Structural Classification of the Exterior-Exponent Collapse Mechanism

Lauri Korpela
MiulusTek Oy, Finland · lauri.korpela@miulustek.com

Paper I, September 2026. The accompanying claim ledger gives the current status of every claim, including the corrections made to Attachment A. Complete constructions belong to Paper II [1].

Abstract

We study a class of mean–interaction collapse architectures organized by the exterior circulation exponent α. Under an explicitly source-free leading-core premise, a strained-core width law, saturated self-similar supply and axial balance, circulation growth and vanishing core energy reduce to the structural window $0 < \alpha < 1/3$. The cut-dependent collapse Reynolds number cancels from this criterion. For the specified interaction-wave family, covariance realizability is governed by a circular cone aligned with the radial shear. Its exterior parameter satisfies $v_s - 2 = \alpha$, connecting the kinematic and realizability descriptions. We express the demanded stress through radial primitives and height-dependent moment identities, retain the full profile state, and state the moment-repair interface with its transcription and uniformity hypotheses. An exact finite perturbation constructs an active midplane layer; a separate linear lift gives a stress-free global midplane preparation under a quantified non-resonance condition. Height continuation is a coupled boundary-value problem. We identify its scalar fold, quasi-elliptic $(2, 6)$ axis core and common modal structure, then prove that the full frozen-window resolvent on its original domain is bounded below by a quantity tending to infinity at large real height frequency. This excludes normalized high-frequency quasimodes without asserting a power-law rate. OpenAI's external construction (OA) maps to $\alpha = 2h < 1/50$ and supplies a realization only conditional on its correctness. Finite-frequency non-resonance, range closure, variable profiles, height continuation and complete Navier–Stokes verification are beyond the results proved here.

Keywords: Navier–Stokes equations; finite-time singularity; axisymmetric vortex; Reynolds stress; covariance realizability; degenerate boundary-value problems; resolvent estimates.

MSC 2020: 35Q30, 76D05, 35B44, 76D17.

1. Contribution and premises

We identify a one-parameter structural class of mean–interaction collapse mechanisms indexed by the exterior circulation exponent α. Within the source-free-core, strained-width and self-similar premises stated below, circulation growth and vanishing core energy reduce to $0 < \alpha < 1/3$. For the specified covariance family, the interaction cone is circular and shear-aligned, and in the exterior satisfies $v_s - 2 = \alpha$. We reduce profile realizability to explicit stress and moment compatibility, construct the global midplane under a quantified non-resonance condition, classify the axis fold and its normal forms, and prove full high-frequency compactness for the frozen fold window on its original domain. Conditional on its correctness, the OA construction [2] realizes the subclass $\alpha = 2h < 1/50$. Its role is a consistency check and conditional existence witness; it does not establish the independently stated reduced-problem theorems or a complete realization at other exponents.

The physical equations, with pressure divided by density, are

$$\partial_t u + (u \cdot \nabla)u + \nabla p = \nu_{\text{phys}} \Delta u + f, \qquad \nabla \cdot u = 0.$$

Write $s = T - t$. The architecture is restricted as follows.

Premise	Content and status
A1	$u = \bar{u} + u'$, with $\bar{u}$ axisymmetric, $\langle u' \rangle_\theta = 0$, both fields divergence-free, and $\tau = \langle u' \otimes u' \rangle_\theta$. **[CLASS]**
A1b	The leading mean swirl source vanishes in the core: $\mathcal{T}[\tau] = 0$, with $\bar{f}^\theta = 0$ there at the same order. The constructions considered below impose zero demanded interaction stress in that region, a stronger reduced-profile requirement. **[CLASS]**
A2	A slender shrinking core has radial width ρ, axial length ℓ, $\ell \gg \rho$, and circulation amplitude $\Lambda = \|\bar{\Gamma}\|_{L^\infty(C_s)}$, $\bar{\Gamma} = r\bar{u}^\theta$. **[ASSUMP]**
A3	The ambient circulation is $\Gamma_{\text{ext}}(r) \sim C_\infty r^{-\alpha}$ at the matching scale. **[ASSUMP]**
A4	Supply is saturated: $\Lambda \asymp \Gamma_{\text{ext}}(\rho)$. This is additional to the maximum-principle upper bound. **[ASSUMP]**
A5	$\|u^z\|/\ell \asymp s^{-1}$ and $\|u^z\| \asymp \|u^\theta\|$ in the leading core. **[BAL]**
SS	The rate identities use an exact self-similar leading model: $\rho^2 = cs$, $\Lambda = C_\Lambda \rho^{-\alpha}$, $\ell = C_\ell s^{(1-\alpha)/2}$, with fixed positive constants, and the width law of §3. For asymptotic profiles, differentiated asymptotics must be proved before asserting instantaneous rate limits. **[CLASS / ASSUMP]**

The supply theorem is exact for the source-free reduced core equation. If A1b holds only at leading order in (NS), applying it to the full field requires bounds on the omitted source. This distinction is part of the construction problem.

There are three distinct levels: exact identities of (NS); theorems for a specified reduced profile system; and complete fields satisfying all of (NS). The profile system in §4 is **[TRANSCRIPTION]** from [2]. Algebraic consequences of the displayed system are exact for that system. Identifying it with the intended OA reduction remains conditional on a correct transcription. Covariance membership is Level A; producing an actual divergence-free interaction field with the required errors and regularity is Level B.

We write a_{grow} for the temporal circulation exponent, a_s for the first shear component, η_{rate} for the normalized axial contraction rate, and η for the profile's height coordinate. They are different quantities.

