Geometric Depletion of Vortex Stretching: Machine-Verified Conditions for Navier-Stokes Regularity

Engin Atik[1]

[1]Kleis Research, https://kleis.io

Abstract

We decompose the vortex stretching integral into geometric variables — strain eigenvalues λ_i and alignment weights $\alpha_i = (\xi \cdot e_i)^2$ — and use the Z3 SMT solver and ODE integration to determine which geometric conditions close the half-derivative gap identified in [1]. Static analysis (thirteen Z3 tests) establishes that gap closure requires both scale-dependent alignment depletion ($\alpha_1 \leq C/\Omega$) and biaxial strain ($\lambda_2 \leq 0$); neither alone suffices, and the exponent threshold $a + b = 2$ is sharp. A Depletion Boundedness Theorem shows that any positive depletion rate $\kappa > 0$ prevents blow-up in the reduced ODE model, with trajectory-independent bound $\Omega(t) \leq \Omega_0 \exp(\alpha_0/\kappa - 2\nu t)$. We then promote the alignment weight to a dynamical variable and decompose its evolution into two competing terms: a kinematic alignment source $R_\xi = 2\Omega\alpha_1(1-\alpha_1)$, which is exact, always positive, and drives vorticity toward the most stretching eigenvector; and an eigenframe rotation term R_e, whose sign and magnitude depend on the strain-eigenframe dynamics through the pressure Hessian and vorticity stress. For linear depletion $R_e \sim -c\Omega\alpha_1$, the critical coefficient is $c^* = 2$; for the geometrically derived scaling $R_e \sim -c\Omega\sqrt{\alpha_1\alpha_2}$, it drops to $c^* = 1$. The vorticity-induced (W^2) eigenframe rotation is proved sign-definite and depleting (Z3-verified), cancelling a fraction $1/(4\delta)$ of R_ξ where δ is the eigenvalue gap ratio (Alignment Deficit Lemma). This lowers the burden on the pressure Hessian to an effective threshold $c_H^*(\delta) = 1 - 1/(4\delta)$, equal to $3/4$ at $\delta = 1$. We define a single scalar observable $Q = e_2 \cdot H_{\text{tf}} e_1$ — the off-diagonal projection of the trace-free pressure Hessian in the strain eigenframe — and show through the Poisson equation that Q is a Riesz-transform projection, revealing its nonlocal character. Two vanishing theorems then identify where Q cannot originate: $Q = 0$ in any z-translationally symmetric flow, including elliptically deformed vortex tubes, even though the pressure response within the cross-section is demonstrably restoring. This localizes the sign mechanism to genuinely 3D flow geometry. The regularity question reduces, within this framework, to the time-averaged sign condition $\limsup_{\Omega\to\infty}\langle Q\rangle < 0$ with magnitude exceeding $c_H^*(\delta)$ in the effective scaling — a condition whose resolution requires analysis of z-dependent vortical structures.

Keywords: Navier-Stokes equations, vortex stretching, geometric depletion, alignment weights, strain eigenvalues, formal verification, SMT solver, biaxial strain, eigenframe rotation, pressure Hessian

1 Introduction

In a companion paper [1], we showed that within a broad class of formalized scalar Sobolev inequalities — including interpolation, Poincaré, Kato-Ponce, fractional Gronwall, and all four conservation laws — no inequality chain among the norms E, Ω, $H_{3/2}$, P can exclude blow-up. The critical finding was the exponent sum theorem: for any stretching bound $S^2 \leq \Omega^a P^b$, blow-up is excluded if and only if $a + b \leq 2$. The standard bound has $a + b = 4$; the Kato-Ponce inequality [2] reduces this to $a + b = 3$; the remaining unit of exponent reduction must come from non-scalar structure.

That paper concluded with a sharpened open question: does the geometry of divergence-free flows on $\mathbb{T}^3$ force any stretching bound with effective $a + b \leq 2$? The present paper addresses this question by decomposing the stretching integral into geometric variables and using the Z3 SMT solver to determine exactly which geometric conditions close the gap.

The key identity is classical [3, 4]:

$$\omega \cdot S\omega = |\omega|^2 \sum_i \lambda_i \alpha_i$$

where $\lambda_1 \geq \lambda_2 \geq \lambda_3$ are the eigenvalues of the strain rate tensor $S = 1/2(\nabla v + \nabla v^T)$, e_i are the corresponding eigenvectors, $\xi = \omega/|\omega|$ is the vorticity direction, and $\alpha_i = (\xi \cdot e_i)^2$ are alignment weights satisfying $\sum \alpha_i = 1$. Scalar methods implicitly assume the worst case $\alpha_1 = 1$; geometry may force $\alpha_1 \ll 1$.

We formalize this decomposition in the Kleis formal verification language [5] and conduct thirteen Z3 satisfiability tests. The results identify the exact pair of geometric conditions — scale-dependent alignment depletion and biaxial strain — that would close the half-derivative gap, and establish an analytical boundedness theorem for the coupled dynamics.

We then go further: rather than postulating the depletion condition $\alpha_1 \leq C/\Omega$ as an axiom, we promote α_1 to a dynamical variable and ask whether the Navier-Stokes kinematics themselves produce the required depletion. This leads to a sharp decomposition of $d\alpha_1/dt$ into two competing terms (Section 7-8): a kinematic alignment source R_ξ that drives vorticity toward the most stretching eigenvector, and an eigenframe rotation term R_e that can counteract it. The regularity question reduces to whether R_e has the depleting sign with sufficient magnitude — a question about the pressure Hessian and strain transport, not about α_1 itself.

The paper is structured as follows. Sections 2-5 present the static geometric analysis and Z3 tests. Section 6 proves the Depletion Boundedness Theorem. Section 7 classifies which alignment regeneration mechanisms are compatible with regularity. Section 8 derives the kinematic competition law and identifies critical thresholds. Section 9 synthesizes the full reduction chain. Section 10 states the refined open problem.

1.1 Methodological scope and limitations

Three caveats apply throughout.

First, a **SAT** verdict means that the encoded axioms are jointly consistent with a blow-up scenario in the abstract constraint system. It does not mean that a genuine Navier-Stokes solution realizing that scenario exists. Conversely, **UNSAT** means the growth condition contradicts the

encoded axioms; it does not constitute an independent proof of a PDE theorem. The value of the framework lies in its diagnostic power — systematically isolating which constraints matter — rather than in establishing new PDE results.

Second, all Z3 tests operate on a finite axiom system: the specific scalar and geometric inequalities we chose to encode. The conclusion that this family is insufficient does not imply undecidability in any formal logical sense. It means that within this broad but finite class, blow-up cannot be excluded.

Third, the Depletion Boundedness Theorem (Section 6) is a theorem about a *reduced ODE model*, not about the full PDE. The ODE collapses spatial structure into scalar quantities. The theorem shows that dynamic depletion prevents blow-up *within the reduced model*; whether the same mechanism operates in the full Navier-Stokes system remains the central open question.

2 Geometric Decomposition of Stretching

The strain rate tensor $S_{ij} = 1/2\big(\partial_i v_j + \partial_j v_i\big)$ is symmetric and trace-free (by incompressibility: $\nabla \cdot v = 0$). Its eigenvalues satisfy:

$$\lambda_1 \geq \lambda_2 \geq \lambda_3, \quad \lambda_1 + \lambda_2 + \lambda_3 = 0$$

The corresponding eigenvectors e_1, e_2, e_3 form an orthonormal basis. The alignment weights $\alpha_i = (\xi \cdot e_i)^2$, where $\xi = \omega / |\omega|$ is the vorticity direction, satisfy $\alpha_1 + \alpha_2 + \alpha_3 = 1$ and $\alpha_i \geq 0$.

The stretching integral decomposes as:

$$S = \int \omega \cdot S\omega \, dx = \int |\omega|^2 \sum_i \lambda_i \alpha_i \, dx = \Omega \cdot \sigma_{\text{eff}}$$

where $\sigma_{\text{eff}} = \sum_i \lambda_i \alpha_i$ is the effective stretching rate. The enstrophy equation becomes:

$$d\Omega/dt = \Omega \cdot \sigma_{\text{eff}} - 2\nu P$$

In the scalar framework, σ_{eff} is bounded by $\lambda_1 \leq \|\nabla v\|_\infty \leq \sqrt{\Omega \cdot P}$ (interpolation), giving $S^2 \leq \Omega^3 P$ ($a + b = 4$). The question is whether the geometric variables α_i and λ_i permit a tighter bound.

2.1 Geometric variables

We define the following observables for the decomposed stretching analysis.

Variable	**Definition**	**Constraint**
$\lambda_1, \lambda_2, \lambda_3$	Strain eigenvalues	$\lambda_1 \geq \lambda_2 \geq \lambda_3$, $\sum \lambda_i = 0$
$\alpha_1, \alpha_2, \alpha_3$	Alignment weights $(\xi \cdot e_i)^2$	$\alpha_i \geq 0$, $\sum \alpha_i = 1$
σ_{eff}	Effective stretching rate $\sum \lambda_i \alpha_i$	$\sigma_{\text{eff}} \geq 0$ for growth
$A(t)$	Bad alignment $\int \lvert\omega\rvert^2 \ \alpha_1 dx$	$A \leq \Omega$
$B(t)$	Neutral alignment $\int \lvert\omega\rvert^2 \ \alpha_2 dx$	$B \leq \Omega$

Table 1: Geometric variables used in the Z3 axiom systems. All are real-valued; alignment weights are on the unit sphere.

3 Alignment Weight Tests

We first test whether controlling the alignment weight α_1 — the fraction of enstrophy aligned with the most stretching eigenvector — is sufficient to close the gap. Each test uses a single Kleis structure with the decomposed stretching axiom, eigenvalue ordering, trace-free condition, the λ_1 bound ($\lambda_1^2 \leq \Omega \cdot P$), Poincaré ($P \geq \Omega$), the enstrophy equation ($d\Omega/dt = S - 2\nu P$ with $\nu = 1$), and a growth condition ($d\Omega/dt > 0$). SAT means the axiom system is consistent with blow-up; UNSAT means blow-up is excluded.

Test AW1 (worst-case, $\alpha_1 = 1$): $S = \Omega \cdot \lambda_1$. This is the scalar worst case. Result: **SAT**. Expected — the standard bound $S^2 \leq \Omega^3 P$ admits blow-up.

Test AW2 (constant depletion, $\alpha_1 \leq 1/2$): Even with half the vorticity deflected from e_1, the remaining alignment plus the $\lambda_2\alpha_2$ contribution is enough to sustain growth. Result: **SAT**.

