

The Lean Code Is Verified; the Paper’s Argument Is Not: A Counterexample to the Pressure-Flux Estimate in “Finite Time Blowup for Navier–Stokes”

Bevan Hunt

DOI: PENDING [CC BY 4.0](#)

Self-published, no affiliation. Released under the Creative Commons Attribution 4.0 International License
(creativecommons.org/licenses/by/4.0).

repository: <https://github.com/truemagic-coder/ns-proof>

Abstract

In October 2026 OpenAI published a 166-page manuscript, “Finite Time Blowup for Navier–Stokes,” claiming Theorem 1.1: for every viscosity, a smooth force, compactly supported in space and time, drives a solution that starts from rest and develops unbounded velocity in finite time while its kinetic energy stays bounded. The proof’s non-existence clause – that no smooth solution with the same data and bounded energy exists – is part of alternative (C) in the statement of the Millennium problem. The manuscript shipped with a Lean 4 formalization, and an independent verification service (FalseGreen) returned a fully green verdict: the fork builds, every theorem re-checks under two independent Lean kernels, and the axiom audit is clean. This paper shows that the green verdict verifies the *code*, not the *paper*. First, the Lean formalization proves a strictly different argument from the manuscript: an A_R -dependent pressure-flux bound where the manuscript states a B_R -only one. Second, the paper’s own argument is invalid at estimate (10.19): its derivation invokes three claims that are false for every admissible choice of cutoff, and the estimate itself is false under its own quantitative hypotheses – we give a smooth, divergence-free, uniformly L^2 -bounded family that satisfies every bound the derivation invokes, with $B_R = 0$ exactly, whose pressure flux is nonzero. Meanwhile the other flagged divergence – the (8.19) torus-inverse estimate – we formalized in Lean and the kernel accepted: it is sound as written. The defect is localized and repairable; the theorem’s conclusion is not thereby refuted. What is refuted is the paper as written.

1 Introduction

The manuscript [9] (166 pages, with a Lean formalization in the `NavierStokesAndEuler` repository) claims:

For every positive viscosity, we construct a solution of the three-dimensional incompressible Navier–Stokes equations that starts from rest and develops unbounded velocity in finite time while maintaining uniformly bounded kinetic energy.

The claim is a millennium-scale one: Theorem 1.1 establishes alternative (C) in Clay’s statement of the Navier–Stokes existence and smoothness problem [1] – the existence of a smooth solution with the same force and zero initial data is ruled out once kinetic energy is assumed bounded. The

manuscript accompanied the claim with a Lean 4 proof of a formal statement, and the natural reading of “formal proof” in 2026 is that the mathematics is *checked*: the Lean kernel either accepts a proof or it does not.

That reading is what makes this paper’s subject important, and it is exactly where it breaks. A green verification verdict on the Lean code establishes only that the *formal statement* is a valid Lean theorem. It says nothing about whether the formal statement matches the natural-language argument on the page. This paper shows that, for the OpenAI manuscript, both halves of that gap matter:

1. **The code and the paper are different arguments.** The Lean formalization proves a pressure-flux bound that depends on the weighted-dissipation quantity A_R ; the manuscript states estimate (10.19), which depends on the L^6 quantity B_R alone (§7).
2. **The paper’s own argument is invalid at (10.19).** The derivation of (10.19) uses three pointwise/norm claims, each false for every admissible choice of cutoff (§4); and the estimate is false under its own quantitative hypotheses – a counterexample family with $B_R = 0$ exactly and nonzero flux (§5).
3. **The failure is localized.** The other flagged divergence, the (8.19) torus-inverse estimate, we formalized in Lean 4.34.1 against mathlib and the kernel accepted it: the paper’s (8.19) argument is sound as written (§6).

The headline is therefore not “OpenAI is wrong about Navier–Stokes.” It is precise: *the Lean result is verified; the paper’s argument is not – and at (10.19) the paper’s argument is invalid as written.*

This paper is structured as a verification report. Every claim is tagged with the evidence that supports it: the verbatim manuscript text (quoted with page numbers and checked mechanically against the PDF), the Lean commits (pinned), the signed verification report (job 31974058-a665-472d-b6f8-00e4e schema 3), and the counterexample script (numerically regenerable; §A).