Status labels. Bold labels state the strength of each claim. **[EXACT]**: proved for the displayed equations, hypotheses and domain. **[CLASS]**: a restriction that defines the class. **[ASSUMP]**: an assumption of the reduction. **[BAL]**: a scale balance. **[COND]**: requires the named assumptions. **[COND-OA]**: conditional on the correctness of the OA construction [2] and on the dictionary of §8. **[TRANSCRIPTION]**: exact for the reduced equations as transcribed from [2], whose identification with that source is conditional. **[MAP]**: a dictionary entry. **[NUM]**: a computation at its stated resolution, not a continuum theorem. The claim ledger gives the complete list. Results cited as "Attachment A, Theorem 9.x" use the numbering of Attachment A; Propositions 9.51–9.57 are proved in Attachment C, and Proposition 9.58, Lemma 9.59 and Theorem 9.60 in Attachment D.

2. Exact angular momentum, mean stress and supply

Proposition 2.1 [EXACT]. With the standard cylindrical basis, $m = ru^\theta$ satisfies

$$(\partial_t + u \cdot \nabla)m = -\partial_\theta p + \nu_{\text{phys}}\left(\Delta m - \frac{2}{r}\partial_r m + \frac{2}{r}\partial_\theta u^r\right) + rf^\theta.$$

Here $\Delta = \partial_r^2 + r^{-1}\partial_r + r^{-2}\partial_\theta^2 + \partial_z^2$ is the scalar Laplacian. Indeed, $r(\Delta - r^{-2})u^\theta = (\Delta - 2r^{-1}\partial_r)m$; multiplying the azimuthal momentum equation by r cancels its geometric transport term against $(u \cdot \nabla r)u^\theta$. The vector-Laplacian cross term has sign $+2\nu_{\text{phys}} r^{-1}\partial_\theta u^r$, as in the Caltech cylindrical-coordinate equations [3].

Thus pressure torque and viscous angular torque both obstruct the source-free swirl maximum principle. Axisymmetric pressure alone removes only the first. Axisymmetric velocity, compatible axisymmetric pressure and zero azimuthal forcing remove both; for smooth axisymmetric solutions the resulting bound $\|ru^\theta(t)\|_{L^\infty} \le \|ru^\theta(0)\|_{L^\infty}$ is classical [4]. The regular axis condition $m = O(r^2)$ and appropriate boundary or decay hypotheses are retained. This corrects the printed full-field identity in Attachment A, §3; its mean equation below is unchanged.

Proposition 2.2 [EXACT]. Angular averaging gives

$$\partial_t \bar{\Gamma} + \bar{u} \cdot \nabla \bar{\Gamma} = \nu_{\text{phys}}(\partial_r^2 - r^{-1}\partial_r + \partial_z^2)\bar{\Gamma} - \mathcal{T}[\tau] + r\bar{f}^\theta, \quad \mathcal{T}[\tau] = \frac{1}{r}\partial_r(r^2\tau_{r\theta}) + r\partial_z\tau_{z\theta}.$$

This follows by incompressibility and the cylindrical divergence of τ. Stress redistributes angular momentum through a flux divergence. Both radial–azimuthal and radial–axial stress components remain in the profile matching problem; the axial momentum equation must also be satisfied.

Theorem 2.3 (supply) [EXACT | A1b and the stated core boundary hypotheses]. A regular solution of the source-free core equation obeys

$$\| \bar{\Gamma} \|_{L^\infty(C)} \le \sup_{\partial_{\text{par}} C} |\bar{\Gamma}|.$$

Apply the scalar parabolic maximum principle [5, Ch. 7] to both signs, with regularity at the axis. The parabolic boundary includes the initial slice and every spatial part through which supply can enter. Comparing its supremum with $\Gamma_{\text{ext}}(\rho)$ requires the matching/supply hypotheses; a radial exterior formula alone does not bound arbitrary axial or initial data. A4 then supplies the matching lower bound.

For $\alpha > 0$ and saturated supply, $\Lambda \to \infty$. A globally axisymmetric, unforced velocity satisfying the maximum-principle hypotheses cannot supply that growth. This is **[COND | growing supply]**, not an unconditional necessity theorem for general Navier–Stokes solutions or a theorem localizing every interaction to one annulus.

3. The structural exponent theorem

Proposition 3.1 (width law) [EXACT for the strained-core model]. Suppose a regular cross-sectional axial vorticity obeys

$$\partial_t \omega = \sigma r \partial_r \omega + 2\sigma\omega + \nu_{\text{phys}}(\partial_r^2 + r^{-1}\partial_r)\omega, \qquad \rho^2 = \frac{\int_0^\infty r^3 \omega \, dr}{\int_0^\infty r\omega \, dr}.$$

Assume the denominator is nonzero, the moments define a positive width, and the integration-by-parts boundary terms vanish. The denominator is constant; differentiating the numerator gives

$$\frac{d\rho^2}{dt} = -2\sigma\rho^2 + 4\nu_{\text{phys}}, \qquad K = \frac{\sigma\rho^2}{\nu_{\text{phys}}}, \qquad \frac{\dot{\rho}}{\rho} = \sigma(2/K - 1).$$

Here σ is the radial compression coefficient in this normalization. This identity is not a consequence of incompressibility alone for an arbitrary core: nonuniform strain, truncated moments or additional vorticity transport produce extra terms. A cut-defined radius must separately satisfy the width balance used by the model. For $\rho^2 = cs$,

$$\sigma = \frac{c + 4\nu_{\text{phys}}}{2cs}, \qquad K = 2 + \frac{c}{2\nu_{\text{phys}}} > 2.$$