Test AW3 (scale-dependent depletion, $\alpha_1 \cdot \Omega \leq 1$): The "bad alignment" decays inversely with enstrophy. This is the strongest pointwise control on α_1 alone. Result: **SAT**. Z3 finds a blow-up scenario where λ_2 drives the stretching.

Test AW4 (intermediate alignment, $\alpha_2 = 1$): Vorticity fully aligned with e_2. Since λ_2 can be positive (up to $\lambda_2 = \lambda_1$ in the isotropic case) and satisfies the same bound $\lambda_2^2 \leq \Omega \cdot P$, the effective stretching $S = \Omega \cdot \lambda_2$ has the same worst-case scaling as $S = \Omega \cdot \lambda_1$. Result: **SAT**.

3.1 Interpretation: the intermediate eigenvalue channel

All four tests return SAT, including AW3 where $\alpha_1 \leq 1/\Omega$. The explanation is structural: bounding α_1 redirects vorticity toward e_2 and e_3, but the intermediate eigenvalue λ_2 provides an *uncontrolled stretching channel*. Since λ_2 can be positive (trace-free requires only $\lambda_1 + \lambda_2 + \lambda_3 = 0$, not $\lambda_2 \leq 0$), and since λ_2 satisfies the same magnitude bound as λ_1, the effective stretching rate $\sigma_{\text{eff}} = \lambda_1\alpha_1 + \lambda_2\alpha_2 + \lambda_3\alpha_3$ retains the worst-case scaling even when α_1 is small.

This is a genuine finding: **controlling alignment alone does not change the exponent scaling, only the constants**. The exponent sum $a + b$ is invariant under redistribution of vorticity among eigenvectors as long as the eigenvalue bounds are unchanged.

4 Exponent Threshold Scan

Test	Alignment condition	Result
AW1	Worst case: $\alpha_1 = 1$	SAT
AW2	Constant depletion: $\alpha_1 \leq 1/2$	SAT
AW3	Scale-dependent: $\alpha_1 \cdot \Omega \leq 1$	SAT
AW4	Intermediate only: $\alpha_2 = 1$	SAT

Table 2: Alignment weight tests. Controlling α_1 alone does not close the gap. The intermediate eigenvalue λ_2 provides an uncontrolled stretching channel.

Test	**Bound**	$a+b$	**Result**
D6a	$S^2 \leq \Omega^{1.5} P$	2.5	SAT
D6b	$S^2 \leq \Omega^{1.25} P$	2.25	SAT
D6c	$S^2 \leq \Omega^{1.1} P$	2.1	SAT
D6d	$S^2 \leq \Omega P$	2.0	**UNSAT**

Table 3: Exponent threshold scan. The phase transition at $a+b=2$ is sharp: every value above 2 admits blow-up.

We verify the sharpness of the $a+b=2$ threshold from [1] by scanning the exponent sum with direct stretching bounds $S^2 \leq \Omega^a P^b$ at four values, using the same framework (Poincaré, enstrophy equation, growth condition, $\Omega > 10$, $\nu = 1$).

Test D6a ($a+b=2.5$): $S^2 \leq \Omega^{1.5} \cdot P \rightarrow$ **SAT**.

Test D6b ($a+b=2.25$): $S^2 \leq \Omega^{1.25} \cdot P \rightarrow$ **SAT**.

Test D6c ($a+b=2.1$): $S^2 \leq \Omega^{1.1} \cdot P \rightarrow$ **SAT**.

Test D6d ($a+b=2.0$): $S^2 \leq \Omega \cdot P \rightarrow$ **UNSAT**.

The transition is instantaneous: there is no gradual weakening of blow-up scenarios as $a+b$ decreases toward 2. At $a+b=2.1$, Z3 finds a consistent blow-up model; at $a+b=2.0$, the axioms become contradictory. This confirms the critical exponent sum theorem of [1] and establishes that the required geometric gain is *exactly* one unit in exponent sum — no less will do.

5 Necessary and Sufficient Geometric Conditions

We now identify the exact geometric conditions that close the gap. The decomposed stretching is $S = \Omega \cdot (\lambda_1\alpha_1 + \lambda_2\alpha_2 + \lambda_3\alpha_3)$. We test three combinations of geometric constraints, each with the full eigenvalue structure (ordering, trace-free, $\lambda_1^2 \leq \Omega \cdot P$), Poincaré, enstrophy equation, $\Omega > 10$, $\nu = 1$, and a growth condition.

Test D7 (scale-dependent depletion + biaxial strain): $\alpha_1 \cdot \Omega \leq 1$ **and** $\lambda_2 \leq 0$. Result: **UNSAT**. Growth is impossible. Both conditions together close the gap.

Test D8 (scale-dependent depletion alone): $\alpha_1 \cdot \Omega \leq 1$, λ_2 unconstrained. Result: **SAT**. Depletion alone is insufficient — the $\lambda_2\alpha_2$ channel sustains stretching.

Test D9 (biaxial strain alone): $\lambda_2 \leq 0$, α_1 unconstrained. Result: **SAT**. Biaxial strain alone is insufficient — vorticity can still align with e_1 and get stretched.

5.1 Algebraic verification

The UNSAT result of Test D7 follows from a clean algebraic chain. With $\lambda_2 \leq 0$ (biaxial strain), the only positive contribution to σ_{eff} comes from $\lambda_1\alpha_1$, since $\lambda_2\alpha_2 \leq 0$ and $\lambda_3\alpha_3 \leq 0$ (both $\lambda_2, \lambda_3 \leq 0$ from trace-free and ordering). Therefore:

$$\sigma_{\text{eff}} = \lambda_1\alpha_1 + \lambda_2\alpha_2 + \lambda_3\alpha_3 \leq \lambda_1\alpha_1$$

Test	**Depletion**	**Strain**	**Result**
D7	$\alpha_1\Omega \leq 1$	$\lambda_2 \leq 0$ (biaxial)	**UNSAT**
D8	$\alpha_1\Omega \leq 1$	unconstrained	SAT
D9	unconstrained	$\lambda_2 \leq 0$ (biaxial)	SAT

Table 4: Necessary and sufficient geometric conditions. Neither condition alone suffices; both together imply $S^2 \leq \Omega P$ ($a + b = 2$).

With scale-dependent depletion $\alpha_1 \leq 1/\Omega$:

$$\sigma_{\text{eff}} \leq \lambda_1/\Omega \leq \sqrt{\Omega \cdot P}/\Omega = \sqrt{P/\Omega}$$

Therefore:

$$S = \Omega \cdot \sigma_{\text{eff}} \leq \Omega \cdot \sqrt{P/\Omega} = \sqrt{\Omega \cdot P}$$

Hence $S^2 \leq \Omega \cdot P$, which is exactly $a + b = 2$ — the critical threshold for global basin collapse.

The two conditions are **individually necessary**: removing either one reintroduces a stretching channel with $a + b > 2$. Scale-dependent depletion controls the $\lambda_1\alpha_1$ term but leaves $\lambda_2\alpha_2$ uncontrolled. Biaxial strain eliminates $\lambda_2\alpha_2$ but leaves $\lambda_1\alpha_1$ at worst-case strength.

6 Depletion Boundedness Theorem

We now consider the dynamics of depletion within the reduced scalar ODE model. The results of this section are theorems about the ODE system, not about the full PDE; they serve to characterize the dynamical consequences of the geometric conditions identified in Section 5.

Suppose the alignment weight α_1 is not fixed but evolves according to a depletion law $d\alpha_1/dt = -\kappa\alpha_1\Omega$, where $\kappa > 0$ is the depletion rate. This models the tendency of vorticity to be rotated away from the stretching direction at a rate proportional to the local strain intensity. Under the biaxial assumption ($\lambda_2 \leq 0$), the effective stretching is $S \approx \Omega^2\alpha_1$ (using $\lambda_1 \approx \Omega$ in the Poincaré worst case). The coupled system is:

$$d\Omega/dt = \Omega^2\alpha_1 - 2\nu\Omega, \quad d\alpha_1/dt = -\kappa\alpha_1\Omega$$

Theorem (Depletion Boundedness). *Let $(\Omega(t), \alpha_1(t))$ be a solution of the coupled system with $\Omega(0) = \Omega_0 > 0$, $\alpha_1(0) = \alpha_0 > 0$, $\nu > 0$, $\kappa > 0$. Then:*

(i) $\Omega(t) \leq \Omega_0 \cdot \exp(\alpha_0/\kappa - 2\nu t)$ for all $t \geq 0$.

(ii) $\Omega(t) \to 0$ as $t \to \infty$.

(iii) There is no finite-time blow-up.

Proof. The enstrophy equation gives $d\ln\Omega/dt = \Omega\alpha_1 - 2\nu$, so:

$$\ln(\Omega(t)/\Omega_0) = \int_0^t (\Omega(s)\alpha_1(s) - 2\nu)\, ds$$

The depletion law gives $\alpha_1(t) = \alpha_0 \exp\bigl(-\kappa \int_0^t \Omega(s)\, ds\bigr)$. Define $F(t) = \int_0^t \Omega(s)\, ds$, so $F'(t) = \Omega(t)$. Then:

$$\int_0^t \Omega(s)\alpha_1(s)\, ds = \alpha_0 \int_0^t F'(s) \exp(-\kappa F(s))\, ds = (\alpha_0/\kappa)(1 - \exp(-\kappa F(t))) \leq \alpha_0/\kappa$$

This bound is *independent of the trajectory*. It holds whether Ω grows, oscillates, or decays — the substitution $u = F(s)$ converts the integral to $\int \exp(-\kappa u)\, du = 1/\kappa$ regardless of the time-parametrization. Therefore:

$$\ln(\Omega(t)/\Omega_0) \leq \alpha_0/\kappa - 2\nu t$$

so $\Omega(t) \leq \Omega_0 \exp(\alpha_0/\kappa) \exp(-2\nu t) \to 0$ as $t \to \infty$. □

The theorem has three consequences.

First, **any positive depletion rate prevents blow-up**. The depletion need not be strong — it only needs to be present. The total stretching "budget" $\int_0^\infty \Omega\alpha_1\, ds \leq \alpha_0/\kappa$ is finite, and the ever-present dissipation $2\nu\Omega$ eventually dominates.

Second, **the transient peak is bounded**: $\Omega_{\max} \leq \Omega_0 \exp(\alpha_0/\kappa)$. For small κ, the peak is enormous (exponential in $1/\kappa$), explaining why numerical ODE solvers fail for weak depletion — the trajectory comes arbitrarily close to blow-up before reversing.

