convert predictor smallness into regularity; and null geometry isolates the branches where the quadratic term can be invisible. The remaining analytic task is to prove the listed source theorems in the concrete Navier-Stokes topology.

References

- [ABP] D. Albritton, T. Barker, and C. Prange, *Epsilon regularity for the Navier-Stokes equations via weak-strong uniqueness*, arXiv:2211.16188.
- [CKN] L. Caffarelli, R. Kohn, and L. Nirenberg, *Partial regularity of suitable weak solutions of the Navier-Stokes equations*, Communications on Pure and Applied Mathematics 35 (1982), 771-831.
- [ESS] L. Escauriaza, G. Seregin, and V. Sverak, *$L_{3,\infty}$ -solutions of Navier-Stokes equations and backward uniqueness*, Russian Mathematical Surveys 58 (2003), 211-250.
- [Leray] J. Leray, *Sur le mouvement d'un liquide visqueux emplissant l'espace*, Acta Mathematica 63 (1934), 193-248.
- [Serrin] J. Serrin, *On the interior regularity of weak solutions of the Navier-Stokes equations*, Archive for Rational Mechanics and Analysis 9 (1962), 187-195.
- [Tao] T. Tao, *Quantitative bounds for critically bounded solutions to the Navier-Stokes equations*, preprint, 2019.
- [Wolf] J. Wolf, *On the local pressure of the Navier-Stokes equations and related systems*, Advances in Differential Equations 17 (2012), 305-338.