Theorem 3.2 (I-S: structural window) [EXACT in the specified similarity class]. Under A1–A5, SS and (W), set

$$G = \frac{\dot{\Lambda}}{\sigma\Lambda}, \qquad \eta_{\text{rate}} = \frac{\dot{\ell}}{\sigma\ell}.$$

Then

$$G = \alpha(1 - 2/K), \quad \eta_{\text{rate}} = -(1-\alpha)(1-2/K), \quad a_{\text{grow}} = \frac{G}{2(G - \eta_{\text{rate}})} = \frac{\alpha}{2},$$

$$K(1 + \eta_{\text{rate}} - G) = 2, \qquad E_{\text{core}} \asymp \Lambda^2 \ell \asymp s^{(1-3\alpha)/2}.$$

Consequently circulation growth and vanishing core energy hold together exactly for

$$\boxed{0 < \alpha < \frac{1}{3}.}$$

Proof. Differentiate $\Lambda = C_\Lambda \rho^{-\alpha}$ and use (W). The velocity amplitude is $\Lambda/\rho \propto s^{-(1+\alpha)/2}$; A5 fixes the axial length exponent $(1-\alpha)/2$. Differentiation gives the second rate, and substitution gives the identities. A fixed nontrivial core shape with comparable axial and swirl amplitudes has energy of order $(\Lambda/\rho)^2 \rho^2 \ell = \Lambda^2 \ell$. Its exponent is positive iff $\alpha < 1/3$; Λ grows iff $\alpha > 0$. □

The common factor $1 - 2/K > 0$ cancels from $-2G - \eta_{\text{rate}} = (1 - 2/K)(1 - 3\alpha)$. Thus a cut-defined K does not select the exponent within this family. This is an algebraic consequence of the similarity premises, not invariance under arbitrary changes of core definition. At $\alpha = 0$ circulation amplitude stays constant; at $\alpha = 1/3$ this core-energy scaling stays of order one. The theorem is neither a universal finite-energy criterion nor existence for every exponent in the window. Comparability $\Lambda \asymp \rho^{-\alpha}$ alone cannot be differentiated to prove (3.2).

4. Covariance geometry and the full profile state

Put $h = \alpha/2$, $A = (1+\alpha)/2$, $D = (1-\alpha)/2$. In profile coordinates $X \propto r^2$, $|\eta| < 1$, define

$$H = \sqrt{2X}\,E, \quad \epsilon = \log E, \quad D_X = X\partial_X, \quad l = D_X \log H, \quad a_s = 2 - 2l, \quad b_s = 2D_X U/E, \quad s_{\text{sh}} = (a_s, -b_s).$$

The profile state is $\Phi = (E, U, \Pi)$, with $E > 0$ for $X > 0$, $H \sim cX$ at the axis, U Cartesian-regular, and

$$\Pi_X = \frac{E^2}{2X}, \qquad \Pi(X,\eta) = -\int_X^\infty \frac{E(y,\eta)^2}{2y}\,dy.$$

4.1 The cone interface I-C [EXACT for the specified covariance family]

For the local interaction family of the transcribed reduction, write the demanded stress as $T = Fw$, $F > 0$, $w = p_s - s_{\text{sh}}$. With $a_s > 0$ and $v_s = |s_{\text{sh}}|^2/a_s > 2$, its open admissible cone is

$$\mathcal{C}(s_{\text{sh}}) = \{T : T \cdot s_{\text{sh}} > 0, \quad \tan^2\angle(T, s_{\text{sh}}) < 2a_s/(|s_{\text{sh}}|^2 - 2a_s)\}.$$

Indeed the covariance coordinates use the matrix $\begin{pmatrix} 1 & -b_s/a_s \\ b_s/a_s & 1 \end{pmatrix}$, whose Gram matrix is $(1 + b_s^2/a_s^2)I$. This conformal change preserves angles and makes the cone circular and shear-aligned. This is not a restriction on all conceivable positive-semidefinite covariance matrices; it is the cone of the specified wave family.

For $s_{\mathrm{sh}}^{\perp} = (b_s, a_s)$, define

$$M_0 = |s_{\mathrm{sh}}|^2 - 2a_s, \quad M_1 = w \cdot s_{\mathrm{sh}}, \quad M_2 = 2a_s(w \cdot s_{\mathrm{sh}})^2 - M_0(w \cdot s_{\mathrm{sh}}^{\perp})^2.$$

Nonzero demand is strictly admissible iff $M_0, M_1, M_2 > 0$, with $a_s > 0$. The sign condition $M_1 > 0$ cannot be dropped after squaring. Zero demanded stress is permitted separately. At a vanishing edge, uniformity is a condition on the factored direction of w, not a positive lower bound on M_1, M_2 themselves.

Proposition 4.1 (pointwise criterion). Given p_s, some admissible shear exists iff $p_{s,2} \neq 0$ or $p_{s,1} > 2$. A positive on-axis demand exists iff

$$p_{s,1} > 0, \qquad \Theta = |p_s|^2 - 2p_{s,1} > 0.$$

Proof. Write $s_{\mathrm{sh}} = q(\cos f, \sin f)$, $\cos f > 0$. The first two inequalities leave an interval of radii iff $(p_{s,1} - 2)\cos f + p_{s,2}\sin f > 0$; near its lower endpoint the aperture inequality is strict. This gives the first criterion. For the on-axis case, $w = ts_{\mathrm{sh}}$ with $t > 0$, hence $p_s = (1 + t)s_{\mathrm{sh}}$ and necessarily $p_{s,1} > 0$. Writing $s_{\mathrm{sh}} = qp_s/|p_s|$, the remaining interval is $2p_{s,1}/|p_s| < q < |p_s|$, equivalent to (4.4). Then $v_s \in (2, |p_s|^2/p_{s,1})$. □

The positive-first-component hypothesis was omitted in Attachment A, Theorem 9.2; for example $p_s = (-1, 0)$ has $\Theta > 0$ but cannot give a positive on-axis demand. The constructed layer in §5 has $p_{s,1} > 2$, so this correction does not change it. Pointwise freedom to choose a shear does not make that shear integrable into a globally compatible profile.