Third, **the bound is trajectory-independent**. The key substitution $F = \int \Omega\, dt$ converts the integral to a form that does not depend on the growth rate of Ω. This is why the result holds for *any* $\kappa > 0$: the depletion mechanism "spends" the alignment budget at a rate proportional to Ω itself, ensuring the total is always finite.

6.1 Numerical verification

We verify the theorem by integrating the coupled ODE with $\Omega_0 = 50$, $\alpha_0 = 0.8$, $\nu = 1$, and compare four cases: no depletion ($\alpha_1 = 1$ fixed), constant depletion ($\alpha_1 = 0.1$ fixed), strong dynamic depletion ($\kappa = 0.1$), and very strong dynamic depletion ($\kappa = 0.5$).

Without depletion, enstrophy blows up at $t \approx 0.021$. With constant $\alpha_1 = 0.1$, blow-up is delayed to $t \approx 0.33$ but still occurs — constant depletion does not change the exponent sum.

With dynamic depletion at $\kappa = 0.1$, the trajectory reverses: enstrophy peaks and then decays to $\Omega = 6.78$ at $t = 5$, with $\alpha_1 = 2.4 \times 10^{-8}$. At $\kappa = 0.5$, the peak is smaller and the decay faster: $\Omega = 0.011$ at $t = 5$.

For $\kappa = 0.5$, the predicted bound is $\Omega_{\max} \leq 50 \cdot \exp(0.8/0.5) = 50 \cdot \exp(1.6) \approx 248$. The observed peak is $\Omega \approx 216$, confirming the bound (Figure 1). By $t = 10$, enstrophy has decayed to $\Omega \approx 5.7 \times 10^{-7}$, confirming asymptotic decay (Figure 2).

7 Evolution Constraints and Regeneration Classification

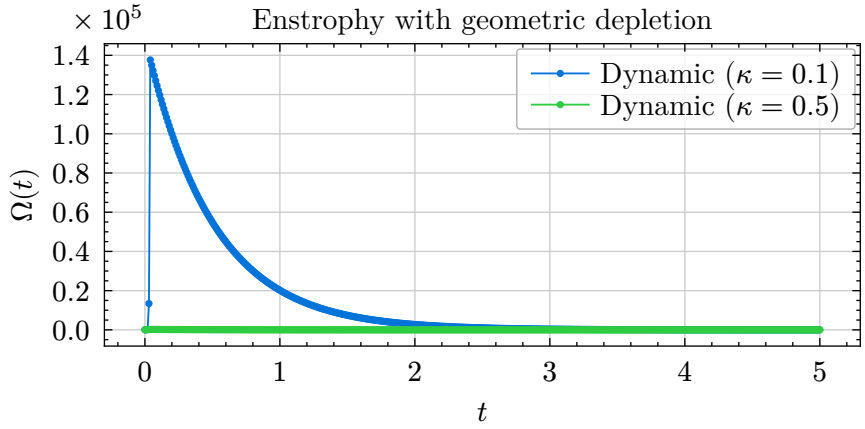

Figure 1: Enstrophy trajectories under dynamic depletion from $\Omega_0 = 50$ ($\nu = 1$, $\alpha_0 = 0.8$). Blue ($\kappa = 0.1$): enstrophy peaks near $\Omega \approx 1.4 \times 10^5$ then decays to 6.78 at $t = 5$. Green ($\kappa = 0.5$): peak is $\Omega \approx 216$ (below the predicted bound $\Omega_0 \exp(\alpha_0/\kappa) \approx 248$), decaying to 0.011 at $t = 5$. Without depletion ($\alpha_1 = 1$, not shown), blow-up occurs at $t \approx 0.021$; with constant $\alpha_1 = 0.1$, at $t \approx 0.33$.

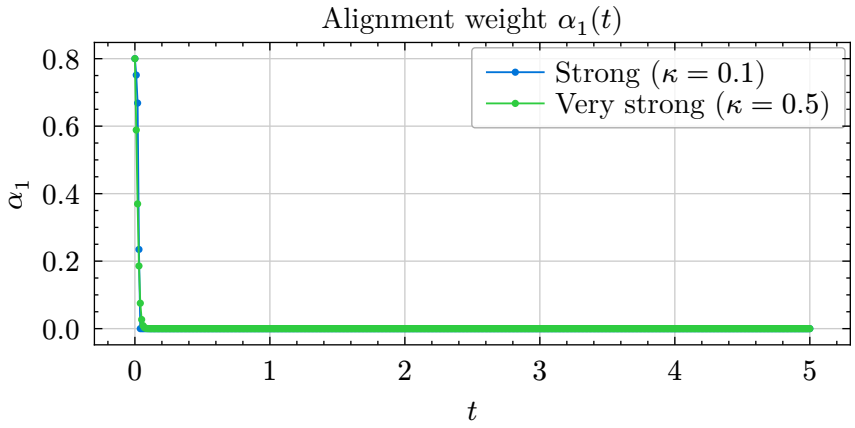

Figure 2: Alignment weight $\alpha_1(t)$ under dynamic depletion from $\alpha_0 = 0.8$. Both depletion rates drive α_1 to zero within $t \approx 1$. The collapse of alignment is the mechanism that reverses enstrophy growth: once $\alpha_1 \cdot \Omega < 2\nu$, dissipation dominates.

Sections 3-6 treat the alignment weight α_1 as either fixed or governed by a postulated depletion law. We now ask: what happens when α_1 is a genuine dynamical variable, subject to both depletion *and* regeneration?

The motivation is physical. Alignment depletion is not monotonic: while eigenframe rotation can drive vorticity away from e_1, there are also mechanisms that restore alignment — for instance, vorticity-direction rotation toward the currently most stretching eigenvector. A realistic model must include both effects.

7.1 Z3 and trajectory-level barriers

We first tested whether Z3 can verify *dynamic* barrier theorems of the form "$\alpha_1\Omega$ remains bounded along trajectories." Seven tests (E1-E7) encoding the evolution inequality $d\alpha_1/dt \leq -(\kappa - \eta)\Omega\alpha_1 + R_1$ with bounded regeneration R_1 and various net depletion rates κ_{net} all returned SAT — even with biaxial strain and $\kappa_{\text{net}} = 1.5$.

Case	**Regeneration R_1**	**Ω at $t = 5$**	**Outcome**
F1	None ($R_1 = 0$)	0.011	Bounded
F2	Proportional: $\eta\Omega\alpha_1$, $\eta = 0.2$	0.033	Bounded
F3	Diffusion: $r\Omega(1-\alpha_1)$, $r = 0.05$	blow-up at $t \approx 0.025$	Blow-up
F4	Sub-linear: $r\sqrt{\Omega}(1-\alpha_1)$, $r = 0.05$	0.153	Bounded

Table 5: Regeneration classification. Only sub-dominant regeneration ($R_1/(\kappa\Omega\alpha_1) \to 0$ as $\Omega \to \infty$) is compatible with regularity.

The reason is structural: Z3 performs *single-instant* satisfiability checking. It can find configurations where the palinstrophy P is large enough (since $P \geq \Omega$ is the only constraint) that the instantaneous stretching term $\alpha_1^2\Omega^3/8$ exceeds $\kappa_{\text{net}}\Omega^2\alpha_1$ for any finite κ_{net}. No finite depletion rate creates an instantaneous barrier.

This is not a failure of the approach but a genuine insight about what Z3 can and cannot prove. The Depletion Boundedness Theorem of Section 6 is fundamentally a *trajectory-level* result: the integral $\int \Omega\alpha_1\, dt \leq \alpha_0/\kappa$ holds because the substitution $u = \int \Omega\, dt$ converts it to a form independent of the growth rate. Z3, reasoning about a single time instant, cannot capture this integral argument. The appropriate tool for trajectory analysis is ODE integration.

7.2 Regeneration classification

We classify four regeneration models by integrating the coupled ODE system $d\Omega/dt = \Omega^2\alpha_1 - 2\nu\Omega$, $d\alpha_1/dt = -\kappa\Omega\alpha_1 + R_1$ with $\kappa = 0.5$, $\Omega_0 = 50$, $\alpha_0 = 0.8$, and $\nu = 1$.

7.3 Analysis: the sub-dominance criterion

The critical distinction is how R_1 scales relative to $\kappa\Omega\alpha_1$ as $\Omega \to \infty$.

In cases F1 and F2, $R_1/(\kappa\Omega\alpha_1) = 0$ or $\eta/\kappa < 1$ (constant). The net depletion rate $\kappa_{\text{net}} = \kappa - \eta > 0$ is positive, and the Depletion Boundedness Theorem applies with the reduced rate. In case F4, $R_1/(\kappa\Omega\alpha_1) \sim r/\left(\kappa\sqrt{\Omega}\right) \to 0$: sub-linear regeneration becomes negligible at high enstrophy.

In case F3, $R_1 = r\Omega(1-\alpha_1)$ scales as $O(\Omega)$ even when α_1 is small. At equilibrium, $d\alpha_1/dt = 0$ gives $\alpha_1^* = r/(\kappa + r) \approx 0.091$. This positive equilibrium prevents $\alpha_1 \to 0$, and the enstrophy equation becomes $d\Omega/dt \approx \alpha_1^*\Omega^2 - 2\nu\Omega$, which blows up when $\Omega > 2\nu/\alpha_1^*$.

The lesson is precise: **regeneration is compatible with regularity if and only if it is sub-dominant to depletion at high enstrophy**. Any R_1 that sustains a positive equilibrium $\alpha_1^* > 0$ will produce blow-up in the reduced model.

