Unfaithfulness in the Navier–Stokes formalization

HumanLean Team

November 2026

What counts as unfaithful?

A formal proof establishes the claim written in code. Faithfulness asks whether that claim and its proof preserve the mathematics of the source text. We assess two dimensions separately:

- **Definition and statement faithfulness:** the formalization preserves the intended objects, assumptions and conclusions. A changed construction or weaker conclusion can alter what the result establishes.
- **Proof faithfulness:** the formalization preserves the substantive argument of the source proof. Routine changes in notation or implementation do not count. Replacing the central argument with another method counts as a proof difference, even when both methods prove the same theorem correctly.

Findings concern the numbered result or its own argument. Differences inherited solely from dependencies are excluded. Direct registrations of existing Mathlib declarations and transparent delegates to them are excluded.

Findings

In the formalization of *Finite Time Blowup for Navier–Stokes*, we identify substantive differences in **24 numbered results**:

- **3** results have definition, construction or statement mismatches.
- **21** results have substantive proof replacements.

These groups do not overlap. The semantic findings concern what the formalization expresses. The proof findings concern how it establishes its conclusion.

The findings below follow the sections and appendices of the paper. Each entry explains the difference and links the original argument and formal code needed to check it.

Section 4. Constructing the leading order flow

Proposition 4.2 — Definition mismatch

The source defines the integrated scalar profiles Q_s, N_s by their smooth extensions through the axis $X = 0$. Equations (4.9)–(4.10) give $Q_s(0,\eta) = S_q(0,\eta)/2$ and $N_s(0,\eta) = S_n(0,\eta)$, which can be nonzero. The generic formal definitions `angularLag` and `axialLag` divide their primitives by XH and X. Lean's division-by-zero convention makes both values zero at the axis, and the linked smoothness theorems exclude that point. A removable singularity needs its limiting value to preserve scalar-profile regularity. The mismatch concerns the generic axis extension attached to the proposition. The linked off-axis identities retain nonzero-radius hypotheses.

Paper p. 26 · Paper p. 27 · ProfileHistories:584–617.

Lemma 4.7 — Proof difference

The source proves moment-matrix invertibility by writing its determinant as an integral of a determinant with constant nonzero sign. The formal proof starts with a vector in the left kernel. Its weighted integrals vanish, so the corresponding power sum has a root in each separated interval. The Rolle-based bound on the number of roots forces the vector to vanish. Both routes use the zero-count property of distinct powers, but the formal proof replaces the central determinant-integral argument with kernel injectivity. This mechanism is shared with Lemma A.1.

Paper p. 34 · PowerMomentMatrix:158–249.

Proposition 4.10 — Proof difference

To preserve the relaxed stress cone during restoration and moment matching, the source controls the forward cumulative moment discrepancies, then chooses tolerances for the individual factors E, Q_s, a, b_s, N_s. The formal argument reconstructs the cumulative integrals backward from exact matching at the right endpoint. It bounds a finite-dimensional vector of field derivatives and integrated quantities, then controls shear speed and the stress projection by compact perturbation estimates. The cone conclusion is reached through a different quantitative stability argument, shared with Proposition B.8.

Paper p. 36 · Paper p. 38 · RepairConeBounds:489–528 · RepairConeBounds:706–784 · RepairConeBounds:1220–1255.

Section 5. Correcting the base flow to every order

Proposition 5.3 — Proof difference

In the source's physical-derivative step, a term $q^b g$ with a transverse, k axial and l time derivatives is bounded by $q^{b-a/2-Dk-l}$ times an explicit profile norm and a polynomial in b. The formal proof uses homogeneity and a compact normalized derivative bound, then estimates all directions by the full dilation operator norm. It obtains q^{b-m} for $m = a + k + l$, with its constant chosen after b and the profile. This replaces anisotropic bookkeeping inside the proof. The numbered residual bound (5.25) allows an N-independent loss K_m and is not itself shown to fail.

Paper p. 55 · Paper p. 56 · PhysicalCoordinateBounds:310–362 · BaseResidual:154–165.

Proposition 5.5 — Statement mismatch

The source promises an asymptotic expansion with the original finite truncation at each fixed order N. Its remainder argument retains the gain $2h(N+1)$. The located formal fixed-prefix interface supplies $h(N+1)$, before the stated leading-power and derivative losses. For $0 < q < 1$, a larger positive exponent gives faster decay, so this is a weaker estimate at the same truncation. Taking a later prefix does not state the rate for the original one. The construction still sums vector potentials before differentiation to preserve divergence. This finding concerns the located interface. It establishes neither a violation of the sharper rate by the actual field nor the absence of another sharp formal estimate.

Paper p. 60 · Paper p. 61 · Paper p. 62 · SlowBorelBase:1646–1686 · BaseResidual:739–768.

