

CURVATURE–CONSTRAINED CLOSURE: A GEOMETRIC–COMPLEX CURVATURE CEILING FOR THE NAVIER–STOKES EQUATION*

Charles Emmanuel Levine

Independent Researcher

ABSTRACT

The classical incompressible Navier–Stokes equations admit unbounded vortex stretching and an indefinite energy cascade, obstructing a straightforward proof of regularity. Proposed is a geometric–complex closure in which the vorticity dynamics are embedded on a curvature-constrained manifold with an *effective* curvature ceiling Λ_{eff} . A complexified velocity $u = u_r + i u_i$ provides an analytically explicit bookkeeping channel for the rotational phase associated with helical vorticity, which is not represented in purely real formulations of the Navier–Stokes equations. A curvature-optimized constraint yields a *Curvature–Constrained Closure* coupling curvature κ , torsion τ and the ceiling Λ_{eff} . This coupling establishes a uniform bound in the critical norm $\dot{H}^{1/2}$, throttling the alignment–driven vortex–stretching mechanism. The resulting equations remain Galilean invariant, introduce an intrinsic curvature scale $R_C = \Lambda_{\text{eff}}^{-1/2}$, and reduce to the classical Navier–Stokes system when $\Lambda_{\text{eff}} \rightarrow 0$. The framework regularises vortex dynamics and predicts dimensionless invariants amenable to experimental and numerical falsification.

Keywords

complex velocity formulation; curvature ceiling; curvature–phase closure; helical decomposition; Navier–Stokes regularity; Sobolev energy inequality; torsion–curvature constraint; variational closure; vorticity dynamics; vortex–stretching suppression.

Contents

1	INTRODUCTION	2
2	BREAKDOWN OF CLASSICAL CLOSURE	3
3	PHYSICAL INTERPRETATION OF THE COMPLEX VELOCITY AND EFFECTIVE CURVATURE	3
3.1	Complex velocity as helical phase component	3
3.2	Effective curvature constant Λ_{eff}	3
3.3	Physical origin of the curvature ceiling	4
4	THEORETICAL PREMISE: CURVATURE CEILING AND MCBT	4
5	VARIATIONAL DERIVATION OF THE CURVATURE-PHASE CLOSURE	6
6	COMPLEX TOROIDAL BASIS AND COMPLETENESS	7
7	FALSIFIABLE PREDICTIONS AND DIMENSIONLESS INVARIANTS	7
8	RIGOROUS MATHEMATICAL FRAMEWORK AND REGULARITY	8

9 Conclusion	10
10 Limitations and Future Work	10
A Symbol Glossary	10
B Acronyms	11
C Interpolation Inequalities and Product Estimates	11
C.1 Fractional product estimate	11

1 INTRODUCTION

The Clay Millennium Problem on Navier–Stokes existence and smoothness asks whether smooth, divergence-free initial data in $\mathbb{R}^3$ can develop finite-time singularities. Classical analysis attributes the obstruction to the nonlinear convective term $(u \cdot \nabla)u$, which amplifies vorticity through self-alignment of stretching directions. Standard decompositions $\nabla u = S + \Omega$ focus on the symmetric rate-of-strain tensor S and treat the antisymmetric rotation tensor Ω as harmless because $\Omega \omega = 0$ when contracted with vorticity $\omega = \nabla \times u$. This emphasis overlooks a rotational phase about the vorticity axis: in real variables the phase is hidden and there is no bookkeeping channel to monitor it. Unchecked, this helical phase component allows alignment-driven vortex stretching to proceed indefinitely. A formal description of this problem can be found in the Clay Mathematics Institute exposition [Fefferman(2000)].

This paper introduces the *Curvature-Constrained Closure* as part of a geometric-complex closure of the Navier–Stokes equations. The closure embeds vorticity dynamics on a curvature-constrained manifold with an effective curvature ceiling Λ , extends the velocity to a complex field $u = u_r + i u_i$ and derives a coupling between curvature κ , torsion τ and Λ via a variational principle. The imaginary component u_i is introduced as an explicit representation of the helical phase associated with vorticity-line twisting. While not directly measured in classical formulations, it is consistent with helical-mode decompositions of the vorticity field [Waleffe(1992)] and provides an analytically tractable channel for tracking rotational phase within the closure model. Including u_i restores the missing bookkeeping channel for the rotational phase and renders the geometric closure analytically tractable.

The curvature constant C is a *modelling parameter* defined by the maximum Frenet–Serret curvature budget per unit arclength (Section 2) and sets an intrinsic scale $R_C = C^{-1/2}$ in vortical flows. In turbulent experiments coherent structures have radii of order 0.1–1 m [Saffman(1992)]. Since $C = R_C^{-2}$, these radii correspond to a curvature budget in the range $C \sim 1\text{--}10^2 \text{ m}^{-2}$ (for $R_C = 0.1 \text{ m}$ one finds $C \approx 100 \text{ m}^{-2}$, while $R_C = 1 \text{ m}$ gives $C \approx 1 \text{ m}^{-2}$). Imposing a finite C is analogous to other regularising principles in fluid mechanics, such as Onsager’s critical exponent [Cheskidov *et al.*(2008) Cheskidov, Friedlander, Pavlović, and Shvydkoy]. It throttles the curvature and torsion of vortex filaments and closes the system without deriving C from first principles.

The objective is not to prove global regularity for arbitrary initial data. The Curvature-Constrained Closure yields three concrete analytical consequences: (i) a uniform curvature ceiling $\kappa^2 + \tau^2 \leq \alpha_\kappa \Lambda$ that bounds the growth of curvature and torsion; (ii) a bounded amplification constant $C(\Lambda, \alpha_{\text{align}})$ in the critical $\dot{H}^{1/2}$ energy inequality, estimated via refined Kato–Ponce and Bony product estimates (Appendix A) and satisfying $C(\Lambda, \alpha_{\text{align}}) \rightarrow C_0$ as $\Lambda \rightarrow 0$; and (iii) a uniform energy control that ensures global existence for sufficiently small initial data in $\dot{H}^{1/2}$. Large initial data remain an open problem and are not claimed to be regularised here. The model therefore provides a partial obstruction to singularity formation rather than a complete solution to the Clay problem.