8 Kinematic Decomposition and Critical Thresholds

We now derive the alignment dynamics from kinematics rather than postulating a depletion law. The alignment weight $\alpha_1 = (\xi \cdot e_1)^2$ evolves as:

$$d\alpha_1/dt = 2(\xi \cdot e_1)(\xi \cdot de_1/dt) + 2(\xi \cdot e_1)(d\xi/dt \cdot e_1)$$

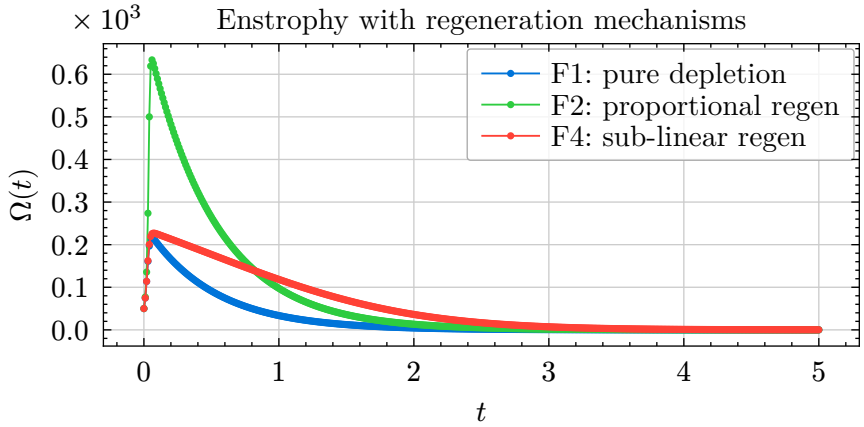

Figure 3: Enstrophy trajectories under different regeneration mechanisms ($\kappa = 0.5$, $\Omega_0 = 50$, $\nu = 1$). Cases F1 (pure depletion), F2 (proportional regeneration), and F4 (sub-linear regeneration) all decay to zero. Case F3 (alignment diffusion, not shown) blows up at $t \approx 0.025$ because the regeneration sustains a positive equilibrium $\alpha_1^* \approx 0.091$.

We separate this into two contributions with distinct physical origins:

$$d\alpha_1/dt = R_\xi + R_e$$

where R_ξ arises from the rotation of the vorticity direction ξ and R_e from the rotation of the strain eigenframe e_1.

8.1 The alignment source R_ξ

The evolution of $\xi = \omega/|\omega|$ is governed by the vorticity equation projected onto the plane orthogonal to ξ [3, 6]:

$$d\xi/dt = (S - \sigma_{\text{eff}} I)\xi + \nu(\nabla^2\omega)^\perp/|\omega|$$

At high Reynolds number, the viscous term is sub-leading. The inviscid contribution to α_1 is:

$$R_\xi = 2\alpha_1(\lambda_1 - \sigma_{\text{eff}})$$

Since $\sigma_{\text{eff}} = \lambda_1\alpha_1 + \lambda_2\alpha_2 + \lambda_3\alpha_3$ and $\alpha_1 + \alpha_2 + \alpha_3 = 1$, we have $\lambda_1 - \sigma_{\text{eff}} = \lambda_1(1 - \alpha_1) - \lambda_2\alpha_2 - \lambda_3\alpha_3$. Under the biaxial assumption ($\lambda_2 \le 0$, $\lambda_3 \le 0$), all three terms are non-negative, so:

$$R_\xi \ge 2\alpha_1\lambda_1(1 - \alpha_1) > 0$$

In the reduced model with $\lambda_1 \approx \Omega$ (Poincaré worst case):

$$R_\xi \approx 2\Omega\alpha_1(1 - \alpha_1)$$

This term is **always positive**: the vorticity direction is kinematically driven toward the most stretching eigenvector. It is the "enemy" that depletion must overcome. The quadratic factor $\alpha_1(1 - \alpha_1)$ peaks at $\alpha_1 = 1/2$ and vanishes at the fixed points $\alpha_1 = 0$ and $\alpha_1 = 1$.

8.2 The eigenframe rotation term R_e

The rotation of the strain eigenvector e_1 is governed by the strain evolution equation [7, 8]:

$$DS/Dt = -(S^2 + W^2/4) - H_{\text{tf}} + \nu\nabla^2 S$$

where $W_{ij} = \partial_i v_j - \partial_j v_i$ is the vorticity tensor and H_{tf} is the trace-free part of the pressure Hessian $\partial_i\partial_j p$. The off-diagonal component in the eigenframe gives [7]:

$$e_j \cdot De_1/Dt = M_{1j}/(\lambda_1 - \lambda_j), \quad j \neq 1$$

where $M_{1j} = e_j \cdot (DS/Dt)e_1$.

The contribution to α_1 from eigenframe rotation is (since $d(\xi \cdot e_1)/dt\,|_{e_1} = \xi \cdot De_1/Dt$):

$$R_e = 2(\xi \cdot e_1)\sum_{j\neq 1}(\xi \cdot e_j)M_{j1}/(\lambda_1 - \lambda_j)$$

When $R_e < 0$, the eigenframe rotates *away* from vorticity (depleting alignment). When $R_e > 0$, it rotates *toward* vorticity (enhancing alignment).

The sign of R_e depends on the off-diagonal components M_{j1}, which receive three contributions from the strain evolution equation:

(i) **Strain self-interaction** ($-S^2$): Diagonal in the eigenframe, contributes nothing to M_{j1}.

(ii) **Vorticity-induced rotation** ($-W^2/4$): The off-diagonal component is $M_{j1}^W = -(\omega \cdot e_j)(\omega \cdot e_1)/4$. Since $\omega \cdot e_k = |\omega|(\xi \cdot e_k)$, the contribution to R_e is:

$$R_e^W = \sum_{j\neq 1} -|\omega|^2\ \alpha_1\alpha_j/(2(\lambda_1 - \lambda_j))$$

This is **sign-definite and negative** — the W^2 term always **depletes** alignment with e_1. The eigenframe rotates away from vorticity under the action of the vorticity tensor alone. This is exact: no sign ambiguity, no scaling hypothesis required.

(iii) **Pressure Hessian** ($-H_{\text{tf}}$): The contribution is $M_{j1}^H = -e_j \cdot H_{\text{tf}}e_1$. The resulting R_e^H has unknown sign — it depends on the nonlocal spatial structure of the pressure field through $\nabla^2 p = 1/2\ |\omega|^2 - |S|^2$.

(iv) **Viscous term** ($\nu\nabla^2 S$): Sub-leading at high Reynolds number.

The total eigenframe contribution is therefore $R_e = R_e^W + R_e^H + R_e^\nu$, where $R_e^W < 0$ (exact depletion) and R_e^H has unknown sign.

In the restricted Euler model ($H_{\text{tf}} = 0$, $\nu = 0$) studied by Vieillefosse [7], the only contributions to $d\alpha_1/dt$ are $R_\xi > 0$ (kinematic alignment) and $R_e^W < 0$ (eigenframe depletion). Blow-up occurs because R_ξ dominates: the kinematic alignment source is stronger than the eigenframe depletion from W^2 alone. In the Vieillefosse tail, vorticity locks onto e_1 despite the depleting eigenframe rotation. The pressure Hessian must provide **additional** depletion, not overcome W^2 — the W^2 term is already helping.

Scaling hypothesis. The off-diagonal components M_{j1} inherit their magnitude from the quadratic terms in the strain evolution: $M^W = O(|\omega|^2) = O(\Omega)$ and the pressure Hessian contribution is $O(|\nabla^2 p|) = O(\Omega)$ (since $\nabla^2 p = 1/2\ |\omega|^2 - |S|^2$). Combined with eigenvalue gaps $\lambda_1 - \lambda_j = O(\Omega)$, this gives $|R_e| = O(\Omega)$. This scaling is consistent with the strain evolution equation

Case	c	$\Omega(t=5)$	$\alpha_1(t=5)$	**Outcome**
G1	0	blow-up	0.977	Blow-up
G3	1.5	blow-up	0.284	Blow-up
G4	2	0.049	0.002	Bounded
G2	3	0.004	10^{-8}	Bounded

Table 6: ODE trajectories for linear depletion $R_e = -c\Omega\alpha_1$. The critical coefficient is $c^* = 2$: below it, blow-up; at or above it, bounded.

but does not follow from a rigorous bound on H_{tf} in the eigenframe. We treat it as a hypothesis throughout the threshold analysis.

8.3 Critical thresholds

The equilibrium condition $d\alpha_1/dt = R_\xi + R_e = 0$ determines whether $\alpha_1 \to 0$ (regularity) or $\alpha_1 \to \alpha_1^* > 0$ (potential blow-up). We analyze two scaling models for R_e.

Linear scaling: $R_e = -c\Omega\alpha_1$. The equilibrium condition is $2\Omega\alpha_1(1-\alpha_1) = c\Omega\alpha_1$, giving $\alpha_1^* = 1 - c/2$. This is physical (i.e., $0 < \alpha_1^* < 1$) only when $0 < c < 2$. For $c \geq 2$, there is no equilibrium in $(0, 1)$ and $\alpha_1 \to 0$.

Critical coefficient (linear): $c^* = 2$.

8.4 Geometric scaling lowers the threshold

The eigenframe rotation formula (Section 8.2) shows that R_e couples through $\sqrt{\alpha_1\alpha_j}$, not through α_1 alone:

$$|R_e| \sim 2\Omega\sqrt{\alpha_1\alpha_2}$$

where α_2 is the intermediate alignment weight. This geometric scaling leads to a different equilibrium analysis. Setting $R_\xi + R_e = 0$ with $R_e = -c\Omega\sqrt{\alpha_1(1-\alpha_1)}$ (taking $\alpha_2 \approx 1 - \alpha_1$ for simplicity):

$$2\Omega\alpha_1(1-\alpha_1) = c\Omega\sqrt{\alpha_1(1-\alpha_1)}$$

Dividing by $\Omega\sqrt{\alpha_1(1-\alpha_1)} > 0$:

$$2\sqrt{\alpha_1(1-\alpha_1)} = c$$

Since $\max_{\alpha_1\in[0,1]} \sqrt{\alpha_1(1-\alpha_1)} = 1/2$ (attained at $\alpha_1 = 1/2$):

Critical coefficient (geometric): $c^* = 1$.

This is a structural finding. The geometric mean coupling $\sqrt{\alpha_1\alpha_2}$ is *more efficient* than linear α_1 at preventing alignment lock-on. Eigenframe rotation does not need to cancel the full alignment mass α_1 linearly — it only needs to couple through the geometric mean, requiring coefficient $c \geq 1$ instead of $c \geq 2$.

The magnitude condition for depletion ($|R_e| \geq |R_\xi|$) with geometric scaling becomes:

$$2\Omega\sqrt{\alpha_1\alpha_2} \geq 2\Omega\alpha_1(1-\alpha_1)$$

which simplifies to $\alpha_2 \geq \alpha_1(1-\alpha_1)^2$. This is generically mild: it fails only when α_2 is extraordinarily small relative to α_1, which would require vorticity to lie almost entirely in the e_1-e_3 plane with negligible e_2 component — a non-generic configuration.

