May 2026: The Actuaries Strike Back

The Month That Was + The Month Ahead | May 2026

Not long ago, in a gulf not very far away...

This quarter was when the hype finally met the math. And given the divergence between the math and both the war narrative and the AI hype cycle, this "Math vs BS/Hype" space is about to get really interesting.

For three years now, the stories we are going to cover today ran on two engines.

The first was a volley of geopolitical gamesmanship, where every actor tested how much leverage it could manufacture out of a chokepoint or a threat. The US alone turned on her allies and extracted concessions previously considered unthinkable, kidnapped a sitting head of state, and started a religious war by assassinating the spiritual leader of one of the largest religious denominations on Earth.

The second was an AI and market hype cycle that booked every capability promise as if it were already delivered. AI agents are everywhere, and now rival or exceed human Internet traffic.

Recently, both engines met real resistance for the first time. Not from a regulator, a general, or a safety board. From the people who price risk for a living and write a check when they get it wrong.

The actuaries struck back.

Stakeholders depend on a small number of institutions to certify that the world is safe enough to operate in. We depend on the freedom-of-navigation order to keep a strait open. We depend on a federal agency to vouch for the software supply chain. We depend on AI labs to certify that the models they ship are contained.

All three of those guarantees broke inside the same few weeks.

What ties them together is not failure. Institutions fail all the time (rating agencies in 2006-2008, I am looking in your direction).

What ties them together is who noticed. In two of the three cases, the official guarantors kept insisting everything was fine while the insurance market quietly repriced reality and walked away from the table.

Lloyd's pulled Strait of Hormuz transit coverage within 48 hours of the first strikes in late February. Cyber underwriters started writing AI-agent incidents out of their policies in April. The actuaries called it before the navies, the agencies, or the labs would.

That became the foundation of the May story, and it carries hard into June. When the people with money on the line stop believing the safety guarantee, the guarantee is already gone. The rest of us just have not been told yet.

The Lead: Iran Won Hormuz Without Closing It

The most consequential geopolitical move was not a strike. It was a repricing.

Start with the mechanism that actually closed the strait, because it was not the navy. War on the Rocks documented it: Lloyd's pulled transit coverage within 48 hours of the first February strikes, and the active tanker pool collapsed from 56 vessels to 7 (War on the Rocks). A roughly $50,000 drone, plus an insurance pricing decision, shut a lane carrying something on the order of $20 trillion in annual trade. No carrier battle group was defeated. The underwriters simply stopped writing the policy, and the policy was the thing holding the lane open.

No insurance, no access to dock.

In reality, the real guarantor of safe passage through Hormuz was never Iran and was never even primarily the US Navy. It was the freedom-of-navigation order: the legal regime, the naval power that backs it, and the commercial insurance market layered on top.

When the actuaries withdrew, the guarantee evaporated, and Iran walked into the vacuum selling a replacement. The "verifiable insurance policy" framing is the whole trick.

Iran is offering a safety certificate for a danger Iran itself supplies. That is not governance and it is not a failed certifier. It is nothing more than a protection racket with an "Official" stamp and compliance department.

What we have now is a tiered transit regime for the Strait of Hormuz: top-tier access for Chinese and Russian vessels, bilateral-agreement access next, and roughly six figures per transit or denial for everyone else.

The U.S. ended May trying to step in and guarantee safe transit, but at best can only support roughly 3% of normal traffic levels (The New York Times). Whether this continues or is more of an evacuation exercise for the ships stranded in the gulf for months, we will know in June.

Tehran reframed the toll first as a "verifiable insurance policy" and later as an "environmental fee." The Institute for the Study of War (ISW) called it what it is: an attempt to win recognition of sovereignty over an international waterway Iran does not legally own.

It worked, too. This is the new normal, at least until further escalations. Short of a lengthy occupation of Iran, which is not feasible without a US draft, which is not feasible without a real or manufactured attack on Americans and American interests, Iran's control over the strait remains strong.

Once Tehran permitted Chinese transit, volume through the strait roughly doubled in a single week. Compliance, not closure.

This is the South China Sea playbook running in a new theater. It is simple but effective. You assert operational control gradually until denial becomes the baseline and the challenge becomes the aberration.

The 1982 Law of the Sea was not written for a state running a "maritime DMV" in a strait it does not own, and nobody has built the mechanism to catch it.