2 Claims

This paper makes four claims about the OpenAI manuscript “Finite Time Blowup for Navier–Stokes” [9]. Each is stated here and evidenced in the section named.

Claim 2.1 (C1 – the code is a different argument). *The Lean formalization that passes independent verification proves a strictly different argument from the manuscript: an A_R -dependent pressure-flux bound where the manuscript states a B_R -only one.*

Claim 2.2 (C2 – the derivation is invalid). *The manuscript’s derivation of estimate (10.19) contains three claims, each false for every admissible choice of cutoff (§4.3).*

Claim 2.3 (C3 – the estimate is false). *Estimate (10.19) is false under its own quantitative hypotheses: there exists a smooth, divergence-free, uniformly L^2 -bounded family satisfying every bound the derivation invokes, with $B_R = 0$ exactly, whose pressure flux is nonzero (§5).*

Claim 2.4 (C4 – the failure is localized). *The other flagged divergence – the (8.19) torus-inverse estimate – is sound as written: formalized in Lean, the kernel accepted it (§6). The paper’s conclusion is not thereby refuted; its proof as written is (§8).*

3 What was claimed, and what the verification covered

3.1 The theorem

manuscript p. 1

#2

Theorem 1.1. For every $\nu > 0$ there exist a force $f \in C_c^\infty(\mathbb{R}^3 \times (0, \infty); \mathbb{R}^3)$, a compact set $K \subset \mathbb{R}^3$, and smooth velocity and pressure fields u, p on $\mathbb{R}^3 \times [0, 1)$ satisfying

$$\partial_t u + (u \cdot \nabla)u - \nu \Delta u + \nabla p = f, \quad \nabla \cdot u = 0, \quad u(\cdot, 0) = 0, \quad (1.1)$$

such that $\text{supp } u(\cdot, t) \cup \text{supp } p(\cdot, t) \subset K$ for every $0 \leq t < 1$,

$$\sup_{0 \leq t < 1} \|u(t)\|_{L^2(\mathbb{R}^3)} < \infty, \quad \limsup_{t \uparrow 1} \|u(t)\|_{L^\infty(\mathbb{R}^3)} = \infty.$$

The proof of the non-existence direction – the part that closes alternative (C) – runs through §10.3: a comparison lemma (Lemma 10.5) showing that any smooth solution v with the same force, zero datum, and bounded energy must agree with the constructed solution u on every $[0, T]$, $T < 1$, and hence on all of $[0, 1)$ – contradicting the blowup (10.21). Lemma 10.5 consumes the pressure-flux estimate (10.19).

manuscript p. 1

#2

Consequently, there is no smooth solution (u, P) on $\mathbb{R}^3 \times [0, \infty)$ with the same force and initial datum whose kinetic energy is uniformly bounded $\sup_{t \geq 0} \frac{1}{2} \int_{\mathbb{R}^3} |u(x, t)|^2 dx < \infty$. This establishes alternative (C) in the Millennium problem statement for Navier–Stokes as stated by Fefferman in [13].

3.2 What the verification job actually verifies

The Lean fork was verified by an independent FalseGreen job (31974058-a665-472d-b6f8-00e4ec9e2d45, production stack, real job image, isolated container; signed report, schema 3; all six contract criteria passed in one attempt, 31.4minutes wall):

1. **build** – the entire fork compiles under Lean 4.34.1 (the prebuilt `.lake` oleans are *replayed*: the Lean kernel re-verifies every olean against its source; zero recompiles);
2. **repaired_estimate** – the repaired estimates of §8 elaborate from source under the Lean kernel;
3. **comparator_ns** – the Comparator independent reference check of `NavierStokes.json`: both Navier–Stokes theorems re-checked by the default Lean kernel *and* by the independent nanoda kernel [7, 8], axioms within the permitted baseline;
4. **comparator_euler** – the same for `Euler.json`;
5. **result_audit_ns** – zero `sorry` in `NavierStokes/`; `#print axioms` of the main theorem within `{propext, Classical.choice, Quot.sound}`;
6. **result_audit_euler** – the same for the Euler full theorem.

Verdict: **accepted**. What this establishes: the Lean statements, as written in the fork, are valid Lean theorems – no gaps, no axioms beyond the permitted baseline, re-checked by a second kernel. What it does *not* establish: that those Lean statements are the argument on the page. That is the alignment question, and it has two independent answers, both adverse.