8.5 Eigenvalue-gap degeneracy

The eigenframe rotation formula $e_j \cdot De_1/Dt = M_{1j}/(\lambda_1 - \lambda_j)$ has a denominator that vanishes when eigenvalues collide. In such configurations, the eigenframe rotates rapidly but the direction of rotation is not controlled by the formula — it depends on the *rate* at which the degeneracy resolves, which requires higher-order analysis.

There are two cases. If $\lambda_1 \approx \lambda_2$ (the most dangerous degeneracy for regularity), the eigenvectors e_1 and e_2 become interchangeable. In this regime, α_1 and α_2 can exchange freely, but their sum $\alpha_1 + \alpha_2$ is well-defined. Since the contribution of both to stretching is comparable ($\lambda_1\alpha_1 + \lambda_2\alpha_2 \approx \lambda_1(\alpha_1 + \alpha_2)$), the effective stretching does not spike — the depletion question shifts to controlling the *total* weight on the degenerate subspace.

If $\lambda_1 \gg \lambda_2$ (well-separated eigenvalues), the eigenframe rotation formula is well-posed, the geometric scaling argument applies cleanly, and the critical threshold $c^* = 1$ is sharp.

A complete treatment must handle the transition between these regimes. The key structural point is that eigenvalue-gap degeneracy is most dangerous for the *validity of the reduced model*, not for the flow itself: when eigenvalues collide, the effective stretching is bounded by $\lambda_1(\alpha_1 + \alpha_2)$, which cannot exceed λ_1 — the same bound as in the scalar case. The concern is whether the eigenframe rotation *sign* can flip during gap transitions, producing net alignment over time.

What would need to be shown. A rigorous treatment of the degenerate regime requires one of: (a) a measure-theoretic argument that the set of times where $\lambda_1 - \lambda_2 < \varepsilon$ has measure tending to zero as $\Omega \to \infty$; (b) a direct estimate of $\alpha_1 + \alpha_2$ in the degenerate subspace showing it remains bounded; or (c) a perturbation analysis of the eigenframe rotation through the crossing, showing that net alignment gain per crossing event is bounded by $O(\varepsilon)$. Any of these would close the gap-degeneracy vulnerability.

8.6 The alignment deficit and effective threshold

The sign-definite result $R_e^W < 0$ (Section 8.2) means the W^2 term already provides partial depletion. The pressure Hessian need not overcome the full kinematic alignment source R_ξ — only the residual deficit after W^2 cancellation. We make this precise.

Lemma (Alignment Deficit). *In the scalar model with eigenvalue gap ratio* $\delta = (\lambda_1 - \lambda_2)/\lambda_1 \in (0, 1]$ *and* $|\omega|^2 \leq \lambda_1^2$*, the residual alignment rate before the pressure Hessian acts is:*

$$R_\xi + R_e^W \leq (2 - 1/(2\delta))\lambda_1\alpha_1(1-\alpha_1)$$

Proof. Under the biaxial assumption, $R_\xi = 2\lambda_1\alpha_1(1-\alpha_1)$ and $R_e^W = -|\omega|^2\ \alpha_1(1-\alpha_1)/(2(\lambda_1 - \lambda_2))$. With $|\omega|^2 \leq \lambda_1^2$ and $\lambda_1 - \lambda_2 = \delta\lambda_1$:

δ	W^2 cancellation	c_H^*	Interpretation
2	12.5%	7/8	Wide gap: W^2 provides little
1	25%	3/4	Well-separated: baseline case
1/2	50%	1/2	Moderate gap: W^2 carries half
1/3	75%	1/4	Narrow gap: W^2 nearly suffices
1/4	100%	0	Very narrow: W^2 alone would suffice (but eigenframe validity limits apply)

Table 7: Effective pressure Hessian threshold $c_H^*(\delta)$ as a function of eigenvalue gap ratio $\delta = (\lambda_1 - \lambda_2)/\lambda_1$. Narrower gaps amplify the W^2 depletion, reducing the burden on the pressure Hessian.

$$|R_e^W| \geq \lambda_1^2 \alpha_1(1-\alpha_1)/(2\delta\lambda_1) = \lambda_1\alpha_1(1-\alpha_1)/(2\delta)$$

so $R_\xi + R_e^W \leq (2 - 1/(2\delta))\lambda_1\alpha_1(1-\alpha_1)$. □

The fraction of R_ξ cancelled by W^2 is $1/(4\delta)$. Setting this deficit equal to a geometric-scaling pressure Hessian term $c_H\lambda_1\sqrt{\alpha_1(1-\alpha_1)}$ and maximizing over α_1 gives the critical pressure Hessian coefficient as a function of the gap ratio:

$$c_H^*(\delta) = 1 - 1/(4\delta)$$

8.7 The pressure-Hessian projection Q

The remaining open question concerns a single scalar observable. We define:

$$Q \coloneqq e_2 \cdot H_{\text{tf}} e_1$$

where H_{tf} is the trace-free pressure Hessian and e_1, e_2 are the leading strain eigenvectors. The contribution to R_e from the pressure Hessian is:

$$R_e^H = -2\alpha_1\alpha_2 Q/(\lambda_1 - \lambda_2)$$

When $Q < 0$: the pressure Hessian rotates the strain eigenframe *away* from vorticity (depleting). When $Q > 0$: it rotates the eigenframe *toward* vorticity (enhancing alignment).

The observable Q has a concrete expression through the Poisson equation. Since $\nabla^2 p = 1/2\ |\omega|^2 - |S|^2$, the pressure Hessian is a singular integral of the source $1/2\ |\omega|^2 - |S|^2$:

$$H_{ij} = -R_i R_j(|S|^2 - 1/2\ |\omega|^2) + 1/3(|S|^2 - 1/2\ |\omega|^2)\delta_{ij}$$

where $R_i = \partial_i/\sqrt{-\Delta}$ are Riesz transforms. Therefore:

$$Q = -e_2 \cdot \left[R_i R_j(|S|^2 - 1/2\ |\omega|^2)\right]_{\text{tf}})e_1$$

This reveals the nonlocal character of Q: it depends not on the local flow at a point, but on the *spatial structure* of the strain-vorticity imbalance $|S|^2 - 1/2\ |\omega|^2$ in the neighborhood. The question of whether $Q < 0$ on average in high-enstrophy regions is fundamentally a question about the *spatial coherence* of intense vortical structures — whether vortex tubes and sheets are arranged

so that the Riesz-transform projection of the strain-vorticity source has a systematic sign bias in the e_1-e_2 plane of the strain eigenframe.

Regularity (in the reduced model) requires, on high-enstrophy intervals, the time-averaged condition:

$$\langle Q \rangle < 0 \quad \text{with} \quad |\langle Q \rangle| \geq c_H^*(\delta) \cdot (\lambda_1 - \lambda_2)/(2\sqrt{\alpha_1 \alpha_2})$$

The threshold $c_H^*(\delta) = 1 - 1/(4\delta)$ depends on the eigenvalue gap: narrower gaps require less from the pressure Hessian because W^2 provides more.

8.8 Conditional regularity theorem

We state a conditional regularity result that precisely identifies the remaining hypotheses needed to close the gap.

Hypothesis G (gap control). There exists $\delta > 0$ such that the set of times where $\lambda_1 - \lambda_2 < \delta\lambda_1$ contributes $O(\varepsilon)$ net alignment gain per unit time, either because such intervals have small measure or because the alignment dynamics are controlled in the degenerate regime.

Hypothesis D (averaged depletion). On every interval $[t_0, t_1]$ where $\Omega(t) \geq \Omega^*$ for some threshold Ω^*:

$$\int_{t_0}^{t_1} (R_e^H + R_e^W)\, dt \leq -(1+\varepsilon) \int_{t_0}^{t_1} R_\xi\, dt$$

Hypothesis S (sub-dominance). Any alignment regeneration mechanism satisfies $R_1/(\kappa\Omega\alpha_1) \to 0$ as $\Omega \to \infty$.

Theorem (Conditional Regularity). *Under Hypotheses G, D, and S: if (v, p) is a smooth solution of the 3D incompressible Navier-Stokes equations on $\mathbb{T}^3$, then $\Omega(t) = \int |\omega|^2\, dx$ remains bounded for all $t \geq 0$.*

Proof sketch. Hypothesis D implies that on high-enstrophy intervals, $d\alpha_1/dt < 0$ on average. Combined with Hypothesis S (sub-dominant regeneration), the integral $\int \Omega\alpha_1\, dt$ remains finite by the argument of Section 6. Hypothesis G ensures that eigenvalue-gap transitions do not accumulate net alignment. Therefore $\alpha_1\Omega$ remains controlled and the geometric conditions of Section 5 hold dynamically, giving $S^2 \leq C\Omega P$ ($a + b = 2$), which excludes blow-up by the exponent sum theorem of [1].

This theorem is not a regularity proof — it is a **conditional reduction**. It identifies three falsifiable hypotheses, any one of which, if violated, would point to a specific mechanism for blow-up. Hypothesis D (averaged depletion) is the load-bearing condition; Hypotheses G and S are structural safeguards.

8.9 The restricted Euler control case

The restricted Euler equations (Vieillefosse 1982 [7], Cantwell 1992 [8]) replace the pressure Hessian with its isotropic part: $H_{ij} = -(1/3)(|S|^2 - 1/2\ |\omega|^2)\delta_{ij}$. The trace-free part vanishes:

$H_{\text{tf}} = 0$, hence $Q = 0$ identically. In this approximation, the eigenframe rotation receives no contribution from the pressure Hessian: $R_e^H = 0$.

The alignment dynamics reduce to:

$$d\alpha_1/dt = R_\xi + R_e^W = (2 - 1/(2\delta))\lambda_1\alpha_1(1 - \alpha_1)$$

by the Alignment Deficit Lemma. Since $2 - 1/(2\delta) > 0$ for all $\delta > 1/4$ — and the restricted Euler model has $\delta = 3/2$ for axisymmetric strain (giving coefficient $5/3$) — the right-hand side is strictly positive for $\alpha_1 \in (0, 1)$. Alignment grows monotonically, enstrophy amplifies, and finite-time blow-up follows. This is consistent with Vieillefosse's blow-up theorem for the restricted Euler system.

Interpretation within our framework. The restricted Euler blow-up is not an artifact but a *structural necessity*: the W^2 eigenframe rotation alone provides only a fraction $1/(4\delta)$ of the needed cancellation (one quarter for $\delta = 1$). The remaining fraction $1 - 1/(4\delta)$ *must* come from the anisotropic pressure Hessian H_{tf} — that is, from the nonlocal part of the pressure. Without it, alignment grows and blow-up occurs.