Since $D_X = \frac{1}{2}r\partial_r$,

$$\alpha_{\mathrm{loc}} = -2l, \qquad a_s = 2 + \alpha_{\mathrm{loc}}, \qquad v_s - 2 = \frac{\alpha_{\mathrm{loc}}(2 + \alpha_{\mathrm{loc}}) + b_s^2}{2 + \alpha_{\mathrm{loc}}}.$$

For $E = c_\infty X^{-A}$, $U = 0$ in the exterior,

$$s_{\mathrm{sh}} = (2 + \alpha, 0), \qquad v_s - 2 = \alpha, \qquad \theta_c = \arctan\sqrt{2/\alpha}, \qquad \Pi = -\frac{c_\infty^2}{2(1 + \alpha)}X^{-1-\alpha}.$$

Thus positive exterior exponent is exactly exterior cone nondegeneracy in this class. Cone existence is still distinct from membership of the demanded stress.

4.2 Sources, stress primitives and moments [EXACT | TRANSCRIPTION]

The reduced system used here is specified by

$$L = 1 - \alpha\eta^2, \quad d = 1 - \eta^2, \quad \mathrm{A}_X U = X^{-1}\int_0^X U(y, \eta)\, dy, \quad H_c = D\eta + dU, \quad W = 1 - 2D\eta\mathrm{A}_X U - d\mathrm{A}_X U_\eta,$$

$$S_q = -H_c\epsilon_\eta - Wl - h(1 - 2\eta U),$$

$$S_n = -H_c U_\eta - WD_X U - A(1 - 2\eta U)U + 4A\eta\Pi + 2\eta D_X\Pi - d\Pi_\eta.$$

Axis regularity fixes the primitive constants:

$$p_{s,1} = \frac{1}{LH}\int_0^X HS_q\, dy, \qquad p_{s,2} = \frac{1}{LE}\int_0^X S_n\, dy.$$

Consequently

$$D_X w_1 + l w_1 = \Sigma_q := XS_q/L - D_X a_s - l a_s, \qquad D_X(E w_2) = \Sigma_n := XS_n/L + 2D_X^2 U.$$

These are driven linear stress equations. Their sources contain height derivatives and pressure primitives, which a state consisting only of $(\alpha_{\mathrm{loc}}, b_s, p_s)$ at one radius does not reconstruct. No finite radial closure is supplied by these equations. In particular, an invariant-region argument for that projected state cannot replace profile compatibility.

The moment conditions are

$$c_Q(\eta) = \int_0^\infty HS_q \, dX = 0, \qquad c_N(\eta) = \int_0^\infty S_n \, dX = 0.$$

They express return of the stress to zero when the exterior itself solves the zero-stress equations. They are identities for each height. Two conditions on one slice are not two scalar conditions on an entire two-dimensional profile. The physical radial primitives likewise read $T_{r\theta} = -r^{-2} \int_0^r q^2 R_\theta \, dq$ and $T_{rz} = -r^{-1} \int_0^r q R_z \, dq$, with the corresponding two total moments zero, in the reduced radial residual matching problem.

Moment-repair interface [COND | TRANSCRIPTION and the hypotheses of Attachment A, Theorem 9.6]. For the repair collar and perturbation class used there, the linearized moment map into a space of two functions of η has weighted diagonal directions

$$\delta c_N|_{\delta U} = \int \partial_X(XW) \, \delta U, \qquad \delta c_Q|_{\delta E} = \int \partial_X(XHW) \, \delta E/E.$$

Under the stipulated height-jet restrictions its fibre matrix is triangular at $\eta = 0$; the lower-left term is $O(\alpha|\eta|)$. Nonzero diagonal weights and a uniform determinant bound give the bounded right inverse stated in that theorem. This export retains the source-transcription, collar, perturbation and smallness hypotheses; it does not assert the claim for arbitrary height-dependent variations, especially when their pressure-height derivatives contribute. A profile-space right inverse must preserve those jet restrictions and the required norm. Repair stays within a strict factored cone margin if its C^1 cost is small; in the repair scaling of Attachment A this is $\mu \ll \delta_0 W_{\mathrm{repair}}^2$ (Attachment A, Theorem 9.5).

5. The constructed midplane interface I-M

All results in this section are **[EXACT | TRANSCRIPTION]** for (4.1), (4.7)–(4.9). They specify a slice and its first height jet, not a solution on a height interval.

5.1 Exact active layer

On $[X_a, X_b]$, take $l = -h$, $v := U(\cdot, 0) = 0$, $E_\eta|_0 = 0$, with the axial flux $\Psi = \int_0^X U_\eta|_0 = 2 + \alpha$ deposited inboard. Then the layer is zero-stress provided $w(X_a, 0) = 0$. Given a smooth flat bump $\chi \geq 0$ supported in the layer, set

$$\delta E = \delta v = 0, \qquad \delta U_\eta|_0 = -\frac{1}{h}(D_X \chi - h\chi)'.$$

Theorem 5.1 (Attachment A, Theorem 9.13). For every real ε, the stress response is exactly

$$w_\varepsilon = w_* + \varepsilon(\chi, 0), \qquad \delta c_Q = \delta c_N = 0.$$

On the zero-stress base, $\varepsilon > 0$ gives strict admissibility on $\{\chi > 0\}$, with unchanged cone and zero demand at the edges.