3.3 Component verdicts

Component of the paper’s argument	Verdict	Evidence
Theorem 1.1 statements (NS/Euler)	faithful in Lean	alignment audit
§8.19, the $m+4$ torus-inverse estimate	SOUND – formalized, kernel-verified	fork 9bbfa7b , §6
§10.3 flux-kernel estimates (K_R , commutator, (10.17))	correct (re-derived)	§4
§10.3 pressure-flux estimate (10.19)	INVALID AS WRITTEN	§4, §5
Non-existence clause (Lemma 10.5 $\rightarrow$ contradiction)	NOT ESTABLISHED – consumes (10.19)	§4.4
Lean formalization (the verified object)	proves a different argument	§7

4 The defect at (10.19): support bookkeeping

4.1 The paper’s setup

The comparison lemma works with the difference $w = v - u$ of two solutions and the pressure difference $\pi = P - p$. The pressure gradient is recovered from the equations via the Riesz transforms (10.16), and the remaining boundary term is the flux $\int \pi w \cdot \nabla \chi_R dx$ through an expanding ball. The paper’s stated purpose for the estimate:

manuscript p. 122

#2

We now bound the pressure contribution at the boundary of an expanding ball in terms of the weighted dissipation inside it.

The cutoff machinery:

manuscript p. 122

#2

We choose $0 \leq \phi \leq 1$ smooth, compactly supported and equal to one on the unit ball. For $R \geq 1$, we write $\phi_R(x) = \phi(x/R)$, $\chi_R = \phi_{8R}$, and set

$$E_R = \left(\int \chi_R |w|^2 \right)^{1/2}, \quad A_R = \left(\int \chi_R |\nabla w|^2 \right)^{1/2}, \quad B_R = \|\phi_{4R} w\|_6.$$

The Sobolev inequality and the product rule imply

$$B_R \leq C \|\nabla(\phi_{4R} w)\|_2 \leq C(A_R + R^{-1} \|w\|_2). \quad (10.17)$$


So the quantities on the table are: E_R and A_R , measured on the ball $\chi_R = \phi_{8R}$ (the “inside”); B_R , an L^6 norm measured on the *smaller* ball $\phi_{4R} \subset B_{4cR}$ where $\text{supp } \phi \subset B_c$, $c > 1$; and the flux, which we now locate.

4.2 The geometric fact

Since $\phi \equiv 1$ on B_1 (so $\phi' = 0$ on B_1) and $\text{supp } \phi \subset B_c$ for some $c > 1$:

$$\nabla \chi_R(x) = \frac{1}{8R} \phi'(x/8R), \quad \text{supp } \nabla \chi_R \subset \{8R \leq |x| \leq 8cR\}.$$

The flux $\int \pi w \cdot \nabla \chi_R dx$ is supported on the **annulus** $[8R, 8cR]$ – outside B_{8R} , exactly “the boundary of the expanding ball.” Every cutoff that appears on the proposed right-hand side of (10.19) is either supported strictly inside the inner edge of that annulus or fails to cover it:



- $\text{supp } \phi_{4R} \subset B_{4cR}$ – for $c < 2$ strictly inside B_{8R} , never covering the annulus;
- $\text{supp } \phi_{3R} \subset B_{3cR} \subset B_{8R}$;
- $\phi_{7R} \equiv 0$ on $\{|x| \geq 7cR\}$, and $7cR < 8cR$.

The quantity the estimate promises to control (the flux) and the quantities it proposes to control it with $(B_R, \|w\|_2)$ live on disjoint regions.

4.3 The three false claims

The paper’s derivation of (10.19), verbatim from p. 123:

manuscript p. 123	#2
Since $\nabla \chi_R \leq CR^{-1} \phi_{7R}$, the two terms in (10.18) pair with $CR^{-1} \phi_{3R} w$. Interpolation yields $\ \phi_{3R}w\ _3 \leq \ \phi_{2R}w\ _3 \leq B_R^{1/2} \ w\ _2^{1/2}, \quad \ \phi_{3R}w\ _4 \leq B_R^{3/4} \ w\ _2^{1/4}.$ Consequently $\left \int \pi w \cdot \nabla \chi_R dx \right \leq C_T R^{-1} (B_R + 1) B_R^{1/2} + R^{-3/4} B_R^{3/4}. \quad (10.19)$	

The two inputs to the pairing are themselves correct, and the paper says so in the right places – the $L^{3/2}$ bound on the first sum of (10.18):

manuscript p. 122	#2
The first sum has $L^{3/2}$ norm at most $C_T(B_R + 1)$, since $\ \phi_{4R}w_iw_j\ _{3/2} \leq B_R \ w\ _2, \quad \ \phi_{4R}w_iu_j\ _{3/2} \leq \ w\ _2 \ u\ _6.$	

and the commutator bound:

Young's convolution inequality bounds the second sum in (10.18) in $L^{4/3}$ by $C_T R^{-3/4}$.