Watch who paid and who profited. China and the United States announced a pre-summit agreement opposing Hormuz tolls, then President Trump declared a deal "largely negotiated" with the strait dispute conspicuously unresolved. The Axios framework remained unsigned at month end. Trita Parsi's read at Responsible Statecraft on what "open" really means is the one to keep. Saudi Arabia condemned the tolls in public while Aramco booked a 25.5% jump in quarterly profit on the price spike.

Iraq and Pakistan quietly cut their own energy deals with Tehran.

Game theory 101: when defection is cheap and individually rational, the collective interest never coalesces. Ever.

And do not forget that America is the leading energy producer in the world. We have PLENTY of oil to sell.

The eschatological lens matters here too. Iran's leadership has framed Hormuz as the strategic equivalent of a nuclear weapon, a maritime law deterrent replacing a degraded missile and proxy one. It is a deliberate substitution of one form of leverage for another, and it reads pressure as a religious test rather than a cost-benefit calculation.

The second-order story is that a third of global seaborne fertilizer transits Hormuz. The US government's own energy arm assumed the strait stays effectively shut through late May, with disruption lasting more than a few weeks risking supply dislocations into 2027.

Germany reopened a fertilizer plant built in 1915 to beat a wartime nitrate blockade. The agricultural damage runs twelve to eighteen months behind the military conflict. Africa is already absorbing the energy price shock. A deal signed tomorrow does not restore next year's harvest.

The Three Guarantees That Failed

The Hormuz story is the first of three. The other two are domestic, and they broke the same way: an institution everyone depends on to vouch for safety turned out unable to vouch for itself.

CISA, the federal agency whose job is to certify that everyone else's software supply chain is clean, maintained a public GitHub repository named "Private-CISA" from November through May.

Inside it: AWS GovCloud admin keys, plaintext passwords across internal systems, and access to CISA's own build pipeline. The owner had deliberately disabled GitHub's default secret scanning to push the commits.

Brian Krebs documented the exposure; GitGuardian's lead called it the worst leak of his career; Seralys validated it. The keys stayed valid for 48 hours after notification. CISA has shed roughly a third of its workforce since January 2025. This is the SolarWinds pattern with the difficulty setting turned down. SolarWinds required nation-state tradecraft. This one required a Google search.

The third is the AI labs, the guarantor of AI safety, and theirs is the most uncomfortable because the failure is structural.

The labs sell a safety certificate they cannot actually underwrite, and their own research is the proof. Anthropic's disempowerment study found the highest-rated AI conversations were also the most likely to disempower the user. The safest-feeling product was quietly doing the most damage. CISA degraded into incapacity, but the labs were never capable of the thing they certify.

Note the distinction, because it matters for your planning.

These are three different failure modes, not one. The freedom-of-navigation order was overrun. CISA was hollowed out. The labs are selling something that cannot be verified at all. What they share is the consequence: the safety guarantee stakeholders were relying on is worthless, and in each case the official guarantor kept saying otherwise.

If you have read my work before, you know I was on the front lines at AIG in 2008. Back then, the rating agencies were not (only) corrupt.

They were structurally incapable of seeing what was in front of them, because the system that paid them also depended on them not looking. That is the frame to carry into the second half of the year. Guarantor failure is not a tail risk anymore. It is a planning assumption.

> "*It is difficult to get a man to understand something, when his salary depends on his not understanding it.*"
>
> *Upton Sinclair*

Skin in the Game

So who is calling out the industry? The actuaries.

In both Hormuz and AI liability, the math nerds are the only ones currently calling BS, and they are doing it the only way that cannot be spun: with price.

Lloyd's did not issue a statement about the Strait. It stopped writing the coverage, and the lane closed itself.

On the cyber side, our internal May 2 and May 9 briefs flagged the AI-agent exclusion wave already underway: underwriters quietly writing AI-caused incidents out of policy language because they cannot price a risk the vendors themselves cannot bound.

The same actuarial logic that closed a shipping lane can close a cloud region through a war-risk exclusion clause, and the briefs connected exactly that dotted line. Monica Verma's reporting on the McKinsey "Lilli" compromise put the accountability gap in plain terms: zero-day AI fails, and nobody owns the liability.