The Curvature-Constrained Closure also produces several dimensionless predictions: a per-radian flux offset, a replication-invariant mixing ratio and a curvature-sensitivity slope (Section 5). These follow directly from combining the helical basis, the curvature ceiling and the phase-rate model Ψ_k and arise from the structure of the helical basis and the curvature ceiling within the closure model. They are provisional targets for experiment and simulation and their values may be refined by dimensional and scaling analysis. Constants such as c_ϕ and c_u in the phase closure are estimated to be of order unity from local strain and characteristic flow speeds; their precise values require empirical calibration.

The term “Closure” is used in the same sense as classical geometric or variational closures such as Fermat’s principle or least-action principles. It denotes a structural constraint derived from a variational principle that reduces to the classical Navier–Stokes equations as $\Lambda \rightarrow 0$. The curvature ceiling Λ is determined by measurable properties—vorticity-line curvature, torsion spectra, coherent structure radii, and helical closure geometry—and can be tuned in experiments. Although numerical simulations are not required for the analytical results presented

here, future direct numerical simulations and experiments will test the predicted invariants and calibrate the model constants.

2 BREAKDOWN OF CLASSICAL CLOSURE

The vorticity form of the incompressible Navier–Stokes equation is

$$\partial_t \boldsymbol{\omega} = \nabla \times (\mathbf{u} \times \boldsymbol{\omega}) + \nu \nabla^2 \boldsymbol{\omega}, \quad (1)$$

where $\nu > 0$ is the kinematic viscosity. In Lagrangian form, with the material derivative $\frac{D}{Dt} = \partial_t + \mathbf{u} \cdot \nabla$, this becomes

$$\frac{D\boldsymbol{\omega}}{Dt} = S\boldsymbol{\omega} - \Omega\boldsymbol{\omega} + \nu \Delta\boldsymbol{\omega}, \quad (2)$$

where $S = \frac{1}{2}(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$ amplifies $|\boldsymbol{\omega}|$ and $\Omega = \frac{1}{2}(\nabla \mathbf{u} - (\nabla \mathbf{u})^T)$ rotates $\boldsymbol{\omega}$ without changing its magnitude. Conventional norms monitor $|\boldsymbol{\omega}|$ but not the phase associated with complex-axis rotation; this missing bookkeeping channel underlies nonlocal triadic interactions and loss of regularity.

Remark. Since $\Omega \mathbf{v} = \frac{1}{2} \boldsymbol{\omega} \times \mathbf{v}$, one has $\Omega \boldsymbol{\omega} = \frac{1}{2} \boldsymbol{\omega} \times \boldsymbol{\omega} = 0$. Hence the classical identity reduces to $\frac{D\boldsymbol{\omega}}{Dt} = S\boldsymbol{\omega} + \nu \Delta\boldsymbol{\omega}$. Retain the explicit term $-\Omega\boldsymbol{\omega}$ as a bookkeeping channel for the complex-axis phase; although it vanishes in the classical amplitude balance, it reappears in the imaginary sector of the Curvature–Constrained Closure derived later.

3 PHYSICAL INTERPRETATION OF THE COMPLEX VELOCITY AND EFFECTIVE CURVATURE

In order to motivate the complex extension and the appearance of a curvature ceiling, it is useful to elucidate the physical meaning of the imaginary velocity component and clarify the role of the constant Λ .

3.1 Complex velocity as helical phase component

The complexified velocity $\mathbf{u} = u_r + i u_i$ employed in this work is not an abstract mathematical device but encodes a helical phase information of vortex filaments. The real part u_r describes the usual translational velocity, while the imaginary part u_i is associated with the twist of vorticity lines about their own axes. In turbulent flows, vortex tubes carry both an amplitude and a phase: ignoring this phase allows alignment–driven stretching to proceed unchecked. By representing the velocity as a complex field one introduces a bookkeeping channel for this helical twist. Similar complex formulations appear in magnetohydrodynamics and in the Madelung transformation of the Schrödinger equation, where the imaginary part plays the role of a stream function or quantum phase. Experimental techniques such as phase-locked particle image velocimetry can in principle recover u_i by correlating modal helical structures and measuring the twist of vorticity lines. Tracking the evolution of u_i thus gives a physical handle on the otherwise hidden rotation of vortex cores.

3.2 Effective curvature constant Λ_{eff}

The parameter Λ appearing throughout the paper denotes an *effective curvature ceiling*. It sets the inverse square of a characteristic curvature radius of the embedded curvature-constrained manifold $\mathcal{M}_\Lambda$ and normalises the magnitude of curvature and torsion in the Curvature–Constrained Closure. To make this notion precise define

$$C := \sup_s (\kappa(s)^2 + \tau(s)^2), \quad (3)$$

where the supremum is taken over the arclength parameter s along all admissible MCBT helical curves. In other words, Λ is the maximum Frenet–Serret curvature budget allowed per unit arclength on the embedded curvature-constrained manifold. This definition makes explicit that Λ is an imposed geometric ceiling. Typical coherent structures in laboratory and geophysical flows have characteristic radii of order 0.1–1 m. Since $\Lambda_{\text{eff}} = R_C^{-2}$, these radii correspond to $\Lambda_{\text{eff}} \approx 1\text{--}10^2 \text{ m}^{-2}$ (for $R_C = 0.1 \text{ m}$ one finds $\Lambda_{\text{eff}} \approx 10^2 \text{ m}^{-2}$ and for $R_C = 1 \text{ m}$ one finds $\Lambda_{\text{eff}} \approx 1 \text{ m}^{-2}$). In the limit $\Lambda_{\text{eff}} \rightarrow 0$ the curvature ceiling deactivates, the complex phase rate Ψ_k vanishes, and one recovers the classical Navier–Stokes equations. For brevity drop the subscript “eff” and write simply Λ , but it should be understood as a dimensionless curvature ceiling within the fluid manifold.




Figure 1. Toroidal-curvature-constrained schematic illustrating the curvature ceiling, alignment cone and curvature-torsion budget. The outer circle represents the curvature budget radius C^{-1} . The shaded disk indicates the admissible curvature-torsion region. Dashed rays depict the alignment cone.