What is false is the pairing step, which contains three claims:

Claim 4.1 (Defect 1 – “ $|\nabla\chi_R| \leq CR^{-1}\phi_{7R}$ ” is false for every admissible ϕ). *The shell $(7cR, 8cR]$ is nonempty ($c > 1$). On it, $\phi_{7R} = 0$ while $\nabla\chi_R = \frac{1}{8R}\phi'(x/8R)$ can be nonzero – and for the standard cutoffs it is, since $\phi' \neq 0$ somewhere on $(1, c)$. No choice of ϕ with $0 \leq \phi \leq 1$, smooth, $\equiv 1$ on B_1 , compactly supported, makes the inequality hold: the transition region of $\nabla\chi_R$ is $[8R, 8cR]$ and the transition region of ϕ_{7R} ends at $7cR < 8cR$.*

Claim 4.2 (Defect 2 – the pairing “with $CR^{-1}\phi_{3R}|w|$ ” is impossible). *$\text{supp } \phi_{3R} \subset B_{3cR} \subset B_{8R}$, while the integrand $\pi w \cdot \nabla\chi_R$ is supported on $[8R, 8cR] \subset (B_{8R})^c$. On the entire support of the integrand, $\phi_{3R}|w| \equiv 0$. A function that vanishes where the integrand lives cannot pair with it; whatever the norm interpolation, the bound cannot be what is claimed.*

Claim 4.3 (Defect 3 – “ $\|\phi_{3R}w\|_3 \leq \|\phi_{2R}w\|_3$ ” is reversed). *For the standard radially nonincreasing cutoff, $\phi(x/3R) \geq \phi(x/2R)$ pointwise, hence $\|\phi_{3R}w\|_3 \geq \|\phi_{2R}w\|_3$ – the opposite of the claimed direction. (The second link, $\|\phi_{2R}w\|_3 \leq B_R^{1/2}\|w\|_2^{1/2}$, does hold: $\|\phi_{2R}w\|_3 \leq \|\phi_{4R}w\|_3 \leq \|\phi_{4R}w\|_6^{1/2}\|\phi_{4R}w\|_2^{1/2} = B_R^{1/2}\|w\|_2^{1/2}$. Only the first link is broken.)*

Each defect holds for *every* admissible cutoff: the argument is not saved by a more careful choice of ϕ .

4.4 Consequence

The right-hand side of (10.19) is controlled by quantities supported strictly inside B_{8cR} (B_R on B_{4cR} ; the global $\|w\|_2$), while the integrand of the flux lives on the annulus $[8R, 8cR]$. The chain “Consequently (10.19)” does not go through; the difference-energy identity is not closed. The paper continues:

Using (10.17) in (10.19), every power of A_R in these flux bounds is at most $3/2 < 2$. Young's inequality therefore gives

$$\frac{1}{2}E'_R + \frac{1}{2}A_R^2 \leq \|\nabla u\|_\infty E_R + \frac{C_T}{2R}$$

for almost every $t \in (0, T)$.

That closure – and with it the Gronwall argument of Lemma 10.5 ($w \equiv 0$ on $[0, 1]$), contradicting the blowup (10.21)) – is not established. **The non-existence clause – and with it the published proof of Theorem 1.1 – is broken at (10.19).**

5 The counterexample: (10.19) is false under its own hypotheses

A broken derivation is an unproved estimate; a stronger statement is that the estimate itself is false. The derivation of (10.19) invokes, quantitatively, only: $\|w\|_2 \leq C_T$; the $L^{3/2}$ bound of the first sum in (10.18) (via $B_R\|w\|_2$ and $\|w\|_2\|u\|_6$); the uniform commutator bound $\|K_R\|_{4/3}^{4/3} \leq CR^{-1}$ (an identity, §6.2); and the three false claims of §4.3. The estimate (10.19) claims the flux is controlled by B_R and $\|w\|_2$ *alone*. It is not.