Proof. At fixed l and $v = 0$, S_q is affine in the changed flux and $S_n = 0$. Equation (4.9) gives $\delta w_1 = \chi$, while $\delta c_Q = \int (H\chi)' = 0$. No quadratic remainder occurs. With $s_{\mathrm{sh}} = (2 + \alpha, 0)$, $w = \varepsilon \chi s_{\mathrm{sh}}/(2 + \alpha)$, so (4.3) is strict precisely on the open active set. □

The scalar base-preparation theorem (Attachment A, Theorem 9.14) supplies an alternative layer input by a monotone one-parameter adjustment. Its uniqueness is only within that chosen family. The global construction below supplies the input with a stress-free inboard preparation instead.

5.2 Stress-free global midplane

Choose smooth l with $l(0^+) = 1$, a single sign crossing at X_0, and $l = -h$ on $[X_s, \infty)$. Choose $v \neq 0$ on $[0, X_s)$, returning flatly to zero at X_s, with the tame cutoff estimates of Attachment A, Theorem 9.26. Write

$$N_0 = X(l+h) - 2Xl' + 2l(1-l), \quad P = 2(Xv')' - (X-\Psi)v' - Av - \Psi' v.$$

The equations for zero midplane stress are

$$l\Psi = N_0 + Xve_1, \quad e_1 = E_\eta|_0/E_0 = XP'/E_0^2, \quad \Pi_\eta|_0 = P, \quad P(\infty) = 0.$$

Eliminating e_1 gives the linear BVP

$$\frac{X^2v^2}{E_0^2}\Psi'' + \left(l - \frac{X^2vv''}{E_0^2}\right)\Psi = N_0 + \frac{X^2v}{E_0^2}\left(2Xv''' + (4-X)v'' - (1+A)v'\right),$$

with $\Psi(0) = 0$ and smooth matching through X_s.

Theorem 5.2 (Attachment A, Theorems 9.25–9.26). If

$$\min_{[0,X_0]} v^2 > Q := \int_0^{X_0} \frac{E_0(x)^2 l(x)}{x}\, dx,$$

then (5.4) has exactly one axis-regular, smoothly matched solution. It satisfies $\Psi = 2 + \alpha$, $E_\eta|_0 = 0$ beyond X_s; the demanded stress and both moments vanish on the prepared slice. Attaching (5.1) at any permitted $X_a \geq X_s$ yields zero demand inboard and outboard and strictly admissible demand in the open active layer.

Proof mechanism. Put $\varphi = \Psi/v$. The homogeneous operator becomes $(v^2\varphi')' + q\varphi$, $q = E_0^2 l/X^2$. Its boundary conditions eliminate the nonregular axis branch and the growing return branch. Since $\varphi(0) = 0$,

$$\int_0^{X_0} q|\varphi|^2 \leq Q\int_0^{X_0} |\varphi'|^2 < \int_0^{X_0} v^2|\varphi'|^2.$$

The negative part of q adds coercivity. Subtracting a smooth exterior particular solution, the variational problem has a unique solution; the tame return estimates give smooth flat matching. The complete singular-endpoint argument is in Attachment A, §9.16.9(h). □

This non-resonance condition concerns the **midplane lift**, not the height-frequency resolvent of §7. Below its threshold the homogeneous Wronskian can vanish; the listed numerical resonance locations are evidence about the tested profile only.

Axial structure is necessary for this global preparation. In the restricted class $E_\eta|_0 = 0$, a stress-free zero of l would require $l' = h/2 > 0$, preventing the needed downward crossing. Allowing $E_\eta|_0 \neq 0$ replaces that obstruction by $ve_1 = 2l' - h$. Attachment A, Theorem 9.18 also excludes the globally admissible preparation with $v \equiv 0$. Thus reflection symmetry $E(X,-\eta) = E(X,\eta)$, $U(X,-\eta) = -U(X,\eta)$ is an excluded global class here, although $v = 0$ remains valid on the active layer. The Return Lemma of the earlier route is redundant, not proved.

6. Height is a boundary-value problem: interface I-F

Setting $\Sigma_q = \Sigma_n = 0$ gives exactly

$$H_c\epsilon_\eta - dl\mathsf{A}_X U_\eta = \frac{2LD_X^2\epsilon - Lla_s}{X} - l(1 - 2D\eta\mathsf{A}_X U) - h(1 - 2\eta U),$$

$$H_c U_\eta + d\Pi_\eta - d(D_X U)\mathsf{A}_X U_\eta = \frac{2LD_X^2 U}{X} - (1 - 2D\eta\mathsf{A}_X U)D_X U - A(1 - 2\eta U)U + 4A\eta\Pi + 2\eta D_X\Pi.$$

Away from $H_c = 0$, the principal radial diffusion coefficient in a height march is $2LX/H_c$. Its sign changes across stagnation. A formal hierarchy at the midplane does not give convergence or arbitrary smooth two-sided continuation. The radial Volterra terms couple horizontal slices across a nonhorizontal stagnation set, so solving independent height evolutions and joining them is not a closed formulation. The global problem requires handover traces and appropriate inflow/outflow conditions. These conclusions are **[EXACT | TRANSCRIPTION]** at the principal and geometric levels of Attachment A, Theorems 9.27–9.35.