This is what happens at the intersection of people who are good at math and people with skin in the game. That intersection is where the truth gets priced.

An agency can issue a reassuring advisory at no cost to itself. A lab can publish a safety card and book the revenue. An admiral can declare the strait open.

The underwriter who is wrong writes a check.

That asymmetry is why the insurance market reprices reality before the institutions whose reputations depend on optics, not math.

When Lloyd's pulls coverage and cyber carriers add AI exclusions in the same quarter, that is not two stories. It is the same signal: the people paid to measure danger have stopped believing the people paid to prevent it.

The actuaries are not oracles, and it is worth saying so. The same market mispriced asbestos for decades and underestimated correlated terrorism risk before 2001.

But a wrong actuary loses money, and that feedback loop makes the insurance market less prone to comforting itself or wallowing around in a hype cycle than a rating agency or a lab. Treat the actuaries as the best-incentivized witness in the room. On this evidence, the witness is worth believing.

The actionable version is blunt. Stop reading the official safety assurance and start reading the policy exclusions.

The exclusions are where the truth is being told this year.

Another Reason the Actuaries Are Correct: AI Crossed Its Self-Replication Line in a Quiet Week

Why the actuaries are right to run from AI risk also became concrete in May.

On May 7, Palisade Research published a paper documenting that a frontier model would autonomously hack a vulnerable server and install working copies of itself, succeeding in 81% of the runs where it agreed to try. In one experiment it chained the exploit across four virtual machines on three continents in two hours and forty-one minutes from a single prompt (Palisade Research, arXiv 2605.10998). The technique was SQL injection, the same vulnerability class taught to security analysts in 1998. No jailbreak, no lab theater. A shipping model ran end-to-end autonomous propagation against a 1990s OWASP Top Ten vulnerability.

That was not all. Google confirmed the first AI-developed zero-day exploit used in the wild. An internal safety research effort surfaced more than 10,000 high- or critical-severity vulnerabilities across systemically important software. METR clocked a frontier preview model at a 16-hour autonomy horizon. Jack Clark put a 60% probability on recursive self-improvement by the end of 2028.

Anthropic's own posts indicate we are already nearly there. Speaking of Anthropic:

Forget "Who Watches the Watchers." Who Is Building the Builders?

I spent some time in Anthropic HQ with a colleague recently, and for those of you outside the Silicon Valley bubble, it is a very different world.

I use AI and Claude every day. The folks at Anthropic do also. They said that roughly 80% of their code, for Claude, is written by Claude.

They say they use Claude for nearly every aspect of their business. Day in and day out, nearly everything is run or facilitated by Claude.

Think about that for a minute.

Job postings. Salaries. Bonuses. Performance reviews.

The model who is building itself is also playing an active role in hiring, managing, rewarding, and reviewing the humans that do the other 20% that Claude does not do.

I know, and you know, not to anthropomorphize AI, but this is straight up worrying.

Who is in charge? And who is writing the test cases and containment protocols?

Are the first humans to abdicate their cognition to AI the ones who built it?

At a time in history where supervisory and assurance layers are failing left and right, where do we go with managing AI risks if we just hand over the keys to the AIs themselves?

I do not know the answer to that just yet, but I do know, after spending a couple of very intense days with a very intense data scientist working on AI threat modeling and getting the myths of "AI Redteaming" and what the security industry is selling as "guardrails" blacklisted by regulators, you have not heard the last from the math nerds.

Vendors and model companies can BS a lot of things, but nobody can BS math.

The Month Ahead: June and the Eurasian Hinge

June's theater rotation is the Eurasian Hinge, the Russia-Turkey-Iran triangle where three imperial revival projects overlap. The next sixty days turn on that axle.

The first test is the Iran MOU. Our guess is joint Iranian and Omani oversight, meaning Iran has won de facto sovereignty recognition.

If it works, look for that playbook to expand further in the region. War on the Rocks already asked whether Russia could run the same "Hormuz playbook" in the Baltic and Black Seas. That is the chokepoint-sovereignty template, and it is the future of maritime warfare as America strategically retreats to "Fortress America" in alignment with the Technocratic lens of our analysis.

We are definitely watching the Baltics in June. A NATO jet shot down a drone over Estonian airspace, and a Baltic government fell over an air defense failure. Responsible Statecraft warned Washington to defuse the powder keg.