Claim 5.1 (Counterexample to (10.19)). *For every $\varepsilon > 0$ there exist smooth divergence-free fields u and $v_\varepsilon = u + \varepsilon\psi$, with $\pi_\varepsilon = \sum_{i,j} R_i R_j g_{ij,\varepsilon}$ (the paper’s (10.16), $g_\varepsilon = v_\varepsilon v_\varepsilon - uu$), such that all quantitative hypotheses of the derivation hold ($\|w_\varepsilon\|_2 = O(\varepsilon)$, $\sum \text{tr} \|g_\varepsilon\|_1 = O(\varepsilon)$, $B_{R,\varepsilon} = 0$) and yet*

$$\left| \int \pi_\varepsilon w_\varepsilon \cdot \nabla \chi_R dx \right| = 1.6529 \times 10^{-2} \varepsilon^2 + O(\varepsilon^3) > 0 = C_T R^{-1} (B_R + 1) B_R^{1/2} + R^{-3/4} B_R^{3/4}$$

for every constant C_T : with $B_R = 0$ the right-hand side of (10.19) is exactly 0.

Construction. Take $R = 1$, $c = 2$ (i.e. $\text{supp } \phi \subset B_2$, so $B_R = \|\phi_4 w\|_6$ sees B_8 , and $\chi_1 = \phi_8$ has $\nabla \chi_1$ supported on the annulus $(8, 16)$):

- $\psi = \text{curl}^\perp f$, with $f(r, \theta) = H(r) (1 + 0.7 \cos 2\theta + 0.5 \sin 3\theta)$ and H a smooth bump with $\text{supp } H \subset (8, 16)$;
- $u = \nabla \times (f_0 e_z + Q e_y)$, with f_0 a mixed-mode bump on $(4, 12)$ and Q a z -annular bump – smooth, divergence-free, $O(1)$;
- $w_\varepsilon = \varepsilon\psi$, $v_\varepsilon = u + \varepsilon\psi$, $g_\varepsilon = v_\varepsilon v_\varepsilon - uu$, $\pi_\varepsilon = \sum_{i,j} R_i R_j (g_\varepsilon)_{ij}$.

Since $\text{supp } \psi \subset (8, 16)$ and $\phi_4 \equiv 0$ for $|x| \geq 8$:

$$B_{R,\varepsilon} = \|\phi_4 w_\varepsilon\|_6 = 0 \text{ exactly.}$$

Numerical verification. FFT-based 3D Riesz transforms on a 64^3 grid in $[-16, 16]^3$, spectral derivatives (script in §A):

quantity	value
flux($\varepsilon = 1.0$)	-1.652947e-02
flux($\varepsilon = 0.5$)	-4.132369e-03 ratio $F(0.5)/F(1) = +0.2500$ ($= \varepsilon^2$, exact to 4 digits)
flux($\varepsilon = 0.1$)	-1.652947e-04 ratio $F(0.1)/F(1) = +0.0100$ ($= \varepsilon^2$, exact to 4 digits)
$\sup \psi / \sup u $	3.1854 / 1.3586
B_R	0 $\Rightarrow$ RHS of (10.19) = 0 for all C_T

The exact ε^2 scaling (the flux is bilinear in w_ε through the cross term $u \otimes \psi + \psi \otimes u$) confirms the flux is nonzero and structural – not a discretization artifact. The numbers above are not hand-typed: they are regenerated by the script and the paper test harness fails if they drift.

Why this is decisive. (10.19) is stated as a *consequence* of the bounds above. The family $(w_\varepsilon, \pi_\varepsilon)$ satisfies every quantitative hypothesis the derivation uses and violates the conclusion for all $\varepsilon > 0$. Hence the step “Consequently (10.19)” does not follow from the hypotheses: the derivation is invalid, and the estimate is false, not merely unproved.

Scope caveat (stated plainly). In the paper’s argument $w = v - u$ is the difference of two *solutions* of (1.1); the counterexample uses general smooth divergence-free fields, showing that (10.19) is not a consequence of the quantitative hypotheses the derivation invokes. Whether a B_R -only flux bound holds for solution-differences is an open question – but the published derivation does not establish it, and the Lean code (§7) does not prove it either.