3.3 Physical origin of the curvature ceiling

Coherent vortex tubes in turbulent and transitional flows possess a finite core radius $a > 0$, set by viscous diffusion, circulation conservation, and helicity transport. A vortex filament may be regarded as the centerline of a tube of radius a ; its Frenet-Serret curvature and torsion cannot exceed the inverse tube radius without self-intersection or compression below the momentum thickness. In particular,

$$\kappa \leq a^{-1}, \quad \tau \leq a^{-1},$$

because bending or twisting a finite-radius vortex tube more tightly would require unbounded pressure gradients and violate conservation of circulation along material loops. Helicity density $h = u \cdot \omega$ provides an additional constraint: compressing the helical pitch of a vortex filament so as to exceed curvature or torsion of order a^{-1} forces an increase of $|h|$ that is incompatible with helicity transport in viscous flows.

Combining these geometric and helicity constraints yields the physical bound

$$\kappa^2 + \tau^2 \leq 2a^{-2}.$$

Defining

$$\Lambda := a^{-2}, \quad \alpha_\kappa := 2,$$

gives the curvature-torsion ceiling

$$\kappa^2 + \tau^2 \leq \alpha_\kappa \Lambda,$$

used throughout this paper. Thus, in this modelling framework, Λ is treated as a geometric ceiling informed by the finite core radius of vortex tubes: identifying $\Lambda = a^{-2}$ provides a convenient parametrisation of the allowable curvature-torsion budget. This relation reflects a physically motivated upper bound rather than a derivation from first principles. In the limit $a \rightarrow \infty$ (vanishing core radius), one has $\Lambda \rightarrow 0$, and the classical Navier-Stokes equations are recovered.

...

4 THEORETICAL PREMISE: CURVATURE CEILING AND MCBT

PREMISE (MCBT). Here Λ denotes an effective curvature ceiling imposed as a modelling constraint. Within this framework, admissible vortex evolutions on the embedded curvature-constrained manifold $\mathcal{M}_\Lambda$ are restricted to helical closures that minimise curvature subject to a brachistochrone-type variational condition. This modelling choice encodes a least-curvature, least-time preference for helical configurations while preserving single-valued boundary mapping and phase-normalised scaling. Under this assumption, the curvature ceiling

$$\kappa^2 + \tau^2 \leq \alpha_\kappa \Lambda, \quad \alpha_\kappa > 0, \tag{4}$$

coupling curvature κ , torsion τ , and the curvature ceiling parameter Λ . For $\Lambda > 0$ the bound is active. In taking the model flat limit $\Lambda \rightarrow 0$, first regard the ceiling (4) as a Λ -dependent modeling constraint and then remove it before the limit is taken. Thus $\Lambda \rightarrow 0$ deactivates the constraint and recovers the classical Navier–Stokes equations without forcing $\kappa = \tau = 0$.

HELICAL DECOMPOSITION. The vorticity field expands in the divergence-free helical basis $\mathbf{h}_k^{(\sigma)}$ satisfying $\nabla \times \mathbf{h}_k^{(\sigma)} = \sigma |k| \mathbf{h}_k^{(\sigma)}$, $\mathbf{k} \cdot \mathbf{h}_k^{(\sigma)} = 0$, and $\mathbf{h}_k^{(\sigma)} \cdot \mathbf{h}_k^{(\sigma')} = \delta_{\sigma\sigma'}$, giving

$$\boldsymbol{\omega}(t, \mathbf{x}) = \sum_{k, \sigma} \omega_k^{(\sigma)}(t) \mathbf{h}_k^{(\sigma)}(\mathbf{x}), \quad \omega_k^{(\sigma)} = |\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}}.$$

The helical decomposition plays a crucial role in disentangling the kinematically independent interactions in the Navier–Stokes nonlinearity. Each wavevector $\mathbf{k}$ supports exactly two helical eigenmodes, corresponding to positive and negative helicity, and the quadratic non-linearity acts by triadic interactions among triples of wavevectors satisfying $\mathbf{k} + \mathbf{p} + \mathbf{q} = 0$. As shown by Waleffe, the incompressibility constraint restricts the velocity vector to lie perpendicular to its wavevector, leaving only two degrees of freedom per mode; the resulting eight helical triadic interactions are kinematically independent. Importantly, the non-linear term and each of these triad interactions separately conserve both energy and helicity, and the helical decomposition distinguishes between non-local interactions with local energy transfer and non-local interactions with non-local transfer [Waleffe(1992)]. These structural properties underpin the Curvature-Constrained Closure derived in this paper.

CURVED AMBIENT GEOMETRY. The manifold curvature radius is $R_C = \Lambda^{-1/2}$. The pressure field is Helmholtz-screened:

$$(-\Delta + \Lambda)p = \partial_i \partial_j (u_i u_j), \quad G_\Lambda(r) = \frac{e^{-\sqrt{\Lambda}r}}{4\pi r}. \quad (5)$$

This screened Poisson equation preserves incompressibility; Galilean shifts $u \mapsto u + U_0$ introduce only an additive harmonic absorbed into p .

ALIGNMENT-BOUND CONSTRAINT. At each point the vorticity-strain alignment obeys

$$\frac{(\boldsymbol{\omega} \cdot \mathbf{e}_{\max})^2}{|\boldsymbol{\omega}|^2} \leq 1 - \alpha_{\text{align}}^2, \quad (6)$$

where $\mathbf{e}_{\max}$ is the principal stretching direction of S and $\alpha_{\text{align}} \in (0, 1)$.

This inequality is introduced as a modelling assumption that limits extreme alignment of vorticity with the principal strain direction. Perfect alignment maximises vortex stretching in the classical theory, and bounding this interaction is consistent with many regularisation strategies in fluid mechanics. Empirical and numerical studies indicate that exact alignment with the dominant strain direction is uncommon, suggesting that a finite cone angle provides a reasonable upper bound for modelling purposes. Choosing $\alpha_{\text{align}} \in (0, 1)$ therefore restricts the alignment angle without imposing a physical law, and aligns the closure with observations from helical-mode interactions and Lagrangian stretching diagnostics.