6.1 Scalar and coupled folds

Near a regular axis zero, use $r = 2\sqrt{X}$, $t = \eta - \eta_*$. Then $H_c = \kappa t + \gamma r^2 + \cdots$. The scalar model is

$$(\kappa t + \gamma r^2)q_t - c\Delta_2 q = f.$$

Under the small-cylinder coercivity and boundary hypotheses of Attachment A, Theorem 9.36, it is well posed. At a source fold $\kappa > 0$, with $\kappa T > |\gamma|R^2$, both height ends are outflow and receive no prescribed data. For $\gamma \neq 0$, the diffusion–drift double commutators generate ∂_t, giving the stated hypoelliptic scalar model. This does not remove the pair coupling in (6.1)–(6.2).

At its own axis weight $t \sim r^3$, the swirl–pressure pair and diffusion give

$$u_{tt} = \frac{4L^2}{d^2\kappa_E}(-\Delta_2)^3 u,$$

a quasi-elliptic principal system of type $(2, 6)$ (Attachment A, Theorem 9.37). The pressure trace remains a radial primitive. The model's radial traces are inherited from u, pressure and swirl, represented by $u, \Delta_2 u, \partial_r \Delta_2 u$; they cannot be prescribed independently of the original system. The factorization check at zero tangential/height frequency is exact. The reported nonzero-frequency singular value for the Lopatinskii–Shapiro complementing condition [6, 7] is **[NUM]** and is not used as a continuum proof here. The global height boundary operator still needs its own range and complementing analysis.

The infinite negative index of a frozen height-mass form does not imply instability of this coupled PDE. Neither does failure of a bounded positive diagonal multiplier: those multipliers omit nonlocal or differentiated structure. These negative results delimit methods, not the solution set.

6.2 The same modes and the normal form

For $l > 0$, the principal height-mass symbol at nonzero radial frequency is

$$\widehat{\mathfrak{A}} = H_c I + \frac{d}{ikX}\begin{pmatrix} 0 & -l \\ E^2 & -XU_X \end{pmatrix}.$$

Put $\varepsilon_m = dE\sqrt{l}/(kX)$, $\zeta = H_c/\varepsilon_m$, $\vartheta = XU_X/(2E\sqrt{l})$. Its two masses are $m_\pm = \varepsilon_m(\zeta \pm \sqrt{1 - \vartheta^2} + i\vartheta)$. The eigenvectors do not depend on H_c, k or ζ. Their balanced condition number is $\sqrt{(1+|\vartheta|)/(1-|\vartheta|)}$ for $|\vartheta| < 1$. Thus the core/outer transition changes scalar masses, not the two modes. Coalescence occurs at $|\vartheta| = 1$, equivalently $\mathrm{Ri} = 1/4$, outside a sufficiently small axis neighbourhood with the stated fixed coefficients (Attachment A, Theorem 9.38).

At a mode zero, the tilt and imaginary shear shift must be retained together. The normalized model is

$$(x + c_*\partial_x^{-1})a_t = \nu_f a_{xx}, \qquad c_* = -\tfrac{1}{2} \pm i\sqrt{\mathrm{Ri} - \tfrac{1}{4}}.$$

On its whole-line realization, the exact gauge $a = |D|^{c_*}b$ gives $xb_t = \nu_f b_{xx}$. A gauge valid on the whole line cannot be assigned to a cut-off window without proving the domain and commutator transfer.

Theorem 6.1 (Attachment A, Lemma 9.50) [EXACT on the specified graph domain]. Let $\nu_f > 0$, $\sigma \neq 0$, $c = -1/2 + i\omega$. For

$$P = -\nu_f \partial_x^2 + i\sigma(x + c\partial_x^{-1}), \quad v = |k|^{-c}\widehat{a}, \quad Bv = -\sigma v' + \nu_f k^2 v,$$

take $\mathcal{H}_* = L^2(\mathbb{R}, |k|^{-1}dk)$ and the maximal domain $v \in \mathcal{H}_* \cap AC_{\mathrm{loc}}(\mathbb{R}\backslash\{0\})$, $Bv \in \mathcal{H}_*$. Then

$$\| Pa \|_2 \geq \nu_f^{1/3} |\sigma|^{2/3} \| a \|_2 .$$

Proof mechanism. The gauge is isometric into $\mathcal{H}_*$. Scaling and reflection reduce to $B = -\partial_k + k^2$. On the positive half-line, integration by parts gives $\|Bv\|_*^2 = \int (|v'|^2/k + k^3|v|^2 + |v|^2)dk$; the weighted Hardy bound (cf. [8, Theorem 330]) controls the missing norm near zero. On the negative half-line the first-order energy is favourable. Finite graph norm forces both Fourier traces at zero to vanish, and graph-core approximation extends the estimate. This proves a lower bound through zero Fourier frequency; it does not prove onto-ness for arbitrary forcing. The unrestricted Airy Green operator has a separate range compatibility. Full proof: Attachment A, Lemma 9.50. □

7. Full frozen-window high-frequency compactness: interface I-HF

Fix $R, \nu_f, d, \kappa_E > 0$ and real $\gamma \neq 0$. These are frozen constants; $\nu_f = 2L$ in the profile normalization, not a redefinition of ν_{phys}. Define

$$H = L^2((0,R), \tfrac{\kappa_E}{4} r^3 dr) \oplus L^2((0,R), r\,dr), \quad \mathcal{L} = -\nu_f \mathrm{diag}(\Delta_4, \Delta_2).$$

Its domain is the radial Dirichlet Laplacian domain in the two corresponding balls, restricted to radial functions, with the regular axis realization and $e(R) = u(R) = 0$. Thus for smooth states the axis expansions are smooth in r^2; the graph domain is their appropriate closure. Define on this same space

$$\Psi(r) = \int_0^r q u(q)\, dq, \quad p(r) = -\frac{\kappa_E}{2} \int_r^R q e(q)\, dq,$$

$$\mathcal{A}_0(e,u) = \left(\gamma r^2 e - \frac{2d\Psi}{r^2}, \quad \gamma r^2 u - 2\gamma\Psi + dp \right), \quad \mathcal{R}(\tau) = \mathcal{L} + i\tau\mathcal{A}_0.$$

The pressure endpoint is $p(R) = 0$; it is not a free constant. The shear term $-2\gamma\Psi$ is retained.