6 What survives: the verified-sound parts

The failure is *localized*. Each standard ingredient of the surrounding argument was re-derived and, where it is a pure estimate, formalized:

6.1 The (8.19) torus-inverse estimate is sound as written

The other divergence flagged in the alignment audit concerns the auxiliary-torus corrections of §8. The paper’s estimate:

manuscript p. 95

#2

Lemma 8.6. Let $N = v_t \cdot \partial y$ on T^2 , with v_t as in (6.2), and let F be a smooth function of zero Haar mean. Then $N\phi = F$ has a unique smooth zero-mean solution $\phi = N^{-1}F$. This solution is given by the following Fourier series and satisfies the displayed norm bound for every integer $m \geq 0$:

$$N^{-1}F = \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} \frac{\hat{F}(k)}{2\pi i v_t \cdot k} e^{2\pi i k \cdot y}, \quad \|N^{-1}F\|_{C^m} \leq C_m \|F\|_{C^{m+4}}. \quad (8.19)$$

manuscript p. 96

#2

The divisor bound (6.7) gives $|v_t \cdot k|^{-1} \leq C(1 + |k|)$. Four more derivatives of F than the requested output leave a summable $(1 + |k|)^{-3}$ bound on the Fourier series in two dimensions. This proves (8.19), smoothness, and uniqueness after fixing the zero Fourier coefficient.

We formalized this argument in Lean 4.34.1 / mathlib (fork commit 9bbfa7b, +159 lines, three files), and the kernel accepted every step:

Lean declaration	content
<code>SmoothFourierData.lean:+93</code>	$\sum (1+n)^{-3/2} < \infty$; the symmetric split $(1+a+b)^{-3} \leq (1+a)^{-3/2}(1+b)^{-3/2}$ from
<code>ParametricTorusInverse.lean:+16</code>	the lattice constant for the $m+4$ weight
<code>SmoothFamilyTorusInverse.lean:+50</code>	(8.19) itself: an r -torus-derivative word of N^{-1} is bounded by $r+4$ input derivat

The Lean kernel accepted the paper’s (8.19) argument as written: it is sound. (The Lean *code* in the OpenAI repository proves a *stronger* bound – $m+5$ with a $(\text{weight})^{-4}$ majorant – which is why a green job verifying the code says nothing about the page; §7.)

6.2 The flux-kernel structure is correct

The commutator kernel $K_R(x - y) = C|x - y|^{-3} \min(|x - y|/R, 1)$ and its radial integration are right:

$$\|K_R\|_{4/3}^{4/3} \leq C \left[R^{-4/3} \int_0^R r^{-2/3} dr + \int_R^\infty r^{-2} dr \right] \leq CR^{-1},$$

and Young’s convolution inequality then gives the $L^{4/3}$ commutator bound $C_T R^{-3/4}$ – matching the manuscript’s own claim (§4.3, quote **q-young**). Likewise (10.17) is correct (Sobolev + product rule, quote **q-setup**). The paper’s failure is not in the summability arguments or the kernel estimates; it is in the support bookkeeping of the flux pairing.

7 The Lean code proves a different argument

The object the verification job actually verified is the Lean code. Its pressure-flux result is:

`exists_uniform_actual_pressure_flux_bound` `NavierStokes/R3/PressureFlux.lean:584`

$$\left| \int \pi w \cdot \nabla \chi_R \right| \leq C_T (B_R^{1/2} + 1) \left(\frac{A_R}{R} + \frac{1}{R^2} \right) + R^{-7/4} B_R^{3/4}, \quad \forall t.$$

An A_R -**dependent** bound. Its proof (`WholeSpaceUniqueness.lean:60`) carries the A_R terms through; they are not dropped. Two facts make the direction of the divergence unambiguous:

1. **(10.17) shows the dependency arrow.** $B_R \leq C(A_R + R^{-1}\|w\|_2)$: B_R is controlled *by* A_R . The other direction – the one a B_R -only flux bound would need – does not exist. The flux, which lives on the annulus where B_R is blind, must be controlled by the *interior dissipation* A_R : exactly the type of bound the Lean code proves.
2. **The Lean bound is “the weighted dissipation inside it” the paper’s own words promise** (quote `q-purpose`): A_R is that dissipation; B_R is not.