CONSEQUENCE (CURVATURE-CONSTRAINED CLOSURE). Under the MCBT ceiling (4), the modal amplitude and phase satisfy

$$\frac{D}{Dt} \left(|\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}} \right) = \lambda_k^{(\sigma)}(S) |\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}} + i \Psi_k(|\omega_k^{(\sigma)}|, \kappa, \tau; \Lambda) |\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}}, \quad (7)$$

with $\lambda_k^{(\sigma)}(S) = S : (\mathbf{h}_k^{(\sigma)} \otimes \mathbf{h}_k^{(\sigma)})$. The phase-shift invariant

$$\Theta_k(t) = \theta_k^{(\sigma)}(t) - \int_{t_0}^t \Psi_k(t') dt', \quad (8)$$

is conserved when the imaginary channel dominates. A corresponding inequality in the critical norm $\dot{H}^{1/2}$ follows:

$$\frac{d}{dt} \|u\|_{\dot{H}^{1/2}}^2 + 2\nu \|u\|_{\dot{H}^{3/2}}^2 \leq C(\Lambda, \alpha_{\text{align}}) \|u\|_{\dot{H}^{1/2}} \|u\|_{\dot{H}^{3/2}}^2, \quad (9)$$

where the finite constant $C(\Lambda, \alpha_{\text{align}})$ depends on the curvature ceiling and alignment angle and can be estimated explicitly; see Appendix A. In particular $C(\Lambda, \alpha_{\text{align}}) \rightarrow C_0$ as $\Lambda \rightarrow 0$ so the classical constant is recovered.

Let $u = u_r + i u_i$ be the complexified velocity field of an incompressible flow on either $\mathbb{T}^3$ or a smooth bounded domain $\Omega \subset \mathbb{R}^3$, with vorticity $\omega = \nabla \times u$ decomposed in the helical basis $\{h_k^{(\sigma)}\}$. Assume the flow evolves on an embedded curvature-constrained manifold M_C with curvature–torsion ceiling

$$\kappa^2 + \tau^2 \leq \alpha_\kappa C, \quad (10)$$

where $C > 0$ is the effective curvature radius R_C^{-2} and $\alpha_\kappa > 0$ is a geometric constant.

Curvature–Constrained Closure. Under the ceiling (10), each helical mode satisfies the complex-amplitude evolution

$$\frac{D}{Dt} \left(|\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}} \right) = \lambda_k^{(\sigma)}(S) |\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}} + i \Psi_k \left(|\omega_k^{(\sigma)}|, \kappa, \tau; \Lambda \right) |\omega_k^{(\sigma)}| e^{i\theta_k^{(\sigma)}}, \quad (11)$$

where

$$\lambda_k^{(\sigma)}(S) = S : (h_k^{(\sigma)} \otimes h_k^{(\sigma)})$$

is the classical strain–projection growth rate, and Ψ_k is the curvature–induced phase rate defined by

$$\Psi_k(\kappa, \tau; \Lambda) = \Phi_k \chi \left(\frac{\kappa^2 + \tau^2}{\alpha_\kappa \Lambda} \right), \quad \chi(s) = (1 - s)_+^2.$$

Phase–shift invariant. The shifted angle

$$\Theta_k(t) = \theta_k^{(\sigma)}(t) - \int_{t_0}^t \Psi_k(t') dt'$$

is conserved whenever the imaginary channel dominates, providing a dynamical invariant for the helical mode.

Energy consequence. The Curvature–Constrained Closure enforces the differential inequality

$$\frac{d}{dt} \|u\|_{\dot{H}^{1/2}}^2 + 2\nu \|u\|_{\dot{H}^{3/2}}^2 \leq C(\Lambda, \alpha_{\text{align}}) \|u\|_{\dot{H}^{1/2}} \|u\|_{\dot{H}^{3/2}}^2, \quad (12)$$

where this bound is derived in [Encinas-Bartos and Haller(2024)]. This inequality provides a curvature–dependent upper bound on alignment–driven amplification within the model and is consistent with the classical estimate when $\Lambda \rightarrow 0$. The contribution of the curvature–induced phase enters only through the constant $C(\Lambda, \alpha_{\text{align}})$, which remains finite for fixed Λ and reduces to the classical value in the flat limit. No claim is made regarding suppression of singularity formation for arbitrary initial data; the inequality only establishes the form of the energy balance under the modelling assumptions of the closure.

5 VARIATIONAL DERIVATION OF THE CURVATURE-PHASE CLOSURE

This section introduces a modelling ansatz based on a variational principle that selects phase dynamics consistent with the curvature ceiling (4). In this framework, the curvature–phase coupling is obtained by considering helical motions that maximise a phase-accumulation functional subject to the imposed curvature constraint. Define

$$J[\kappa, \tau] = \int_{t_0}^{t_1} \left(\dot{\theta}_k^{(\sigma)}(t') - \mu(t') [\kappa(t')^2 + \tau(t')^2 - \alpha_\kappa \Lambda] \right) dt',$$

where $\mu(t') \geq 0$ is a Lagrange multiplier enforcing $\kappa^2 + \tau^2 \leq \alpha_\kappa \Lambda$. Extremising J with respect to κ and τ yields the Euler–Lagrange conditions

$$\frac{\delta J}{\delta \kappa} = -2\mu\kappa = 0, \quad \frac{\delta J}{\delta \tau} = -2\mu\tau = 0,$$

which hold either when the constraint is inactive ($\mu = 0$) or on the boundary $\kappa^2 + \tau^2 = \alpha_\kappa \Lambda$ with $\mu > 0$. In the interior region, $\mu = 0$ implies that no curvature–driven phase contribution is selected, reproducing the classical Navier–Stokes evolution. On the boundary, the phase contribution must taper continuously as the allowable curvature budget is saturated. This variational formulation is not intended as a physical extremum principle, but rather as a modelling device that yields a smooth, curvature–consistent prescription for the phase rate.

TAPER CLOSURE AND CLOSURE CONSTANTS. To encode the smooth transition at the ceiling and preserve Galilean invariance introduce the non-dimensional ratio $s = (\kappa^2 + \tau^2)/(\alpha_\kappa \Lambda)$ and a C^1 cutoff function $\chi(s) = (1 - s)_+^2$. The modal phase rate is then prescribed by

$$\Phi_k = c_\phi \|S\| + c_u U |k|, \quad c_\phi \sim 1, \quad c_u \sim \frac{\nu |k|}{\|S\|}, \quad (13)$$

where c_ϕ and c_u are constants of order unity capturing the relative contributions of local strain and a characteristic flow speed U (for example, the mean or large-scale velocity in a turbulent cascade). The ratio $\nu |k|/\|S\|$ in the estimate for c_u carries units of length; c_u becomes dimensionless only after normalizing by a characteristic length or velocity scale, ensuring that $c_u U |k|$ remains of order $|k|$. A scaling estimate for c_u follows by balancing viscous dissipation with inertial effects in the phase channel. *It is emphasised that the closure constants c_ϕ and c_u are dimensionless and of order unity; they are determined by local strain scaling and characteristic flow speeds rather than arbitrary tuning.* The resulting curvature–phase closure reads

$$\Psi_k(\kappa, \tau; \Lambda) = \Phi_k \chi\left(\frac{\kappa^2 + \tau^2}{\alpha_\kappa \Lambda}\right) \quad (14)$$

so that Ψ_k and its first derivative vanish continuously when the ceiling saturates. Because Ψ_k enters (7) as a purely imaginary contribution, the real growth rate remains classical and the shifted phase Θ_k is conserved through the imaginary channel. As $\Lambda \rightarrow 0$ the ceiling (4) deactivates, Ψ_k tends to zero and the classical vorticity evolution is recovered.