Theorem 7.1 (Attachment A, Theorem 9.60) [EXACT for (7.1)–(7.2)].

$$\inf_{\substack{z \in D(\mathcal{L}) \\ \| z \|_H = 1}} \| \mathcal{R}(\tau) z \|_H \to \infty \quad (|\tau| \to \infty).$$

Equivalently, if $|\tau_n| \to \infty$ and $\|z_n\|_H + \|\mathcal{R}(\tau_n) z_n\|_H$ is bounded, then $\|z_n\|_H \to 0$.

Proof structure; complete proofs are in Attachments C and D, numbered as in Attachment A. Set $t = |\tau|$, $b = d\sqrt{\kappa_E}$, $\rho_t = (\nu_f b^3/(7\gamma^4 t))^{1/7}$. The five steps estimate actual original-domain states throughout.

1. **Axis concentration.** Compactness of the primitive part of $\mathcal{A}_0 = \gamma r^2 I + K_0$, its trivial kernel, and the exact resolvent pairings force every hypothetical bounded-residual normalized sequence to lose mass on each fixed annulus (Proposition 9.51).
2. **Reflection and transport.** The pressure endpoint relates the two oscillatory primitive coefficients. Transfer through the outer viscous layer and compactness of corrected amplitudes on r/ρ_t pass this relation to the transport limits. With that relation the limits vanish; without it the transport model has a decaying kernel (Propositions 9.52–9.56).
3. **Physical reconstruction.** An exact differential identity reconstructs the physical state from those coefficients. It proves $\|1_{r \geq \epsilon\rho_t} z_n\|_H \to 0$ for every fixed $\epsilon > 0$, with the derivative control needed at an averaged endpoint (Proposition 9.57).

4. **Axis core.** In isometric coordinates $y = ((\sqrt{\kappa_E}/2)re, u)$, a regular self-adjoint radial Dirac operator satisfies

$$(\nu_f D_0^3 + i\tau b)y + i\tau\gamma r^2 D_{\mathrm{m}} y = D_0 \widehat{f}.$$

Its cubic inverse and controlled localization yield, for every fixed $B > 0$,

$$\| 1_{r \le Bt^{-1/6}} z \|_H \le C_B(t^{-2/3} \| f \|_H + t^{-1/12} \| z \|_H)$$

(Proposition 9.58).

5. **The intervening interval.** The exact six-component system retains both primitives, both physical components and both radial derivatives. Its signature estimate, with inner endpoint observation and outer endpoint reconstruction, removes the interval between $Bt^{-1/6}$ and $\epsilon\rho_t$ (Lemma 9.59 and Theorem 9.60). No nonzero transport-scale limit is assumed. Together the regions exhaust the original norm. □

The independent model-transfer proposal (Q) is superseded **for this compactness objective**. The proof attaches the original pressure condition to the original state and closes its full norm; it does not infer a window theorem from a whole-line gauge alone.

The exact small-frequency estimate of Attachment A, Proposition 9.44 is $\|\mathcal{R}(\tau)z\|_H \ge \lambda_1\|z\|_H/2$ for $|\tau| \le \lambda_1/(2\|S\|)$, where S is the skew part of the bounded frozen mass and λ_1 is the bottom Dirichlet diffusion eigenvalue. The remaining compact frequency band is a separate non-resonance problem. Neither estimate supplies surjectivity by itself.

Full-window calculations support growth near $|\tau|^{4/7}$ **[NUM]**. Theorem 7.1 asserts divergence of the lower bound, with no proved exponent and no uniformity over varying frozen parameter families unless separately established.

8. OA as a conditional realization

The external construction called "OA" here is OpenAI's *Finite Time Blowup for Navier–Stokes* [2]. It has the stated parameter $0 < h < 1/100$ and the following dictionary. These are comparisons with that source; this paper does not certify its full theorem or make a priority claim. A Lean formalization accompanies the construction [9]; the results here do not use it. Attachment E lists the Lean declarations that correspond to the entries of the dictionary below.

Item	Class calculation	OA comparison and status
Exterior circulation	$\Gamma_{\mathrm{ext}} \sim r^{-\alpha}$	Stated far field r^{-2h} gives $\alpha = 2h < 1/50$. **[MAP]**
Circulation and lengths	$a_{\mathrm{grow}} = \alpha/2$, $\rho \sim s^{1/2}$, $\ell \sim s^{(1-\alpha)/2}$	Substitution gives $a_{\mathrm{grow}} = h$, $\ell \sim s^{1/2-h}$. **[EXACT | MAP and similarity premises]**
Core energy	$s^{(1-3\alpha)/2}$	$s^{1/2-3h}$. **[EXACT | MAP]**
Exterior cone	$v_s - 2 = \alpha$	$v_s - 2 = 2h$. **[EXACT for the displayed profile and cone]**
Source-free core and interaction field	Must be verified for an actual field	OA is a realization and a nonemptiness witness only conditional on its correctness and the dictionary. **[COND-OA]**

Item	Class calculation	OA comparison and status
Full moment repair and oscillatory realization	Level B has obligations beyond the two Level A moments	The external implementation is not independently discharged here. **[COND-OA]**

The exterior exponent, stress return, cone geometry and fold models are structural within the stated premises. The progressively finer pulse hierarchy of [2] may be universal, restricted to part of the class, or specific to OA's implementation. The present results do not decide between those possibilities. Its necessity is a Paper II question, not an added class axiom.