Interpretation, stated as reading rather than claim about intent: the Lean code implements the correct-type, consistently-scaled version of the flux argument; the manuscript’s B_R -only (10.19) is a different – and false – statement. A fully green verdict therefore verifies the code, not the paper’s argument. The same phenomenon at the torus inverse (Lean $m+5$ vs. page $m+4$) is the alignment audit’s other divergence; there the page side is the *sounder* one.

8 What a repair takes

The minimal repair of (10.19) is a consistent radius hierarchy: e.g. $\chi_R = \phi_{32R}$, $B_R = \|\phi_{16R} w\|_6$, pairing cutoffs covering the annulus $[32R, 32cR]$, and a flux bound of the Lean code’s type – $(B_R^{1/2} + 1)(A_R/R + 1/R^2) + R^{-7/4} B_R^{3/4}$ – with the A_R terms absorbed into the A_R^2 dissipation in the difference-energy identity. All ingredients are standard and individually verified correct (§6).

Status: the repair is formalized and kernel-verified. Fork commit 24b1944 (`NavierStokes/R3/Paper101`) proves both corrected estimates in Lean:

- **(10.19)’ – `exists_uniform_flux_bound_corrected_radii`:** the fork’s own `exists_uniform_actual_press` (identical hypotheses, identical proof obligation) with B_R replaced by $\tilde{B}_R = \|(\text{cutoff at } 2R)^4 w\|_6$ and A_R by $\tilde{A}_R = (\int (\text{cutoff at } 2R)^8 |\nabla w|^2)^{1/2}$ – the L^6/L^2 norms measured on the ball that *contains* the flux annulus. The $R^{-1}/R^{-7/4}$ structure is unchanged.
- **(10.17)’ – `covering_sobolev`:** $\tilde{B}_R \leq C(\tilde{A}_R + R^{-1}\|w\|_2)$.

What remains open (stated honestly): whether the corrected flux bound closes the difference-energy inequality and the full Gronwall/ $R \rightarrow \infty$ chain in the paper’s structure. The closure uses the fork’s `eq_of_pressure_flux_bound` machinery with (10.17)’ substituted; it is type-level consistent but is a separate formalization step.

9 Consequences

1. **The published proof of Theorem 1.1 is broken at (10.19).** The non-existence clause (Lemma 10.5 $\rightarrow v = u$ on $[0, 1) \rightarrow$ contradiction) consumes (10.19); (10.19) is not a valid estimate as written, and its proof as written does not go through.
2. **The theorem’s conclusion is not thereby refuted.** The defect is repairable (a radius rescaling; the repaired estimates are formalized), so the theorem may still be true. What is refuted is the *paper as written*: its central estimate is false under its own hypotheses, and its “verified” Lean formalization proves a different argument.
3. **For verification practice:** this is precisely the gap a code-verifier cannot see. Independent verification jobs verify the Lean code; the paper’s argument is a natural-language claim about the mathematics. Aligning paper $\leftrightarrow$ code, and formalizing the paper’s *actual* argument step by step (the kernel as referee), is the only way to verify the paper itself – and it is what found the defect here. The companion papers [4–6] develop the mapping, the walkthrough, and the technical detail.

10 Related work

Bastounis, Circelli, and Hansen [3] study the same OpenAI manuscript and document the same phenomenon at a different level: “lost in translation” between the natural-language proof and its autoformalization, with worked examples of semantic drift in the Lean code. Their examples show that the code can quietly prove something the page did not claim. This paper complements that work in two ways. First, we identify where the *page itself* fails – not merely where the translation drifts: (10.19) is invalid as a natural-language argument, and is false under its own quantitative hypotheses, with an explicit counterexample family. Second, we formalize the sound parts ((8.19), and the repaired flux estimates) so that the kernel, not the reader, decides which steps hold. The division is: they characterize the translation gap; we close it for the flux argument, in both directions – showing the page’s bound is false and giving the bound the code proves.

The precedent for publishing a counterexample to a claimed Navier–Stokes result as a standalone short paper is long-established [2]. On the mechanics of what “verified” may and may not claim for a formalized manuscript, the Lean ecosystem’s own tooling – kernel re-verification of compiled artifacts, dual-kernel reference checking [7, 8], axiom audits – is summarized in the verification report (Appendix B).