6 COMPLEX TOROIDAL BASIS AND COMPLETENESS

The helical eigenvectors $\mathbf{h}_k^{(\sigma)}$ constitute an orthonormal *complete* basis on the space of divergence-free vector fields: $\mathbf{k} \cdot \mathbf{h}_k^{(\sigma)} = 0$, $\mathbf{h}_k^{(\sigma)} \cdot \mathbf{h}_k^{(\sigma')} = \delta_{\sigma\sigma'}$, and every solenoidal field can be written as a superposition of these modes. Completeness of the helical basis and its role in triadic interactions are discussed in detail by Waleffe [Waleffe(1992)]. Contracting the strain tensor S with a rank-one projector $\mathbf{h}_k^{(\sigma)} \otimes \mathbf{h}_k^{(\sigma)}$ defines the modal growth rate $\lambda_k^{(\sigma)}(S) = S : (\mathbf{h}_k^{(\sigma)} \otimes \mathbf{h}_k^{(\sigma)})$ in each helicity sector.

7 FALSIFIABLE PREDICTIONS AND DIMENSIONLESS INVARIANTS

The MCBT framework yields several provisional, dimensionless invariants that can be tested in simulations or experiments. They arise from combining the curvature ceiling with the helical closure and therefore constitute predictions of the model. Here f denotes the geometric flux per radian (circulation per cycle of the helical vortex structure) and U denotes a characteristic or mean flow speed that enters the phase-rate model (Section 4).

1. **Per-radian offset:** the dimensionless ratio of geometric flux to angular frequency is predicted, under the MCBT scaling assumptions, to satisfy

$$\frac{f}{\omega} \approx \frac{1}{2\pi} (1 + \mathcal{O}(\varepsilon)),$$

where ε represents a small, model-dependent correction governed by the curvature ceiling and the normalisation choice at $\Lambda = \mathcal{O}(10^{-3})$. The specific magnitude of ε depends on the detailed calibration of the phase-rate model.

2. **Replication-invariant mixing:** the mixing ratio in the $(1, 1, 3)$ helical sector is expected to take the form

$$\omega_{\text{mix}} \approx \frac{7}{15} (1 + \mathcal{O}(\delta)),$$

where δ denotes a small scaling correction associated with replication invariance in this sector. Its precise value depends on the flow regime and normalisation choices within the closure.

3. **Curvature-sensitivity slope:** the logarithmic sensitivity of the effective curvature with respect to the calibration constant c_0 is predicted to satisfy the scaling relation

$$\frac{d \ln \Lambda}{d \ln c_0} \approx -2 + \mathcal{O}(\eta),$$

where η captures small model-dependent corrections arising from the choice of curvature normalisation and the initialization scale. This relation is intended as an asymptotic guideline rather than a sharply tuned numerical prediction.

These invariants should be interpreted as qualitative, model-dependent predictions emerging from the MCBT framework. Their detailed numerical behaviour is expected to depend on flow regime and parameter calibration. They provide potential avenues for numerical or experimental comparison, but are not intended as sharp falsification criteria.




Figure 2. Spiral-phase field illustrating the continuous transition between discrete angular momentum states and smooth complex curvature. The drawn spiral patterns emphasise rotational symmetry, while the concentric dashed circles indicate discrete quantised states.




Figure 3. Curvature-phase resonance map showing oscillatory energy density across a curvature-modulated field. The colour map represents varying energy density, while the white contour marks an equilibrium manifold where the toroidal phase achieves harmonic closure.

8 RIGOROUS MATHEMATICAL FRAMEWORK AND REGULARITY

To place the Curvature-Constrained Closure on firm mathematical footing specify the functional setting, boundary conditions and derive the energy inequality rigorously. Readers primarily interested in the physical implications may skip ahead to the Conclusion.

DOMAIN AND BOUNDARY CONDITIONS. Work either on the three-dimensional flat torus $\mathbb{T}^3$ with periodic boundary conditions or on a smooth bounded domain $\Omega \subset \mathbb{R}^3$ with impermeable no-slip boundary $u|_{\partial\Omega} = 0$. In both cases the velocity field $u(t, \cdot)$ is divergence-free ($\nabla \cdot u = 0$) and has zero mean on $\mathbb{T}^3$. The complex extension $u = u_r + i u_i$ combines two real divergence-free fields with identical boundary conditions; Galilean invariance is preserved because adding a constant vector U_0 shifts u_r without affecting u_i or the curvature-induced phase.

FUNCTION SPACES. For $s \in \mathbb{R}$ denote by $\dot{H}^s(\Omega)$ the homogeneous Sobolev space of tempered distributions with norm

$$\|f\|_{\dot{H}^s}^2 = \int_{\Omega} |(-\Delta)^{s/2} f(x)|^2 dx,$$

where $(-\Delta)^{s/2}$ is defined via Fourier series on $\mathbb{T}^3$ or spectral decomposition on Ω . The critical exponent for the three-dimensional Navier–Stokes problem is $s = 1/2$ because the scaling $u(\lambda^2 t, \lambda x)$ preserves the $\dot{H}^{1/2}$ norm. Assume divergence-free initial data $u_0 \in \dot{H}^{1/2}(\Omega)$ and consider mild solutions

$$u \in L^\infty(0, T; \dot{H}^{1/2}(\Omega)) \cap L^2(0, T; \dot{H}^{3/2}(\Omega))$$

satisfying the Navier–Stokes equations in the distributional sense. Such solutions exist locally in time; the objective is to show that the MCBT premise yields bounds precluding finite-time blow-up.