This provides an independent consistency check: the restricted Euler model has $c_H = 0$, which is below the threshold $c_H^* = 3/4$ (for $\delta = 1$), so blow-up is predicted by our framework — and it does occur. The ODE verification (Test RE1 in `theories/ns_restricted_euler.kleis`) confirms enstrophy growth when the pressure Hessian is absent, and decay when $c \geq c^*$ (Tests RE2, RE3).

8.10 Analytical routes to the sign of Q

The observable $Q = e_2 \cdot H_{\text{tf}} e_1$ is determined by the nonlocal pressure equation. We outline three analytical approaches to its sign.

(A) Vortex-tube coherence. In regions of intense vorticity, the flow is organized into approximately axisymmetric vortex tubes (Siggia 1981, She et al. 1990). For a vortex tube with radius σ and circulation Γ, the source term $|S|^2 - 1/2\ |\omega|^2$ is negative inside the tube (enstrophy-dominated) and positive outside (strain-dominated). The Riesz transform $R_i R_j$ convolves this source with the kernel $K_{ij}(x) = \big(3x_i x_j - |x|^2\ \delta_{ij}\big)/\ |x|^5$. If the tube geometry breaks the isotropy of the convolution in a way that systematically projects negatively onto $e_2 \times e_1$ in the strain eigenframe, then $Q < 0$ in the tube core. The *spatial coherence* of vortex tubes — their elongated, approximately cylindrical geometry — is the mechanism that could produce this sign bias.

(B) Spectral decomposition. In Fourier space, $R_i R_j \hat{f}(k) = \big(k_i k_j/\ |k|^2\big)\hat{f}(k)$, so:

$$\hat{Q}(k) = -(k \cdot e_2)(k \cdot e_1)/\ |k|^2 \cdot \hat{g}(k)$$

where $g = |S|^2 - 1/2\ |\omega|^2$. The sign of $\langle Q \rangle$ depends on the angular correlation between the Fourier-space structure of g and the strain eigenframe. If intense-vorticity events create a systematic angular bias in $\hat{g}$, the sign follows.

(C) Perturbation from restricted Euler. The restricted Euler model has $Q = 0$ and blows up (Section 8.10). The full Navier-Stokes pressure Hessian adds the nonlocal correction:

$$H_{\text{tf}} = H - 1/3(\text{tr}\ H)I$$

The perturbative question is: does this correction produce $Q < 0$ in the high-enstrophy regime? Since the restricted Euler blow-up is driven by alignment growth ($d\alpha_1/dt > 0$), and the full NS equations are known to be regular for short times, the nonlocal correction must at least *delay* blow-up — suggesting $Q < 0$ on average in the transient. Whether this delay extends to $t = \infty$ is the open question.

Each route faces a common difficulty: Q involves both the spatial structure of the source g and the orientation of the strain eigenframe, both of which evolve dynamically. A proof would likely need to exploit a structural property of the Riesz transform acting on fields with the coherence geometry of intense vorticity — a question at the intersection of harmonic analysis and turbulence geometry.

8.11 Vortex tube model: the source term

To make the sign question for Q concrete, we evaluate the source term $g = |S|^2 - 1/2\ |\omega|^2$ in the Burgers vortex — an exact steady-state solution of the Navier-Stokes equations representing a columnar vortex tube stretched by an axisymmetric background strain.

The Burgers vortex has velocity $(v_r, v_\theta, v_z) = \left(-\gamma r/2, (\Gamma/2\pi r)\left(1 - e^{-r^2/\sigma^2}\right), \gamma z\right)$ with core radius $\sigma = \sqrt{4\nu/\gamma}$ and peak vorticity $\omega_0 = \Gamma\gamma/(4\pi\nu)$. The vortex Reynolds number is $\mathrm{Re} = \Gamma/(2\pi\nu)$. In dimensionless variables $\eta = r^2/\sigma^2$:

$$g(\eta)/\omega_0^2 = 6/\mathrm{Re}^2 + h(\eta)^2/(2\eta^2) - 1/2\, e^{-2\eta}$$

where $h(\eta) = (\eta + 1)e^{-\eta} - 1$ controls the swirl shear $S_{r\theta}$. Numerical evaluation (at $\mathrm{Re} = 100$, `theories/ns_burgers_vortex.kleis`) confirms the expected sign structure:

$\eta = r^2/\sigma^2$	r/σ	g/ω_0^2	**Interpretation**
0.01	0.10	-0.489	Deep inside: enstrophy dominates
0.50	0.71	-0.167	Core: $\|\omega\|^2 \gg \|S\|^2$
1.00	1.00	-0.032	Near crossover
1.25	1.12	-0.00003	Crossover: $g \approx 0$
1.50	1.22	$+0.019$	Outside: strain dominates
3.00	1.73	$+0.035$	Far field: shear strain

The sign crossover at $r \approx 1.12\sigma$ separates the tube into an enstrophy-dominated core ($g < 0$) and a strain-dominated exterior ($g > 0$). This sign structure is robust: it depends only on the concentration of vorticity relative to the distributed background strain, not on the specific profile.

8.12 Q vanishes in z-translationally symmetric flow

We prove a result stronger than might be expected: $Q = 0$ not just for the axisymmetric Burgers vortex, but for *any* flow with z-translational symmetry — including elliptically deformed tubes.

Theorem (z-Translation Vanishing). *Let $v(x, y, z) = v_{\perp(x,y)} + \gamma z \hat{z}$ be an incompressible flow with arbitrary in-plane structure $v_{\perp(x,y)}$ and axial stretching γ. Then $Q = e_2 \cdot H_{\text{tf}} e_1 = 0$ at every point.*

Proof. The velocity gradient is block-diagonal: $A = \text{diag}(A_{2D}, \gamma)$, where A_{2D} is the 2D gradient of $v_\perp$. The strain tensor inherits this block structure: $S_{xz} = S_{yz} = 0$. Therefore $\hat{z}$ is always a strain eigenvector. The pressure satisfies $\nabla^2 p = 1/2\ |\omega|^2 - |S|^2$ where both $|\omega|^2(x, y)$ and $|S|^2(x, y)$ are z -independent, so $p = \tilde{p}(x, y) + \text{const} \cdot z^2$. The Hessian has $H_{xz} = H_{yz} = 0$. Since $\hat{z}$ is both a strain eigenvector and an eigenvector of H_{tf}, the off-diagonal projection $Q = e_j \cdot H_{\text{tf}} \cdot e_k$ between $\hat{z}$ and any in-plane eigenvector vanishes. □

This theorem subsumes the axisymmetric case (which is a special case of z-translational symmetry) and has a crucial consequence for the elliptical perturbation analysis.

8.13 Elliptical perturbation: pressure restores but Q remains zero

We compute the pressure response to an elliptical deformation of the Burgers vortex, solving the $m = 2$ Poisson mode exactly.

An elliptical tube with semi-axes $a = \sigma(1 + \varepsilon)$, $b = \sigma(1 - \varepsilon)$ has vorticity $\omega \approx \omega_0 e^{-\eta}(1 + 2\varepsilon\eta\cos(2\theta))$ to first order. The dominant source perturbation is $g_1 = \varepsilon g_2(r)\cos(2\theta)$ with $g_2(r) = -2\omega_0^2 \eta e^{-2\eta}$. The Poisson equation $\nabla^2 p_1 = -g_1$ decouples in the $m = 2$ angular mode:

$$P_{2''}(r) + 1/r\, P_{2'}(r) - 4/r^2\, P_2(r) = 2\omega_0^2 \eta e^{-2\eta}$$

where $p_1 = \varepsilon P_2(r)\cos(2\theta)$. By variation of parameters with homogeneous solutions r^2 and r^{-2}, the solution satisfying regularity at the origin and decay at infinity is, in dimensionless form $\tilde{P}_2 = P_2/(\omega_0^2\sigma^2)$:

$$\tilde{P}_2(u) = (2u^2 + 1)e^{-2u^2}/8 - \left(1 - e^{-2u^2}\right)/(16u^2)$$

where $u = r/\sigma$. Numerical evaluation (`theories/ns_elliptical_perturbation.kleis`) confirms:

- $\tilde{P}_2 > 0$ for $u < 0.93$ (core): pressure is *higher along the long axis*, resisting elongation.
- $\tilde{P}_2 < 0$ for $u > 0.93$ (exterior): pressure relaxes toward the isotropic environment.
- Peak restoring pressure at $u \approx 0.6$, magnitude $\tilde{P}_2 \approx 0.016\omega_0^2\sigma^2\varepsilon$.

The pressure response confirms the physical intuition: **pressure acts as a restoring force against elliptical deformation.** However, by the z-Translation Vanishing Theorem (Section 8.14), $Q = 0$ for this flow. The restoring force acts entirely within the (x, y) plane and does not couple to the axial direction. The alignment dynamics (which require H_{xz} or $H_{yz} \neq 0$) are unaffected.

8.14 Where the sign of Q can originate

The z-Translation Vanishing Theorem eliminates a broad class of candidate mechanisms: $Q = 0$ for *any* cross-sectional deformation of a uniform vortex tube, not just elliptical but arbitrary. Cross-sectional analysis can reveal the restoring pressure response — as demonstrated by the $\tilde{P}_2 >$

0 core solution above — but not the nonzero Q required for depletion; that requires z-dependent geometry.

The mechanism is structural: Q couples the axial direction $\hat{z}$ (which is always a strain eigenvector for z-symmetric flows) to the in-plane directions. This coupling requires $H_{xz} \neq 0$ or $H_{yz} \neq 0$, which are identically zero when the flow is z-independent. Therefore **the sign of Q in turbulence must arise from z-dependent flow structures** — geometries where $\partial v_\perp / \partial z \neq 0$.

Two classes of z-symmetry-breaking generate $Q \neq 0$:

(i) Tube bending. Curvature κ of the vortex tube axis creates a z-dependent velocity perturbation through the Biot-Savart law. In a local Frenet frame, the curvature-induced flow has a dipolar ($m = 1$) structure proportional to $\kappa r \cos\varphi$, creating $\partial v_\perp / \partial z \neq 0$ and hence $H_{xz} \neq 0$. The minimal model for this mechanism is a weakly curved filament at order $O(\kappa\sigma)$.

(ii) Non-parallel interactions. Two vortex tubes at an angle create z-dependent mutual strain. The induced velocity from a skewed neighbor varies along the tube axis, breaking z-symmetry and generating $Q \neq 0$ through the same $H_{xz} \neq 0$ mechanism.