Section 6. Auxiliary torus and separation of oscillatory supports

Lemma 6.1 — Proof difference

The source greedily colors the bounded-degree overlap graph, then selects rational torus centers from the open dense complement of finitely many forbidden equalities. The formalization assigns colors by modular arithmetic on the band and grid indices, and constructs rational centers with explicit denominators. Algebraic bounds exclude their forbidden covering relations. Both constructions separate the supports, but the existence mechanism changes from graph coloring and generic selection to an explicit arithmetic construction.

Paper p. 65 · Paper p. 66 · SlotColoring:143–188 · SlotGeometry:95–207.

Section 7. Oscillatory realization and correction of the residual stress

Proposition 7.2 — Proof difference

The source factors the harmonic propagator exactly into the first-harmonic propagator and the scalar damping factor in (7.18). Duhamel's formula then controls the forced solution. The formal argument derives an energy inequality directly for each nonzero harmonic, using its extra damping only through an inequality. Exact comparison of propagators is replaced by a bound that discards surplus damping. This changes the core comparison mechanism while retaining the intended uniform control of higher harmonics.

Paper p. 77 · Paper p. 79 · PrimaryODE:378–396 · ActualParticularControl:44–57.

Corollary 7.3 — Proof difference

The source invokes the exact harmonic factorization (7.18) in this corollary's own proof to keep one domain for all finite harmonic sets. The formal construction uses direct per-harmonic energy control with the same rate and error constant for every nonzero harmonic. Uniform dissipativity supplies the common threshold without identifying each propagator with a scalar multiple of the first. This is a local comparison-step replacement shared with Proposition 7.2, rather than a flag inherited solely from that proposition.

Paper p. 80 · ActualParticularControl:44–57 · ScaledActualParticularControl:139–167.

Lemma 7.4 — Proof difference

The source bounds the ratio of decaying to growing coordinates while the growing coordinate is positive. Its logarithmic envelope estimate prevents a first zero. The formal proof first uses backward uniqueness for the homogeneous linear ordinary differential equation to show that the whole vector never vanishes. An invariant-cone argument and continuity then give positivity of the growing coordinate and its ratio bound. Whole-vector uniqueness replaces the source's envelope-based continuation step, although both arguments control the same growing mode.

Paper p. 80 · Paper p. 81 · GrowingMode:43–67 · GrowingMode:128–204 · GrowingMode:365–426.

Section 8. Compactly supported mean corrections

Lemma 8.2 — Proof difference

Two local mechanisms differ within this lemma's proof. **Edge weights:** the source bounds a flat radial primitive by eventual monotonicity of its exponential weight and bounded integration length. The formal proof changes variables, factors out the flat weight, proves smoothness of the remaining improper integral by dominated differentiation, and bounds that factor on a compact interval.

Cutoff remainder: the source's Fourier-multiplier bound uses $m + p + 3$ torus derivatives and $I_R + p$ radial derivatives at the requested slow orders. Repeated formal inverse estimates instead use full joint derivatives through $m + 6p$, or $m + 6p + 4$ after mean subtraction. Here p is the number of integrations by parts, distinct from pressure. The quantitative derivative argument changes. The numbered lemma requires finite derivative dependence and leaves the exact budgets unspecified. Smooth-input flatness is not refuted, and no uniformity in p is required.

Paper p. 90 · Paper p. 91 · Paper p. 92 · WeightedRadialPrimitive:93–120 · FlatPrimitiveFactor:128–178 · FlatPrimitiveFactor:233–280 · UniformFourierAlias:804–881 · UniformFourierAlias:1209–1228.

Proposition 8.3 — Construction mismatch

The source fixes a unit-integral pressure bump supported in the reserved mean patch I_m before defining (8.12). The formal `commonGauge` passes the active shell's endpoints to pressure reconstruction. Its normalized bump is centered in that full interval, rather than confined to the prescribed mean patch. Unit integral still cancels the integrated defect, but normalization alone does not enforce the source's spatial placement. The mean patch is reserved because the base has an exact power-law form there. The mismatch concerns the chosen pressure density, rather than the algebraic zero-integral cancellation.

Paper p. 92 · Paper p. 93 · PressureStream:28–59 · CorrectionInitialization:4255–4262 · CommonBaseContext:118–125 · CorrectionState:270–274 · PrimaryTargetBounds:667–678.

Lemma 8.7 — Proof difference

The source geometrically dilates one bump so its power moments form explicit Vandermonde blocks. The formalization chooses separated intervals on an additive mesh and uses a localized moment inverse. Its nonsingularity proof extracts roots from a kernel relation and applies the Rolle-based zero-count argument. Thus the five-row solve retains its power-law patch and distinct-power structure, while replacing both the bump construction and the central rank argument. Constants are allowed to

depend on fixed positive λ, without uniformity as λ tends to zero.