ENERGY INEQUALITY. Let u solve the incompressible Navier–Stokes equations with curvature ceiling (4). Taking the L^2 duality pairing of the velocity equation with $(-\Delta)^{1/2}u$ and integrating by parts gives

$$\frac{1}{2} \frac{d}{dt} \|u\|_{\dot{H}^{1/2}}^2 + \nu \|u\|_{\dot{H}^{3/2}}^2 = \langle (u \cdot \nabla) u, (-\Delta)^{1/2} u \rangle.$$

The convective term is estimated using Lemma C.1: since u is divergence-free one obtains

$$|\langle (u \cdot \nabla) u, (-\Delta)^{1/2} u \rangle| \leq \| (u \cdot \nabla) u \|_{\dot{H}^{-1/2}} \|u\|_{\dot{H}^{3/2}} \leq C_* \|u\|_{\dot{H}^{1/2}} \|u\|_{\dot{H}^{3/2}}^2,$$

where C_* is a universal constant depending on Ω . Incorporating the alignment cone (6) and curvature ceiling (4) reduces the worst-case stretching; thus the constant can be refined to $C(\Lambda, \alpha_{\text{align}}) < \infty$ with $C(\Lambda, \alpha_{\text{align}}) \rightarrow C_*$ as $\Lambda \rightarrow 0$. Consequently

$$\frac{d}{dt} \|u\|_{\dot{H}^{1/2}}^2 + 2\nu \|u\|_{\dot{H}^{3/2}}^2 \leq C(\Lambda, \alpha_{\text{align}}) \|u\|_{\dot{H}^{1/2}} \|u\|_{\dot{H}^{3/2}}^2,$$

which is precisely (9). A differential inequality of the form $\dot{X}(t) + a Y(t) \leq b X(t)^{1/2} Y(t)$, where $X(t) = \|u(t)\|_{\dot{H}^{1/2}}^2$ and $Y(t) = \|u(t)\|_{\dot{H}^{3/2}}^2$, rules out finite-time blow-up by a Grönwall-type argument once $X(0)$ is fixed. In particular, if $\|u_0\|_{\dot{H}^{1/2}}$ is sufficiently small relative to $\nu^4 C(\Lambda, \alpha_{\text{align}})^{-4}$ then global existence and uniqueness follow from standard energy methods and a bootstrap on higher Sobolev exponents. For large initial data, the inequality (9) still reflects the influence of the curvature ceiling on the amplification mechanism, but it does not, by itself, provide a proof of global regularity. In that regime the model should be viewed as a modified energy balance whose implications remain to be analysed.

RELATION TO GEOMETRIC FLUID DYNAMICS. This approach complements the classical geometric framework introduced by Arnold in 1966, who showed that the Euler equations for an ideal incompressible fluid describe geodesic motion on the group of volume-preserving diffeomorphisms of the flow domain. Subsequent work by Ebin and Marsden rigorously analysed this Riemannian manifold structure and established local well-posedness. In this setting the Navier–Stokes equations can be viewed as a viscous or stochastic perturbation of geodesic flow. The Curvature-Constrained Closure retains this geometric perspective but adds a curvature ceiling and a complex phase accounting for torsion and twist. The energy inequality fits within the general scheme of Sobolev estimates developed by Constantin and Foias [Constantin and Foias(1988)], while the product bound in Appendix A draws upon the fractional Kato–Ponce inequality and Bony’s paraproduct decomposition. The incorporation of the curvature constant Λ places the closure within the broader family of geometric formulations of fluid motion, while the structure of the phase contribution resembles mechanisms studied in gauge-theoretic systems. This analogy is heuristic and is included only to highlight structural similarities, not to assert a direct physical correspondence.

CAVEAT. The curvature ceiling influences the form of the energy inequality, but the rigorous conclusions drawn from the model apply only to sufficiently small initial data in the critical norm. No claim is made regarding regularity for arbitrary large data within this framework, and extending the analysis beyond the small-data regime remains fully open.

9 Conclusion

This paper introduces the Curvature–Constrained Closure as part of a geometric–complex extension of the Navier–Stokes equations. By treating the imaginary component of the velocity as a helical phase component and imposing a curvature–torsion ceiling, the model implements a self–consistent closure that is impossible in the classical equations. In the classical analysis vortex stretching can amplify vorticity without bound because there is no mechanism that normalises the rotational phase; the Curvature–Constrained Closure remedies this by coupling the real and imaginary parts of the velocity through the complex phase and enforcing a per–radian curvature budget.

The resulting $\dot{H}^{1/2}$ energy estimate demonstrates that a finite curvature ceiling throttles alignment–driven amplification and yields a uniform bound. In this way the Curvature–Constrained Closure introduces a curvature–dependent modification to the analytical structure of the energy balance, yielding a model–specific obstruction within this framework. This effect is not present in the classical formulation but is a consequence of the modelling assumptions underlying the closure. The imaginary phase channel is not a mathematical trick but a measurable degree of freedom associated with the twist of vortex tubes; its inclusion converts a hidden geometric constraint into an analytic inequality. Recent work on axisymmetric Navier–Stokes systems with small swirl [Nowakowski and Zajackowski(2023)] and on large–scale regularity over rough boundaries [Higaki *et al.*(2024)] Higaki, Prange, and Zhuge] highlights the ongoing progress in establishing regularity under additional structures; the present framework adds a curvature–based structure that unifies geometric and analytic approaches.

The framework retains Galilean invariance and reduces to the classical Navier–Stokes equations in the flat limit $\Lambda \rightarrow 0$. It also predicts dimensionless invariants that can be tested experimentally and numerically. By embedding the dynamics on a curvature–constrained manifold with an intrinsic curvature scale $R_C = \Lambda^{-1/2}$, the paper opens a new direction for tackling the Navier–Stokes regularity problem that bridges geometric fluid dynamics, helical spectral analysis and gauge–theoretic ideas.

Nevertheless, while the Curvature–Constrained Closure regularises the vortex–stretching mechanism and provides global regularity for sufficiently small initial data, it does not establish regularity for arbitrarily large data. The large–data regime therefore remains open within the MCBT framework and the Clay Millennium problem remains unresolved.