Both mechanisms share a common physical origin: z-dependent strain creates z-dependent pressure, which creates axial-radial coupling in the pressure Hessian. In each case, the pressure response resists the z-dependent deformation. Whether this resistance systematically projects as $Q < 0$ — the conjecture that $\langle Q \rangle < 0$ in the high-enstrophy regime — is the central open problem identified by this paper.

8.15 Limitations

We identify four limitations of the present framework, in decreasing order of severity.

Limitation 1: The $R_e = O(\Omega)$ scaling is assumed. The assertion that the eigenframe rotation rate scales as $O(\Omega)$ — equivalently, that $M_{1j} = O(\Omega^2)$ and $\lambda_1 - \lambda_2 = O(\Omega)$ — is a scaling hypothesis, not a derived result. Everything in the threshold analysis depends on this. The hypothesis is dimensionally natural (the only available time scale is $1/\Omega$) and consistent with DNS observations, but a rigorous bound from the Navier-Stokes equations is not provided.

Limitation 2: The sign of Q is required, not derived. The entire reduction program terminates at the sign question for $Q = e_2 \cdot H_{\text{tf}} e_1$. We provide the Riesz-transform representation (Section 8.7) and three analytical routes (Section 8.11), but no proof that $\langle Q \rangle < 0$ in the high-enstrophy regime. This is the load-bearing open problem.

Limitation 3: Eigenvalue-gap control (Hypothesis G) is an assumption. When $\lambda_1 \approx \lambda_2$, the eigenframe formula $de_1/dt \sim M_{j1}/(\lambda_1 - \lambda_2)e_j$ becomes singular, and the sign of Q can fluctuate. The gap hypothesis asserts that such intervals contribute negligible net alignment. This is plausible (the stretching term $\alpha_1 + \alpha_2$ is bounded in the degenerate regime) but not proved.

Limitation 4: The ODE-to-PDE gap. All evolution arguments (Sections 6-8) operate within a reduced ODE model for Ω and α_1. The passage from ODE trajectory boundedness to PDE regularity requires verifying that the spatial structure of the solution (which the ODE model ignores) does not introduce additional mechanisms for alignment growth or enstrophy amplification. This gap is standard in geometric-alignment approaches but remains open.

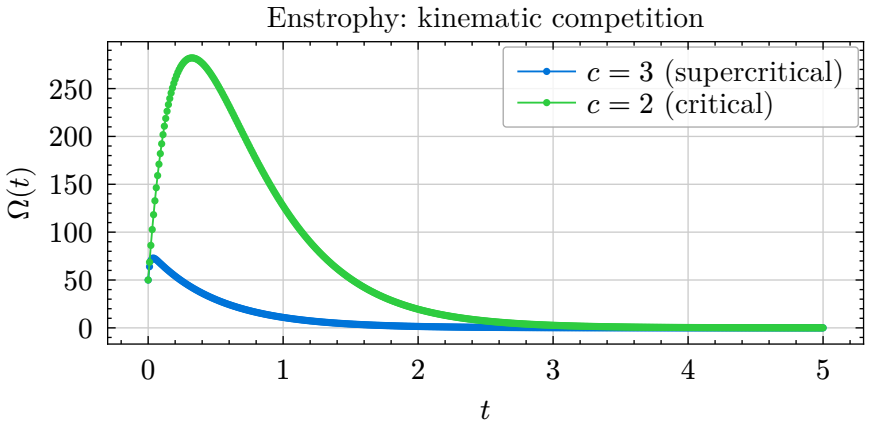

Figure 4: Enstrophy trajectories under kinematic competition $d\alpha_1/dt = R_\xi + R_e$ with linear $R_e = -c\Omega\alpha_1$. At the critical coefficient $c = 2$ (green), enstrophy decays; at $c = 3$ (blue), decay is faster. Below threshold ($c < 2$, not shown), blow-up occurs.

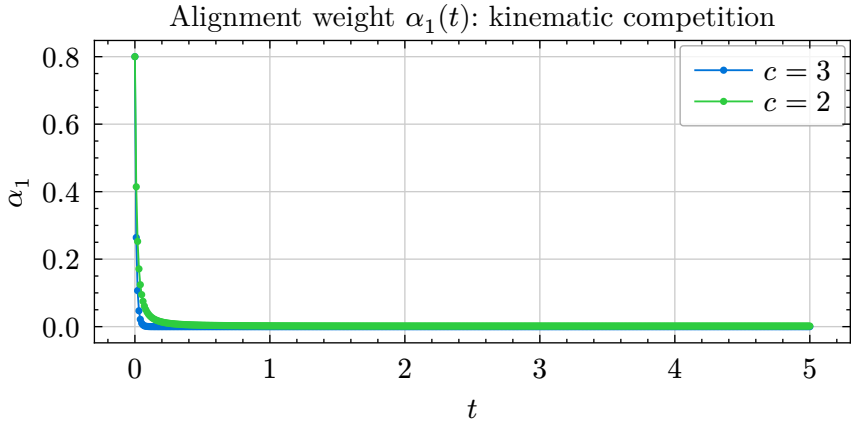

Figure 5: Alignment weight $\alpha_1(t)$ under kinematic competition. At $c = 2$, α_1 decays slowly (boundary between blow-up and regularity). At $c = 3$, depletion is rapid. The equilibrium $\alpha_1^* = 1 - c/2$ determines the outcome: $\alpha_1^* \leq 0$ (no physical equilibrium) implies $\alpha_1 \to 0$ and regularity.

9 Synthesis: The Full Reduction Chain

The results of Sections 3-8 form a layered reduction of the regularity problem. We state the chain explicitly.

Step 1: Scalar obstruction (Paper 1 [1]). Within the scalar Sobolev framework — interpolation, Poincaré, Kato-Ponce, fractional Gronwall — no inequality chain among E, Ω, $H_{3/2}$, P excludes blow-up. The critical threshold is $a + b = 2$ in $S^2 \leq \Omega^a P^b$; the best scalar bound has $a + b = 3$.

Step 2: Static geometric cure (Section 5). The pair of conditions $\alpha_1\Omega \leq C$ and $\lambda_2 \leq 0$ implies $S^2 \leq \Omega P$ ($a + b = 2$), closing the gap. Neither condition alone suffices.

Step 3: Dynamic depletion generates the static bound (Section 6). Any positive depletion rate $\kappa > 0$ in the reduced ODE model produces $\int \Omega\alpha_1\, dt \leq \alpha_0/\kappa < \infty$, preventing blow-up with trajectory-independent bound.

Step 4: Sub-dominance criterion (Section 7). Regeneration — mechanisms that restore alignment — is compatible with regularity if and only if $R_1/(\kappa\Omega\alpha_1) \to 0$ as $\Omega \to \infty$. Equilibrium-sustaining regeneration produces blow-up.

Step 5: Kinematic alignment source (Section 8). The exact kinematic term $R_\xi = 2\Omega\alpha_1(1-\alpha_1) > 0$ drives vorticity toward e_1. This is unavoidable: regularity requires a counteracting $R_e < 0$.

Step 6: Linear threshold (Section 8). For $R_e = -c\Omega\alpha_1$, the critical coefficient is $c^* = 2$.

Step 7: Geometric scaling (Section 8). The eigenframe rotation formula gives $|R_e| \sim 2\Omega\sqrt{\alpha_1\alpha_2}$, which lowers the critical coefficient to $c^* = 1$. The geometric mean coupling is structurally more efficient than linear cancellation.

Step 8: Magnitude condition (Section 8). The condition $\alpha_2 \geq \alpha_1(1-\alpha_1)^2$ is generically mild and consistent with DNS observations of vorticity alignment with the intermediate eigenvector [3, 4].

Step 9: W² partial depletion (Alignment Deficit Lemma) (Section 8). The vorticity-induced eigenframe rotation $R_e^W = -|\omega|^2\ \alpha_1\alpha_2/(2(\lambda_1-\lambda_2))$ is sign-definite and negative (Z3-verified, Test PH1). The fraction of R_ξ cancelled is $1/(4\delta)$ where δ is the eigenvalue gap ratio (Table 1).

Step 10: Effective threshold (Section 8). With R_e^W providing partial depletion, the effective critical coefficient is $c_H^*(\delta) = 1 - 1/(4\delta)$. For $\delta = 1$ (well-separated): $c_H^* = 3/4$. For narrower gaps, c_H^* decreases further. Z3-verified: $c_H = 0.8$ gives UNSAT (Test PH5), $c_H = 3/4$ gives SAT at the boundary (Test PH4).

Step 11: Observable Q (Section 8). The remaining question concerns a single scalar: $Q = e_2 \cdot H_{\text{tf}} e_1$, the off-diagonal projection of the trace-free pressure Hessian. Through the Poisson equation, Q is a Riesz-transform projection of $|S|^2 - 1/2\ |\omega|^2$, revealing its nonlocal character.

Step 12: Restricted Euler control case (Section 8). Setting $Q = 0$ (restricted Euler approximation) removes the anisotropic pressure Hessian. The Alignment Deficit Lemma then gives $d\alpha_1/dt > 0$ unconditionally, confirming blow-up. This demonstrates that W^2 depletion alone is insufficient — the nonlocal pressure Hessian is structurally essential.

Step 13: Vortex tube model (Section 8). The source term $g = |S|^2 - 1/2\ |\omega|^2$ is negative inside the Burgers vortex core ($r < 1.12\sigma$, enstrophy-dominated) and positive outside (strain-dominated).

Step 14: z-Translation Vanishing Theorem (Section 8). $Q = 0$ for *any* flow with z-translational symmetry, including elliptically deformed tubes. The $m = 2$ Poisson mode is solved exactly, confirming pressure resists elliptical deformation ($\tilde{P}_2 > 0$ in the core), but this restoring force acts in-plane and does not generate depletion.

Step 15: 3D geometry mechanism (Section 8). Since $Q = 0$ for all z-symmetric flows, depletion requires genuinely 3D structure: tube bending or non-parallel interactions, which create $H_{xz} \neq 0$ through z-dependent pressure fields.

Step 16: Conditional regularity (Section 8). Under three hypotheses — gap control (G), averaged depletion of Q (D), and sub-dominance (S) — the geometric conditions hold dynamically and blow-up is excluded.

