

THE HAND IN THE WATER

OpenAI opened door C.
The world reported door A.

Just before midnight on Monday, a mathematician at New York University hit post.

Tristan Buckmaster had spent months on a problem worth a million dollars. He had used OpenAI's coding tool the whole time. On Saturday, OpenAI announced it had finished the job.

What he wrote that night is the second most interesting thing about this story.

The first is that almost nobody who reported it told you which problem was solved.

The Navier–Stokes equations describe how fluids move. They are the reason a weather model works at all, and the reason turbulence is still the last open wound in classical physics. In 2000 the Clay Mathematics Institute put a million dollars on one question about them and made it one of seven Millennium Prize Problems.

Ninety years ago, Jean Leray proved something unsettling. He showed that solutions to these equations always exist in a weak sense, and then admitted he could not rule out that they tear themselves apart. Nobody has closed that gap since.

On Tuesday the world was told a machine closed it in eighty-eight hours.

That is not what happened. What happened is more interesting, and it turns entirely on a single word that no headline had room for.

THE CLAIM

OpenAI, 8 Sep 2026: an internal system produced a proof that Navier–Stokes dynamics **can develop a singularity in finite time**.

THE PAPER

166 pages. Title: Finite Time Blowup for Navier–Stokes. Author line reads only **OpenAI**. No academic co-authors.

THE PRIZE

OpenAI says it does **not intend to claim** the Millennium Prize. Clay has not commented and still lists the problem as unsolved.

THE AGE

Headlines say a 90-year-old problem. The equations are roughly **200** years old. The 90 is Leray, 1934.

02 FOUR DOORS

Charles Fefferman wrote the official rules in 2000. Most people think the problem asks one question. It asks four, and you only have to answer one of them to be paid.

$f = 0$ · NOBODY TOUCHES THE WATER



OPEN SPACE STAYS
SMOOTH FOREVER

A LOOPED BOX STAYS
SMOOTH FOREVER

$f \neq 0$ · A HAND IS ALLOWED



OPEN SPACE CAN BE
TORN APART

A LOOPED BOX CAN BE
TORN APART

OPENAI WENT THROUGH THIS ONE

CLAY MILLENNIUM PRIZE · NAVIER-STOKES · THE FOUR ADMISSIBLE ANSWERS

Doors A and B are the ones everybody pictures. They say water left completely alone stays smooth forever, and they ask you to prove it. Nothing pushes. Nothing stirs.

Doors C and D are different. They let you apply an outside force to the fluid and ask only that you make something break. Build the push, break the water, collect the money.

You never had to answer the question everyone is asking. You only had to answer one of the four.

OpenAI opened the third door, the one marked in red above. That is a legitimate answer to a real Millennium question and it deserves the attention it is getting. It is also not the question anyone outside the field thinks was being asked, and nobody reading a headline could reasonably be expected to know the difference.

03 THE HAND IN THE WATER

A forcing term is a hand.

It is an outside push applied to the fluid at every point, smooth, and in this case switched off outside a small patch of space and a short window of time. Stir a cup of coffee and your spoon is a forcing term.

OpenAI's theorem builds one very particular push. Under that push, in finite time, the speed of the fluid runs away to infinity while its total energy stays bounded. The water tears. The hand is what tore it.

Door A asks whether water tears itself. Door C asks whether water can be torn.

That is the whole story, and it is the one thing a headline cannot fit.

DOES IT COUNT

Yes. Fefferman's statement explicitly permits a smooth force for C and D. OpenAI's force is **compactly supported**, which exceeds what the rules demand.

DOES IT SETTLE THE FAMOUS ONE

No. A and B require **no force at all**. They are untouched.

SAID BETTER

Scientific American: the Clay problem as written is solved, but the Clay problem as most experts imagine it lacks the very piece the method depends on.


WHY NINETY YEARS OF ATTEMPTS BROKE HERE

In two dimensions, water behaves. In three, it might not, and the difference has a name. A spinning tube of fluid in three dimensions can be pulled thin. Pull it thin and it spins faster, the way a skater pulls in her arms. Faster spin pulls harder. It feeds itself. Two dimensions forbid the move entirely.

In 2014 Terence Tao built a fake Navier–Stokes, close enough to the real one to share its energy budget, and proved the fake one blows up. The point was never the fake equation. The point was that any proof leaning on the energy budget alone is dead on arrival, because a doomed equation passes the identical test.

He called it the supercriticality barrier. It is the wall every serious attempt has hit.