Paper p. 97 · Paper p. 98 · FiveRowRank:24–114 · LocalizedMomentRepair:89–100 · PowerMomentMatrix:216–249.

Section 9. Residual improvement and the local field

Lemma 9.7 — Proof difference

The source's own common-domain proof explicitly uses (7.18) to compare every harmonic with the first, then counts the finitely many input derivatives needed at each stage. The formal modal-control construction uses direct energy bounds for each harmonic instead of exact harmonic factorization. It retains the useful distinction between one common domain and stage-dependent constants. The harmonic-comparison replacement is shared with Proposition 7.2 and Corollary 7.3.

Paper p. 111 · Paper p. 112 · ActualParticularControl:44–57 · ScaledActualParticularControl:139–262.

Section 10. Compact forcing and whole-space breakdown

Proposition 10.1 — Proof difference

The source cuts off a representation $u_{\mathrm{loc}} = \operatorname{curl} A + Be_\theta$ whose potential A vanishes in the heat exterior. The actual compact candidate localizes the raw potential instead. If J is the axial component moved into the direct field by rebundling, the cutoff product rule adds $\chi_t \nabla \chi_x \times J$ relative to localization of the rebundled representation. The local velocities agree before cutoff, but their transition terms differ. A faithful exterior-zero rebundling is also implemented in `LocalPotentialRebundle`; the finding concerns the representative used by the compact candidate.

Paper p. 117 · Paper p. 118 · R3CompactCandidate:201–205 · LocalPotentialRebundle:8–37 · PaperLocalization:26–50.

Lemma 10.2 — Proof difference

In the exterior part of the source proof, $A = 0$ leaves only the heat velocity and its pressure to extend through the terminal time. The formal compact construction also controls the raw axial heat primitive. `TailGaugePotential` anchors that primitive at physical $s = 1$, proves its smooth extension, and passes its one-sided endpoint extension into the summed potential. This is an additional analytic step required by the chosen representative, shared with the localization difference in Proposition 10.1.

Paper p. 118 · Paper p. 119 · TailGaugePotential:180–204 · TailGaugePotential:329–344 · LocalScheduleWitness:43–99.

Lemma 10.5 — Proof difference

The source's uniqueness proof localizes the stress tensor with ϕ_R^4, estimates it in $L^{3/2}$, and pairs the pressure contribution with an L^3 test. Its flux bound (10.19) is expressed in $B_R = \|\phi_R^4 w\|_6$. The formal proof uses ϕ_R^2-localized stress in $L^{6/5}$ and an L^6 test controlled by its L^2 derivative. This introduces A_R/R

and R^{-2} terms, where A_R is weighted dissipation's square root. Both routes absorb flux into dissipation, but the Hölder/Sobolev mechanism differs inside the proof.

Paper p. 121 · Paper p. 122 · Paper p. 123 · PressureFlux:289–353 · PressureFlux:576–598 · WholeSpaceUniqueness:56–69.

Appendix A. Matching radial moments and constructing the heat exterior

Lemma A.1 — Proof difference

The source integrates a constant-sign generalized-power determinant to prove the moment matrix nonsingular, then substitutes $x = e^y$ for the logarithmic version. The formal proofs use kernel injectivity: zero weighted averages supply one root per separated interval, and the zero-count theorem forces all coefficients to vanish. Both the power and exponential versions implement this root-extraction route. The central proof difference is shared with Lemma 4.7. Compact-family inverse bounds do not assert uniformity when distinct exponents coalesce.

Paper p. 126 · Paper p. 127 · PowerMomentMatrix:158–249 · ExponentialMomentMatrix:40–63.

Lemma A.8 — Proof difference

The source proves that the total weighted tangential residuals vanish by conserving a renormalized angular moment, canceling transport and viscosity boundary terms, and using Fubini for pressure. It then rewrites the stress as a backward integral. The formal proof starts from explicit forward stress primitives, identifies their moment expressions, and proves that they tend to zero along the heat tail. Integration of their derivatives backward from infinity gives the terminal stress. Explicit primitive identities and tail limits replace the source's global conservation calculation. The comparison is for the constructed matched profile and its exact moment conditions.

Paper p. 140 · Paper p. 141 · TerminalHistoryBridge:928–1000 · TerminalHistoryBridge:1172–1223.

Appendix B. Analytic profiles near the axis and their continuation

Proposition B.2 — Proof difference

The source obtains the mixed-error estimate (B.13) from the coefficient norm and radial Cauchy estimates on a smaller complex disk. The formalization differentiates the coefficient series term by term and bounds its derivatives by summable polynomial-geometric majorants. `evaluated_mixed_error` transfers the coefficient-space error through these bounds. The derivative estimate follows a different analytic mechanism while retaining a smaller evaluation radius $R < 20$ and the original scaling $\partial_X^k = \Lambda^k \partial_Y^k$.