9. The theorem boundary

Paper I supplies a structural classification within explicit premises, exact cone and stress identities, a constructed midplane for the transcribed system under (NR), fold-model results on their specified domains, and full frozen-window high-frequency compactness. Each interface can be used and tested without assuming that a complete singular solution exists.

Deferred beyond Paper I: finite-frequency non-resonance on the remaining compact band; surjectivity and range closure where required; restoration of the source-fold term κt; height outflow and transmission traces; transfer to variable profile coefficients; global inboard patching; and the interaction field, regularity, forcing and final Navier–Stokes verification needed for complete constructions. Paper II imports these interfaces and must close these obligations without identifying Level A covariance realizability with Level B field realization.

Attachments

- **Claim ledger:** the current status of every claim in this paper and in Attachments A–D, with the corrections made to Attachment A.
- **Attachment A**, *The Exterior Exponent of a Collapsing Vortex:* the complete technical manuscript, cited here by its theorem numbers. It presents the analysis in the order it was carried out, with superseded statements kept and marked.
- **Attachment B**, *Cone Realizability and the Midplane Construction for Lemma H:* covariance algebra and profile-state analysis.
- **Attachment C**, *Global Quasimodes of the Frozen Axis-Fold Operator:* primitive reflection, transport and physical reconstruction.
- **Attachment D**, *Compactness Closure for the Full Frozen Axis-Fold Window:* axis-core and interval estimates and the complete proof of Theorem 9.60.
- **Attachment E**, *Correspondence with the Lean Formalization of the OA Construction:* where the OA-dependent statements of §8 appear in the Lean formalization [9].

Attachments A and B compare with the OA construction [2] and its Lean formalization [9], and Attachment B cites D. Gondauri's audit [10] by section. Attachment C cites DLMF §§10.5, 10.6 and 10.17 [11]. Attachment A cites the decomposition rules of [12].

Code availability

The scripts for the symbolic and numerical audits cited in this paper and its attachments are available at https://github.com/LarCorps/ns-exterior-exponent-calculations [13]. Numerical results are audits and are not used in place of proofs.

Use of AI tools

Large language model tools, including Anthropic's Claude, assisted with derivations, symbolic and numerical audits, source verification, and the drafting and editing of the text. The author directed the research, checked the results and takes full responsibility for the content.

References

1. L. Korpela, *Paper II — Constructive Realizations of the Exterior-Exponent Collapse Class*, companion manuscript in preparation, 2026.
2. OpenAI, *Finite Time Blowup for Navier–Stokes*, preprint, 166 pp., September 2026. https://cdn.openai.com/pdf/32d9f210-8b73-45e0-91bc-82a30aef8a9a/navier-stokes.pdf. Version compared: PDF dated 2026-09-08, SHA-256 `0e779481c4da40bd28d1e642e1d8ca57447d129610df28dfa5a11e9af8ae228f`.
3. "Navier-Stokes Equations" (cylindrical and spherical coordinates), course handout for ME 19b Fluid Mechanics (instructor C. E. Brennen), California Institute of Technology, February 2010. https://www.its.caltech.edu/~mefm/me19b/handouts/Equations.pdf.
4. T. Y. Hou and C. Li, Dynamic stability of the three-dimensional axisymmetric Navier–Stokes equations with swirl, *Comm. Pure Appl. Math.* 61 (2008), no. 5, 661–697. https://doi.org/10.1002/cpa.20212.
5. L. C. Evans, *Partial Differential Equations*, 2nd ed., Graduate Studies in Mathematics 19, American Mathematical Society, Providence, RI, 2010. https://doi.org/10.1090/gsm/019.
6. Ya. B. Lopatinskii, On a method of reducing boundary problems for a system of differential equations of elliptic type to regular integral equations, *Ukrain. Mat. Zh.* 5 (1953), no. 2, 123–151 (Russian).
7. Z. Ya. Shapiro, On general boundary problems for equations of elliptic type, *Izv. Akad. Nauk SSSR Ser. Mat.* 17 (1953), 539–562 (Russian).
8. G. H. Hardy, J. E. Littlewood and G. Pólya, *Inequalities*, 2nd ed., Cambridge University Press, Cambridge, 1952.
9. OpenAI, *NavierStokesAndEuler: Lean certificates accompanying Navier–Stokes and Euler results*, GitHub repository, Apache-2.0. https://github.com/openai/NavierStokesAndEuler, accessed 2026-09-15 at commit `f9e8bc5`.
10. D. Gondauri, *Critical Audit of OpenAI's Navier–Stokes Proof: Scope of the Mathematical Audit, Independent Audit and Certification Status of the Lean Formalization, and Procedural Readiness for the Millennium Prize*, Zenodo, 10 September 2026. https://doi.org/10.5281/zenodo.22694089.
11. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl and M. A. McClain (eds.), *NIST Digital Library of Mathematical Functions*, Release 1.2.7 of 2026-06-15, https://dlmf.nist.gov/; Chapter 10, Bessel Functions, by F. W. J. Olver and L. C. Maximon.
12. L. Korpela, *RRFT Working Rules: A practical anti-flattening protocol for reasoning, modelling, and research*, unpublished working document, 2026.
13. L. Korpela, *Supporting calculations: the exterior-exponent collapse class for Navier–Stokes*, GitHub repository, 2026. https://github.com/LarCorps/ns-exterior-exponent-calculations.