TWO DIMENSIONS



TRAPPED FLAT.
CANNOT STRETCH.

THREE DIMENSIONS



PULLED THIN. SPINS FASTER.
IT FEEDS ITSELF.

88

hours of wall clock, start to finish

10,000

concurrent agents on this
one problem

341,000

lines of Lean 4 formalisation

2.7M

messages passed between
them

17

further hours to machine-
check it

130B

output tokens burned

~50

hours and 100 agents for the
Euler warm-up

166

pages of finished
manuscript

4

days between
announcement and rumour
that started it

The loudest number is not the interesting one. Before this run, roughly a hundred agents spent about fifty hours and took down the regularity problem for the Euler equations, which strips out viscosity entirely. Inside the building that was treated as a warm-up.

On cost, OpenAI said only **millions of dollars**. Every dollar figure you have seen since, from six million to forty, is an outsider multiplying token counts by public API prices. Some of those estimates price the whole multi-problem sprint, not this proof. Treat them as arithmetic, not disclosure.

05

WHAT A PROOF CHECKER CHECKS

Three hundred and forty one thousand lines of Lean 4 sit on GitHub. They compile. The machine agrees the argument is valid.

Here is the part every engineer already knows in their bones. A proof assistant returns exactly one bit. It tells you the theorem follows from the assumptions you wrote down. It does not tell you that the theorem you wrote down is the theorem you meant.

The tests pass. Nobody has published a review of the spec.

Somebody credible has to read the formal statement, line by line, and confirm it encodes Fefferman's door C and not a cousin of it. As of this writing no named expert in partial differential equations or in Lean has put their name to that confirmation in public.

This is not a reason to assume the proof is wrong. It is a reason to notice that the loudest evidence being waved around, the green compile, is evidence for a narrower claim than the one being celebrated.

Anthropic hit a version of the same wall four days earlier and handled it differently. Its model produced a machine-checked Fermat's Last Theorem, thirteen million lines of Lean, and Kevin Buzzard at Imperial College went on the record about what it did and did not establish. A named human vouched for the statement. That step is the whole ballgame and it has not happened here yet.

ONE BIT

Lean answers: **does this follow?** It never answers: **is this the right question?**

THE GAP

Formal verification moves trust from the proof to the **statement**. The statement is still read by people.

WHAT IS PUBLIC

Repository **openai/NavierStokesAndEuler**. What is not public is a named expert's audit of what it asserts.

NOT THE SAME THING

A green build is a **fact**. A solved Millennium Problem is a **judgment**, and Clay reserves it for humans.

06

TELEPHONE

Four caveats leave OpenAI's own page intact. The force is required. Nobody has verified it. Clay has not accepted it. There is a priority fight. Watch which ones survive the trip.

	THE FORCE	UNVERIFIED	NOT CLAY'S	THE FIGHT
OPENAI				
QUANTA				
SCIENTIFIC AMERICAN				
BBC				
CNBC				
TRADE AND VERTICAL PRESS				

FILLED = CARRIED · HOLLOW = NOT CARRIED · A READING OF THE COVERAGE, NOT A METRIC

Nobody lied. The BBC says plainly that the work is unverified and that Clay has not accepted it. The trade piece that pivoted to healthcare warns its readers against rebuilding anything around an unverified headline.

Look at the column that empties first. It is the one that needs a sentence of mathematics to explain. Every other caveat fits inside a clause, so every other caveat survived. That is not editorial bias. That is a word budget behaving exactly as a word budget behaves.

**Caveats do not die of malice.
They die of word count.**

Notice what grew to replace it. By the time the story reached the vertical press it had sprouted a section on healthcare research, which is speculation bolted to a result that changes nothing in any laboratory on earth.

In 2014 the Kazakh mathematician Mukhtarbay Otelbaev announced he had solved Navier–Stokes. Serious people took it seriously. Then he emailed Stephen Montgomery-Smith.

ALSO IN THE GROUND

Penny Smith, 2006. Claim withdrawn for an error.

SURVIVAL RATE

Announced Millennium solutions that went the distance: **one**. Poincaré. That is the entire list.

To my shame, on page 56 the inequality (6.34) is incorrect therefore the proposition 6.3 (p. 54) isn't proved. I am so sorry.

MUKHTARBAY OTELBAEV, 2014

Grigori Perelman posted his Poincaré proof in 2002. It was right. Three separate teams then spent years writing out the verification so the field could check it, and Clay paid in 2010.

Eight years, for a proof that was correct the day it appeared.