Paper p. 146 · Paper p. 148 · AxisEvaluation:57–131 · AxisContraction:693–708 · NaturalEntrance:970–1006.

Proposition B.3 — Proof difference

The source tracks χ-sensitive logarithmic derivatives in (B.20)–(B.21), then uses Young's inequality to absorb the square-root error into the growing $\Lambda L \chi$ term. The formal argument packages the source into

a continuous finite-derivative model over compact bounded data. It proves a strict positive margin on $\chi = 0$, uses compact absorption away from that zero set, and controls small perturbations uniformly. This replaces the explicit gradient estimate. The zero-set check is essential because the growing term gives no help there. The mechanism is shared with Lemma B.4 and Proposition B.5.

Paper p. 148 · Paper p. 149 · Paper p. 150 · ReferenceBounds:447–509 · NaturalEntrance:86–140.

Lemma B.4 — Proof difference

The source carries the axis-profile gradient bounds through the reference cutoff and absorbs the $\sqrt{\chi}$ and χ errors to preserve a positive angular source. The formal reference-bound chain uses compact models of finitely many field derivatives, a positive source on $\chi = 0$, and uniform perturbation bounds. The resulting positivity argument replaces the explicit logarithmic-gradient absorption within this lemma's continuation proof. The comparison concerns that local estimate, with the reference family's width-uniform bounds remaining part of the source's stated construction.

Paper p. 150 · Paper p. 151 · ReferenceBounds:447–509 · NaturalEntrance:86–140.

Proposition B.5 — Proof difference

Two mechanisms differ in this continuation proof. **Source control:** the source reuses the explicit logarithmic-gradient absorption to preserve the relaxed cone after reducing the shear. The formal proof uses compact zero-set models and small-perturbation estimates for the corresponding source bound.

Parameter choices: the source uses one coupled reference/activation width. It chooses damping κ_0 uniformly over that reference family before choosing the width. The formal construction fixes a reference width, obtains bounds for its integrated quantities, and then chooses a separate activation width and damping. The changed order and coupling replace the source's continuation argument. These are proof findings within the existential proposition. The source's flat stress factorization and its distinction between finite prescribed smallness and higher-order finite bounds remain material to the construction.

Paper p. 151 · Paper p. 152 · Paper p. 153 · ActivationContinuation:262–325 · ActivationContinuation:1680–1738 · ReferencePath:38–67 · OutgoingSchedule:24–39.

Proposition B.8 — Proof difference

The source estimates the forward normalized moment discrepancies, then controls individual factors in (B.36) to preserve the relaxed cone during restoration and five-moment correction. The formal proof recovers the integrated quantities backward from their exact right-end match and compares a finite-dimensional vector of field derivatives and integrals with the outgoing model. Compact perturbation bounds directly control shear speed and stress projection. This replaces the cone-stability proof, while preserving the exact five-row matching and a tolerance chosen from outgoing data before the incoming axis stage. The mechanism is shared with Proposition 4.10.

Paper p. 155 · Paper p. 156 · Paper p. 157 · RepairConeBounds:489–528 · RepairConeBounds:742–784 · RepairConeBounds:1203–1255.

Appendix C. Realizing the admissible stress cone

Lemma C.1 — Proof difference

The source first proves pointwise divergence of the variance with raw amplitude μ, uses a finite cover to choose $\mu_{\max}$, and chooses a cone margin for every $0 \le \mu \le \mu_{\max}$. It then solves for the required variance. The formal proof constructs the global variance inverse first. It bounds the inverse's compact image and chooses the cone margin over the family of solved variance roots. Compactness is applied to a differently parametrized family and in a different order. This replaces the intermediate margin argument (C.7)–(C.10). No failure of the final cone construction or circularity is established.

Paper p. 158 · Paper p. 159 · Paper p. 160 · LoopVariance:902–918 · TrueConeLoop:164–194 · TrueConeLoop:459–499.

Sources and verification

The report compares the source paper with formalization revision `f9e8bc5`. Findings concern the linked implementations; alternative faithful implementations have not been exhaustively ruled out. This is a source-level comparison, without a new Lean compilation or kernel check.

Conventions. Paper page numbers match the PDF pages. Code links fix both revision and line range. We retain the paper's symbols: $q > 0$ is the shrinking scale, X the scaled radius, η the axial parameter, and $h > 0$ a fixed exponent. A field is *flat* if every fixed derivative decays faster than every prescribed power of q. A finite collection of derivatives is also called a *jet*. The stress cone describes which momentum fluxes can be realized by the oscillatory construction.