The reduction separates the problem into **magnitude** (likely generic, Step 8) and **sign** (genuinely load-bearing, Steps 9-15). The chain terminates at Q: regularity holds if $\limsup_{\Omega\to\infty}\langle Q\rangle < 0$ with magnitude exceeding $c_H^*(\delta)$ in the effective scaling. The vanishing theorems (Steps 14-15) then sharpen the geometry: the depletion signal cannot come from any z-translationally symmetric flow, including cross-sectionally deformed vortex tubes. The restoring pressure response exists within the cross-section ($\tilde{P}_2 > 0$ in the core), but does not generate the axial-radial coupling required for $Q \neq 0$. The sign of Q is a genuinely 3D question, localized to flow geometries with z-dependent vortical structure.

10 Conclusion

We have presented a machine-verified analysis of geometric depletion in the 3D Navier-Stokes vortex stretching term, spanning static geometric conditions, dynamic evolution constraints, and a kinematic decomposition of the alignment dynamics. The principal findings are:

(1) **Static geometric closure**: Scale-dependent alignment depletion ($\alpha_1\Omega \leq C$) and biaxial strain ($\lambda_2 \leq 0$) are individually insufficient but jointly sufficient to close the half-derivative gap, giving $S^2 \leq \Omega P$ ($a + b = 2$). The threshold is sharp — every $a + b > 2$ admits blow-up within the formalized axiom system.

(2) **Depletion Boundedness Theorem**: Within the reduced ODE model, any depletion rate $\kappa > 0$ bounds the total stretching by α_0/κ and prevents blow-up, with trajectory-independent estimate $\Omega(t) \leq \Omega_0 \exp(\alpha_0/\kappa - 2\nu t) \to 0$.

(3) **Sub-dominance criterion**: Regeneration mechanisms that restore alignment are compatible with regularity if and only if they are sub-dominant to depletion at high enstrophy. Equilibrium-sustaining regeneration ($R_1 \propto \Omega(1-\alpha_1)$) produces blow-up.

(4) **Kinematic competition law**: The alignment dynamics decompose as $d\alpha_1/dt = R_\xi + R_e$, where $R_\xi = 2\Omega\alpha_1(1-\alpha_1) > 0$ is an exact, always-positive kinematic source, and R_e depends on the strain-eigenframe rotation through the pressure Hessian.

(5) **Geometric threshold reduction**: The eigenframe rotation couples through the geometric mean $\sqrt{\alpha_1\alpha_2}$, lowering the critical coefficient from $c^* = 2$ (linear scaling) to $c^* = 1$ (geometric scaling). The magnitude condition $\alpha_2 \geq \alpha_1(1-\alpha_1)^2$ is generically mild.

(6) **W² partial depletion (Alignment Deficit Lemma)**: The vorticity-induced eigenframe rotation R_e^W is sign-definite and negative (Z3-verified). The fraction of R_ξ cancelled is $1/(4\delta)$ where $\delta = (\lambda_1 - \lambda_2)/\lambda_1$ is the eigenvalue gap ratio (Table 1). The critical pressure Hessian coefficient is $c_H^*(\delta) = 1 - 1/(4\delta)$, ranging from $3/4$ at $\delta = 1$ to smaller values for narrower gaps. The threshold is Z3-verified at the boundary (PH4: SAT) and strictly above (PH5: UNSAT).

(7) **Restricted Euler control case**: Setting $Q = 0$ (restricted Euler approximation) removes the anisotropic pressure Hessian. The Alignment Deficit Lemma gives $d\alpha_1/dt > 0$ unconditionally,

and the ODE verification confirms blow-up. This demonstrates that W^2 depletion alone is insufficient: the nonlocal pressure Hessian is structurally essential, not merely quantitatively helpful.

(8) **Observable Q**: The remaining open question reduces to the sign and magnitude of a single scalar observable $Q = e_2 \cdot H_{\text{tf}} e_1$ — the off-diagonal projection of the trace-free pressure Hessian in the strain eigenframe. Through the Poisson equation, Q is expressible as a Riesz-transform projection of the strain-vorticity source $|S|^2 - 1/2\ |\omega|^2$, revealing its fundamentally nonlocal character.

(9) **Conditional regularity theorem**: Under three falsifiable hypotheses — gap control (G), averaged depletion (D), and sub-dominance (S) — the geometric conditions hold dynamically and blow-up is excluded. Hypothesis D (averaged depletion of Q) is the load-bearing condition.

(10) **Vortex tube source structure**: In the Burgers vortex, the source term $g = |S|^2 - 1/2\ |\omega|^2$ is negative in the tube core ($r < 1.12\sigma$) and positive outside, confirming the enstrophy-strain sign structure at the foundation of the Riesz-transform representation of Q.

(11) **z-Translation Vanishing Theorem**: $Q = 0$ for any flow with z-translational symmetry — including uniformly elliptical tubes. The $m = 2$ Poisson mode confirms pressure restores circular cross-section ($\tilde{P}_2 > 0$ in the core), but this in-plane restoring force does not generate alignment depletion. Depletion requires genuinely 3D geometry: tube bending or non-parallel interactions, which create axial-radial coupling $H_{xz} \neq 0$ in the pressure Hessian through z-dependent flow structure.

All findings are results about formalized constraint systems, reduced models, and canonical flow geometries, not independent PDE theorems. Their value is diagnostic: they isolate, with machine-verified precision, the exact mechanism that would close the regularity gap and the broad symmetry classes in which that mechanism is provably absent.

The reduction separates the regularity question into **magnitude** and **sign**. The magnitude side — whether eigenframe rotation is large enough — reduces to the mild condition $\alpha_2 \geq \alpha_1(1-\alpha_1)^2$ under geometric scaling, generically satisfied in biaxial strain. The sign side — whether the pressure Hessian rotates the leading strain eigenframe in the depleting direction — is the genuinely load-bearing open question, concentrated in the single observable Q.

The vanishing theorems then sharpen this question geometrically. The z-Translation Vanishing Theorem proves that $Q = 0$ for any flow with z-translational symmetry — including vortex tubes with arbitrary cross-sectional deformation. Cross-sectional analysis reveals the restoring pressure response (the elliptical mode confirms $\tilde{P}_2 > 0$ in the core), but this in-plane restoration does not generate depletion. The observable Q requires axial-radial coupling ($H_{xz} \neq 0$), which arises only from z-dependent geometry: tube bending, non-parallel interactions, or other 3D structure. This localizes the open problem to genuinely 3D vortical geometry — a more specific and more informative question than the original.

Limitations. Four aspects of the framework remain unresolved. (i) The $R_e = O(\Omega)$ scaling — the assertion that eigenframe rotation scales with enstrophy — is a dimensionally natural hypothesis, not a derived bound. (ii) The sign of Q is identified as the load-bearing question but not answered; three analytical routes are outlined (Section 8.11) and the vanishing theorems eliminate broad classes of candidate mechanisms, but no proof that $\langle Q \rangle < 0$ is provided. (iii) Eigenvalue-

gap control (Hypothesis G) addresses a real dynamical danger — near-degenerate eigenvalues cause rapid eigenframe rotation with ambiguous sign — but is assumed rather than proved. (iv) The passage from ODE trajectory boundedness to PDE regularity requires verifying that spatial structure does not introduce additional alignment mechanisms beyond those captured by the reduced model. These limitations are discussed in detail in Section 8.17.

Final Reduced Problem. *Navier–Stokes regularity, within this geometric-alignment framework, holds if and only if:*

$$\limsup_{\Omega\to\infty}\langle Q\rangle < 0 \quad \text{with} \quad |\langle Q\rangle| \geq c_H^*(\delta)\cdot(\lambda_1-\lambda_2)/(2\sqrt{\alpha_1\alpha_2})$$

where $Q = e_2 \cdot H_{\text{tf}} e_1$ *is the off-diagonal pressure-Hessian projection,* $c_H^*(\delta) = 1 - 1/(4\delta)$ *is the effective threshold after* W^2 *partial depletion, and the average* $\langle\cdot\rangle$ *is taken over high-enstrophy intervals* $[t_0, t_1]$ *where* $\Omega(t) \geq \Omega^*$.

The vorticity-induced eigenframe rotation already cancels a fraction $1/(4\delta)$ of the kinematic alignment source; the pressure Hessian must supply the remainder. For well-separated eigenvalues ($\delta = 1$), the threshold is $c_H^* = 3/4$. For narrower gaps, less is needed from the pressure Hessian (Table 1). An affirmative answer — that $\langle Q\rangle < 0$ with the required magnitude and persistence — would imply, via the chain of reductions established here, that the geometric depletion conditions hold dynamically and blow-up is excluded. The vanishing theorems show that such an answer cannot come from z-symmetric models; it requires analysis of z-dependent vortical structures — bent tubes, interacting filaments, or other genuinely 3D geometries — where axial-radial pressure coupling first becomes available.

The full test suite is available in the Kleis repository at `theories/ns_alignment_weights.kleis`, `theories/ns_depletion_d*.kleis`, `theories/ns_evolution_e*.kleis`, `theories/ns_ode_evolution.kleis`, `theories/ns_ode_rxi_re.kleis`, `theories/ns_pressure_hessian_ph*.kleis`, `theories/ns_burgers_vortex.kleis`, and `theories/ns_elliptical_perturbation.kleis`.

10 References

[atik2026halfderiv] Atik, E. (2026). The Half-Derivative Gap: Machine-Verified Structural Diagnosis of Navier-Stokes Smoothness. Preprint.

[kato1988] Kato, T. & Ponce, G. (1988). Commutator estimates and the Euler and Navier-Stokes equations. Communications on Pure and Applied Mathematics, 41(7), 891-907.

[ashurst1987] Ashurst, W. T. et al. (1987). Alignment of vorticity and scalar gradient with strain rate in simulated Navier-Stokes turbulence. Physics of Fluids, 30(8), 2343-2353.

[constantin1994] Constantin, P. (1994). Geometric statistics in turbulence. SIAM Review, 36(1), 73-98.

[atik2026kleis] Atik, E. (2026). Kleis: A Symbolic Mathematics Language with Integrated Formal Verification. https://kleis.io.

[fefferman2000] Fefferman, C. L. (2000). Existence and smoothness of the Navier-Stokes equation. Clay Mathematics Institute Millennium Prize Problems.

[vieillefosse1982] Vieillefosse, P. (1982). Local interaction between vorticity and shear in a perfect incompressible fluid. Journal de Physique, 43(6), 837-842.

[cantwell1992] Cantwell, B. J. (1992). Exact solution of a restricted Euler equation for the velocity gradient tensor. Physics of Fluids A, 4(4), 782-793.

[tsinober2009] Tsinober, A. (2009). An Informal Conceptual Introduction to Turbulence. Springer, 2nd edition.