Clay's own rules require refereed publication, then a two year wait, then general acceptance by the mathematical community. Run that clock from a paper that has not been submitted anywhere.

The earliest this can be an accepted Millennium result is 2028, and only if it is right.

PERELMAN • POSTED TO PAID • 8 YEARS



VERIFICATION TAKES THE TIME IT TAKES

This is a sequence, not an accusation. Read it and decide for yourself what it means.

In March, OpenAI closed the largest private technology financing on record: a hundred and twenty two billion dollars at an eight hundred and fifty two billion valuation. In June it filed confidentially for an offering aimed at a trillion. In May, Anthropic passed it on paper at nine hundred and sixty five billion.

On the fourth of September, Anthropic announced a machine-checked proof of Fermat's Last Theorem. On the eighth, OpenAI announced Navier–Stokes.

There is a precedent for how this company handles a scoreboard. In July 2025 it announced a gold-medal score at the International Mathematical Olympiad, graded by a panel it assembled itself, published before the closing ceremony and reportedly against the organisers' wishes. Google DeepMind posted the same score and had it certified by the olympiad. Same number. Different epistemics.

Announce fast. Grade yourself. Let the correction arrive later, to fewer people.

MAR 2026

\$122B raised at **\$852B**. Largest private tech round ever closed.

MAY 2026

Anthropic valued at **\$965B**, ahead of OpenAI.

JUN 2026

Confidential S-1 filed. Reported target: **\$1T**.

4 SEP 2026

Anthropic ships Fermat's Last Theorem in Lean. **13M lines**.

8 SEP 2026

OpenAI announces Navier–Stokes.

None of this makes the mathematics wrong. Incentives do not refute theorems. They do tell you how hard to squint at a claim that arrived with a press call attached and no referee in sight, and how much weight to put on a superlative that has not yet been checked by anybody outside the building.

09 THE INPUT WAS A RUMOUR

Terence Tao has not assessed OpenAI's manuscript. He praised different work, by other people, and said clearly that it had not reached Navier–Stokes. If you saw his name attached to an endorsement this week, you saw a misreading.

What he did say is the part that should keep the builders awake.

STARTED WHEN

OpenAI says the effort began **1 September**, after hearing a rumour that two Millennium problems had been resolved.

THE DISPUTED PART

Buckmaster says he learned on 3 September that word of his progress reached OpenAI. OpenAI says it used no specific user data, while conceding it **cannot rule out** that de-identified usage data helped train the model.

UNRESOLVED

Both accounts are public. Neither has been established. Treat it as **contested**, not settled.

Even the rumor of someone working on a problem can trigger a massive amount of AI-powered effort to flatten it before the original research project has time to reach its full potential.

TERENCE TAO, SEPTEMBER 2026

Sit with what OpenAI volunteered about its own timeline. The trigger was not a paper. It was not a public preprint. It was a rumour that somebody was close.

Run that forward. If saying out loud what you are working on summons ten thousand agents to get there first, the rational move is to stop saying it. Open science runs on the exact opposite habit, and it has run on it for four hundred years.

The risk is not that machines do the mathematics. It is that people stop talking.

10 WHAT WOULD CHANGE MY MIND

Three things would move this from serious claim to settled result, in order of how much they should move you.

01 A named expert in partial differential equations or in Lean confirms in public that the formal statement encodes Fefferman's door C.

02 An outside group reproduces the check independently.

03 A refereed journal accepts it.

Two things would move it the other way. A specific error surfacing in the construction, which is exactly how Otelbaev ended, on page fifty six. Or a demonstration that the method cannot survive losing the hand. Stan Palasek at Princeton has already flagged that viscosity may simply eat the growth mechanism once the force is removed.

Silence on all five, six weeks from now, is itself information.

DOES THIS CHANGE YOUR FLIGHT

No. Real water is made of molecules. These singularities live inside the idealisation, not inside the ocean. Quanta put it plainly: **no immediate practical consequences.**


SO WHAT DID HAPPEN

A machine wrote 166 pages of frontier mathematics in under four days, formalised it, and published the code. **That part is not in dispute** and it is enormous.

So hold both. The capability is real and it is the fastest thing that has ever happened in this field. The result is unverified, narrower than reported, and contested by the people whose method it was built on.

Everything in the gap between those two sentences is a preposition.

WATER THAT TEARS ITSELF. WATER THAT CAN BE TORN.



**The hand in the water is the whole story.
It is the one thing the headline could not fit.**