References

- [1] Clay Mathematics Institute. The Millennium prize problems: Navier–Stokes existence and smoothness. URL <https://www.claymath.org/millennium-problems/navier-stokes-equations>.
- [2] R. C. Hallock. A counterexample to a proposed proof of the navier–stokes existence and smoothness conjecture. Precedent: counterexamples to claimed NS results publish as their own papers.
- [3] Hansen et al. Navier–stokes lost in translation. arXiv:2610.08144, 2026. Semantic faithfulness of AI autoformalisation; worked examples on the OpenAI NS proof.

- [4] Bevan Hunt. The proof as the machine checked it: The verified Lean formalization of the navier–stokes blowup argument. Companion paper, this series, 2026. DOI: PENDING.
- [5] Bevan Hunt. The navier–stokes proof, step by step: What OpenAI claimed, what the machines checked, and where the argument breaks. Companion paper, this series, 2026. DOI: PENDING.
- [6] Bevan Hunt. Inside the verified proof: A step-by-step walkthrough of the Lean navier–stokes formalization. Companion paper, this series, 2026. DOI: PENDING.
- [7] Sander R. J. J. Janssen et al. nanoda: a standalone Lean kernel. URL <https://github.com/ammkrn/nanoda>.
- [8] OpenAI. Comparator: independent reference checking for Lean proofs (landrun + nanoda). Dual-kernel reference checks: default Lean kernel + independent nanoda kernel, under the landrun sandbox.
- [9] OpenAI. Finite time blowup for navier–stokes. <https://cdn.openai.com/pdf/32d9f210-8b73-45e0-91bc-82a30aef8a9a/navier-stokes.pdf>, 2026. Manuscript with Lean formalization; Lean fork: NavierStokesAndEuler.

A The counterexample, reproducible

The script is `counterexample/counterexample-1019.py` in the ns-proof repository (<https://github.com/truemagic-coder/ns-proof>); dependencies: `numpy` only. It builds the fields of §5 on a 64^3 grid in $[-16, 16]^3$, computes the Riesz transforms by FFT ($\widehat{R_i f}(\xi) = i\xi_i/|\xi| \hat{f}(\xi)$), and evaluates the flux $\int \pi_\varepsilon w_\varepsilon \cdot \nabla \chi_1 dx$ at $\varepsilon \in \{1, 0.5, 0.1\}$, printing the ε^2 ratios. The paper’s numbers are regenerated from its output (`counterexample/gen_numbers.py`) – the test harness regenerates and diffs, so a hand-typed number cannot survive a commit.

The exact- ε^2 scaling in the output is structural, not fitted: π_ε is bilinear in $(u, \varepsilon\psi)$ through $g_\varepsilon = 2\varepsilon u \otimes \psi + \varepsilon^2 \psi \otimes \psi$, and the flux pairs π_ε with $w_\varepsilon = \varepsilon\psi$; the leading term is the ε^2 cross term.

B The verification report

Job 31974058-a665-472d-b6f8-00e4ec9e2d45 (FalseGreen, production stack; the full fork at commit 24b1944; the real Lean job image; isolated container; no network beyond the approved build hosts). Verdict **accepted**; signed report (schema 3) with signature 9f1834f2...; single attempt; 31.4minutes wall. The six contract criteria are listed in §3. The report embeds the frozen `falsegreen.toml` contract (verified against the job’s bound `contract_sha256` and `criteria_sha256`) and, per evidence row, the typed `JobResult` (per-criterion pass/fail, per-command argv/exit/duration). The report is self-contained: its `evidence_sha256` commitments can be re-derived from the embedded rows, and its signature – an HMAC-SHA256 over the canonical report digest, keyed with the deployment’s report key – can be re-verified offline (`falsegreen-canary verify-report`) without contacting the platform.

The verification pipeline’s independence layers: (i) the *build* criterion re-runs `lake build` in an isolated container – the Lean kernel re-verifies every prebuilt olean against its source (olean replay), so a corrupted or backdoored `.lake` cannot pass; (ii) the *comparator* criteria re-check the target theorems with *two* independent kernels (the default Lean kernel and nanoda) under the landrun sandbox [7, 8]; (iii) the *result-audit* criteria count `sorry` and print the axioms actually used. Nothing in this pipeline checks the *page*; that is §4.