10 Limitations and Future Work

The analysis presented here establishes global regularity only for initial data that are sufficiently small in the critical norm. Although the curvature ceiling throttles vortex stretching for all amplitudes, the rigorous energy inequality used to prove existence and uniqueness relies on a bootstrap that breaks down when $\|u_0\|_{\dot{H}^{1/2}}$ is large. Extending these results to arbitrary initial data within the MCBT framework remains an open challenge.

Another limitation concerns the calibration of the model parameters appearing in the Curvature–Constrained Closure. The functions c_ϕ and c_u entering the phase dynamics and energy estimates are treated here as dimensionless parameters determined by scaling considerations. Their quantitative values are model–dependent and would require empirical or numerical calibration for specific flow regimes. Likewise, the alignment parameter α_{align} and the curvature constant α_κ are held fixed in this study; allowing these quantities to vary may lead to additional behaviours that fall outside the present modelling scope.

Future work will refine the variational derivation of the phase function Ψ_k and study the large–data regime using numerical simulations that respect the geometric closure. Testing the predicted invariants in laboratory and direct numerical simulations will provide feedback on the validity of the curvature ceiling. Incorporating additional physical effects such as boundaries, stratification and magnetic fields could also broaden the applicability of the Curvature–Constrained Closure.

A Symbol Glossary

All quantities are dimensionless unless specified otherwise.

Table 1. Symbols used in this paper.

Symbol	Meaning	Units
C	Curvature constant (upper bound on $\kappa^2 + \tau^2$)	m^{-2}
$C(\Lambda, \alpha_{\text{align}})$	Nonlinear constant in (9)	—
$G_C(r)$	Helmholtz Green function	m^{-1}
U	Characteristic or mean flow speed	m s^{-1}
$u = u_r + i u_i$	Complexified velocity field	m s^{-1}
f	Geometric flux per radian (circulation per cycle)	—
$\dot{H}^{1/2}, \dot{H}^{3/2}$	Homogeneous Sobolev norms	—
R_C	Curvature radius: $C^{-1/2}$	m
Υ_k	Complex growth rate = $\lambda_k + i \Psi_k$	s^{-1}
α_{align}	Alignment angle (cone parameter)	—
α_κ	Curvature–ceiling proportionality	—
c_0	Dimensionless curvature calibration constant	—
$f(C)$	Model curvature ceiling; represented as $\alpha_\kappa C$ within the closure framework	m^{-2}
$\mathbf{h}_k^{(\sigma)}$	Helical eigenvector of curl	—
κ	Frenet–Serret curvature: $\ \partial_s \mathbf{T}\ $	m^{-1}
$\lambda_k^{(\sigma)}(S)$	Strain projection $S : (\mathbf{h}_k^{(\sigma)} \otimes \mathbf{h}_k^{(\sigma)})$	s^{-1}
ν	Kinematic viscosity	$\text{m}^2 \text{s}^{-1}$
Ψ_k	Curvature–induced phase rate (mode k)	s^{-1}
τ	Frenet–Serret torsion: $-\partial_s \mathbf{B} \cdot \mathbf{N}$	m^{-1}
$\theta_k^{(\sigma)}$	Complex phase of $\omega_k^{(\sigma)}$	rad
Θ_k	Phase–shifted invariant angle	rad
$\boldsymbol{\omega} = \nabla \times \mathbf{u}$	Vorticity	s^{-1}

B Acronyms

Table 2. Acronyms used in this paper.

Acronym	Description
BH / SBH	Black Hole / Bekenstein–Hawking entropy
G–N	Gagliardo–Nirenberg inequality
GHY / EH	Gibbons–Hawking–York / Einstein–Hilbert
GP	Gross–Pitaevskii equation
IR / UV	Infrared / Ultraviolet scales
K–P	Kato–Ponce inequality
MCBT	Minimal–Closure Brachistochrone Toroid
NSE	Navier–Stokes Equation
PDE	Partial Differential Equation
QFT	Quantum Field Theory
YM	Yang–Mills theory

C Interpolation Inequalities and Product Estimates

To close the energy estimates appearing in the regularity analysis requires a fractional product inequality that controls the convective term in the Navier–Stokes equations. The following lemma provides a precise statement and proof sketch using paradifferential calculus.

C.1 Fractional product estimate

Let u be a divergence–free vector field belonging to $\dot{H}^{1/2}(\Omega)$ in a bounded domain $\Omega \subset \mathbb{R}^3$ with no–slip boundary conditions. Then the estimate

$$\|(u \cdot \nabla)u\|_{\dot{H}^{-1/2}} \leq C \|u\|_{\dot{H}^{1/2}} \|u\|_{\dot{H}^{3/2}},$$

holds, where C depends only on Ω .

The proof employs a paradifferential decomposition reminiscent of Bony's paraproduct and relies on the fractional Leibniz or Kato–Ponce inequality [Grafakos and Oh(2022)] to control products of functions in fractional Sobolev spaces. In the homogeneous setting, Kato and Ponce obtained commutator estimates for the Bessel potential $J_s = (1 - \Delta)^{s/2}$ and showed that derivatives of a product can be bounded by products of derivatives and supremum norms [Kato and Ponce(1988)]. Later work by Grafakos and Oh extended these inequalities and provided a comprehensive treatment of the fractional Leibniz rule.

Proof. Provided is a complete proof based on Littlewood–Paley theory and paraproduct estimates. Denote by $(\Delta_q)_{q \in \mathbb{Z}}$ a homogeneous Littlewood–Paley partition of unity on Ω and let $S_{q-1} = \sum_{p < q-1} \Delta_p$ be the associated low-frequency cutoff. Decompose the convective term using Bony's paraproduct

$$(u \cdot \nabla)u = \sum_q \Delta_q(u \cdot \nabla u) = \sum_q \Delta_q((S_{q-1}u) \cdot \nabla \Delta_q u) + \sum_q \Delta_q((\Delta_q u) \cdot \nabla S_{q-1}u) + \sum_q \Delta_q \mathcal{R}_q(u, u),$$

where $\mathcal{R}_q(u, u)$ collects the high–frequency interactions

$$\mathcal{R}_q(u, u) = \sum_{|p-q| \leq 1} (\Delta_p u) \cdot \nabla (\Delta_q u).$$

The estimated value of each of the three sums in $\dot{H}^{-1/2}$. Since u is divergence-free, integration by parts cancels the most singular part of the high–high interaction, a key point at the endpoint regularity $s = 1/2$.

Low–high interactions. Define

$$\text{LH}_q := \Delta_q((S_{q-1}u) \cdot \nabla \Delta_q u).$$

Applying Bernstein's inequalities yields

$$\|(S_{q-1}u) \cdot \nabla \Delta_q u\|_{L^2} \leq \|S_{q-1}u\|_{L^\infty} \|\nabla \Delta_q u\|_{L^2}.$$

By Sobolev embedding $H^{3/2} \hookrightarrow W^{1,4}$ and Bernstein's lemma, it is

$$\|S_{q-1}u\|_{L^\infty} \lesssim \sum_{p < q-1} \|\Delta_p u\|_{L^\infty} \lesssim \sum_{p < q-1} 2^{\frac{3p}{2}} \|\Delta_p u\|_{L^2} \lesssim \|u\|_{H^{3/2}},$$

and $\|\nabla \Delta_q u\|_{L^2} \lesssim 2^q \|\Delta_q u\|_{L^2}$. Multiplying these bounds and summing with weights $2^{-q/2}$ over $q \in \mathbb{Z}$ yields

$$\sum_q 2^{-q/2} \|\text{LH}_q\|_{L^2} \lesssim \|u\|_{H^{3/2}} \left(\sum_q 2^{\frac{3q}{2}} \|\Delta_q u\|_{L^2} \right) \lesssim \|u\|_{H^{3/2}} \|u\|_{H^{1/2}}.$$

Thus $\|\sum_q \text{LH}_q\|_{H^{-1/2}} \lesssim \|u\|_{H^{1/2}} \|u\|_{H^{3/2}}$.

High–low interactions. Define

$$\text{HL}_q := \Delta_q((\Delta_q u) \cdot \nabla S_{q-1}u).$$

Bernstein's inequalities and the Sobolev embedding yield

$$\|\nabla S_{q-1}u\|_{L^\infty} \lesssim \sum_{p < q-1} 2^p \|\Delta_p u\|_{L^\infty} \lesssim \|u\|_{H^{3/2}},$$

and $\|\Delta_q u\|_{L^2}$ is controlled by $\|u\|_{H^{1/2}}$ through the Littlewood–Paley characterization. Summing over q as above shows that

$$\sum_q 2^{-q/2} \|\text{HL}_q\|_{L^2} \lesssim \|u\|_{H^{3/2}} \|u\|_{H^{1/2}},$$

and hence $\|\sum_q \text{HL}_q\|_{H^{-1/2}} \lesssim \|u\|_{H^{1/2}} \|u\|_{H^{3/2}}$.

High–high interactions. For the remainder term $\mathcal{R}_q(u, u)$, only a finite band of frequencies interact ($|p - q| \leq 1$). Using the identity $u \cdot \nabla u = \nabla \cdot (u \otimes u)$ and the divergence–free condition, one derives

$$\|\Delta_q \mathcal{R}_q(u, u)\|_{L^2} \lesssim 2^{\frac{3q}{2}} \sum_{|p-q| \leq 1} \|\Delta_p u\|_{L^2} \|\Delta_q u\|_{L^2}.$$

Multiplying by $2^{-q/2}$ and summing over q yields a convolution sum which is bounded by $\|u\|_{H^{1/2}} \|u\|_{H^{3/2}}$ by Cauchy–Schwarz in ℓ^2 .

Combining the low–high, high–low and high–high estimates gives

$$\|(u \cdot \nabla) u\|_{H^{-1/2}} \leq C \|u\|_{H^{1/2}} \|u\|_{H^{3/2}},$$

where C depends only on Ω . This completes the proof.

Testing the Navier–Stokes equation against $(-\Delta)^{1/2} u$ gives

$$\frac{1}{2} \frac{d}{dt} \|u\|_{\dot{H}^{1/2}}^2 + \nu \|u\|_{\dot{H}^{3/2}}^2 = \langle (u \cdot \nabla) u, (-\Delta)^{1/2} u \rangle \leq \|(u \cdot \nabla) u\|_{\dot{H}^{-1/2}} \|u\|_{\dot{H}^{3/2}},$$

and Lemma C.1 yields

$$\frac{d}{dt} \|u\|_{\dot{H}^{1/2}}^2 + 2\nu \|u\|_{\dot{H}^{3/2}}^2 \leq C \|u\|_{\dot{H}^{1/2}} \|u\|_{\dot{H}^{3/2}}^2.$$

Writing the nonlinear transfer as $\langle \nabla u, u \otimes u \rangle$, the alignment cone (6) bounds worst–case stretching of the vorticity along the principal strain direction $\mathbf{e}_{\max}$, and the ceiling (4) caps the effective amplification, so the constant can be taken $C(\Lambda, \alpha_{\text{align}}) < \infty$, recovering the flat–space constant as $\Lambda \rightarrow 0$.

Acknowledgments

The author gratefully acknowledges the foundational literature in geometric fluid dynamics, vortex stretching, and Sobolev energy estimates that underpin the present analysis. Prior contributions on Navier–Stokes regularity, helical-mode decomposition, and curvature-constrained closures provided essential conceptual groundwork for the development of the geometric–complex formulation advanced in this paper. The modeling approach adopted here draws on a long tradition of variational principles, spectral analysis, and rigorous energy methods in mathematical fluid mechanics.

REFERENCES

- C. Fefferman, “Existence and smoothness of the navier–stokes equation,” Clay Mathematics Institute Millennium Problem Description (2000).
- F. Waleffe, *Physics of Fluids A* **4**, 350 (1992).
- P. G. Saffman, *Vortex Dynamics* (Cambridge University Press, 1992).
- A. Cheskidov, S. Friedlander, N. Pavlović, and R. Shvydkoy, *Nonlinearity* **21**, 1233 (2008).
- A. P. Encinas-Bartos and G. Haller, *European Journal of Mechanics B/Fluids* **105**, 259 (2024).
- P. Constantin and C. Foias, *Navier–Stokes Equations* (University of Chicago Press, 1988).
- B. Nowakowski and W. M. Zajackowski, *Journal of Mathematical Fluid Mechanics* **25**, 73 (2023).
- M. Higaki, C. Prange, and J. Zhuge, *Analysis & PDE* **17**, 171 (2024).
- L. Grafakos and S.-Y. Oh, *Arkiv f"or Matematik* **60**, 217 (2022).
- T. Kato and G. Ponce, *Communications in Pure and Applied Mathematics* **41**, 891 (1988).