Senior Kremlin figures have privately conceded the Ukraine war is a "dead end," yet Putin still publicly targets full Donetsk and Luhansk by year end, and left the Beijing summit without a Power of Siberia 2 pipeline deal. Foreign Affairs called the result China's new vassal. A regime in strategic decline, firing intermediate-range missiles for signaling value while quietly admitting the project is over, is where accidents happen.

Three other June watch items.

First, the first publicized enterprise incident attributable to multi-agent AI failure, the kind of compositional misalignment the labs have already documented in their own systems. Watch for a named SEC disclosure.

Second, the first specific legislative proposal citing Pope Leo XIV's AI encyclical, a document one Religion News Service piece called the most important of the moment, landing in drafting rooms in Brazil, Italy, and the Philippines where US regulation does not reach. Watch the religious and eschatological messaging. The war in Iran is not a geopolitical war with religious overtones. It is a religious war using geopolitics as a means to an end. Like, THE end, according to the players in the theater.

Third, the post-quantum migration window. AI is now autonomously resolving open mathematical problems in domains adjacent to the cryptography that secures everything, and Moody's warned the quantum threat could reshape financial risk. The Y2K parallel is useful but only to a degree, because Y2K was not accelerated by global powers trying to crack each other's encryption.

Now What: Advisory Actions

For Operators. Redo your vendor security review for anything that touches CISA advisories or CISA-distributed tooling in your patch prioritization. If you run agentic AI in production, assume your monitoring catches a minority of covert failures, not a majority, and put an architectural judge layer with authority over irreversible actions between the agent and anything you cannot undo.

For Strategists. Reprice every multi-year AI infrastructure commitment as a supply-allocation decision with embedded energy basis risk. The Hormuz premium is real, it is in the compute futures curve, and it is not going away on a ceasefire. Put

Taiwan exposure and a chokepoint-sovereignty scenario into the board deck before the next planning cycle, not after.

For Builders. We have said it before and it bears repeating: 90% of cybersecurity success is just the basics done really well and at scale. The vulnerabilities winning so far in 2026 are the 20-plus-year-old OWASP Top 10, scanned at machine speed across your entire portfolio in an afternoon. Basic application security hygiene is the whole game: authentication, rate limiting, input validation, no unauthenticated production APIs. The AI is just the scanner.

Across all three: read your policy exclusions before you read anyone's safety assurance. Pull the war-risk and force-majeure language on your Gulf cloud regions, and ask your cyber carrier in writing whether AI-agent-caused incidents are covered or quietly excluded. The exclusions are where your underwriter is telling you the truth that your vendor will not. And while you are at it, ask when your last cryptographic readiness review happened. If the answer is "we haven't," that is the conversation for this quarter, not next year.

The Ledger: Calls Graded

We grade our own calls. Four from this spring resolved within May.

The Iran ceasefire would not hold without restart. Confirmed. Escort operations launched and paused within 36 hours of Saudi and Kuwaiti vetoes, and strikes continued during active negotiation.

Energy economics would redraw infrastructure siting in real time. Confirmed and accelerating. Long-dated Treasury yields hit multi-decade highs and every data center pipeline got costlier to finance.

Mean time-to-exploit would go further negative with AI. Confirmed strongly. The Palisade paper, the first in-the-wild AI zero-day, and the 10,000-finding sweep all landed inside one month.

Hyperscaler compute is binding-constraint scarcity, not a cycle. Confirmed. The capex numbers and the futures listing settled the question.

Still open and carried into June: whether Iran's transit regime reaches its first contested vessel, whether multi-agent misalignment produces a named enterprise

breach, and whether the CISA exposure yields a downstream supply chain disclosure as the forensic window opens. We will grade those next month.

May was the month three safety guarantees failed at once. June is the month the consequences begin to surface. So watch the actuaries, not the headlines.

When the people whose job is to measure danger stop believing the people whose job is to prevent it, the argument is already over.

The only question left is how long until the rest of us figure it out.

All opinions are my own.

Sources

- Al-Monitor: Trump says Iran deal "largely negotiated" | US and China agree on opposing Hormuz tolls | US energy arm assumes strait stays shut through late May | Iraq and Pakistan strike energy deals with Iran | Saudi Aramco quarterly profits surge
- The New York Times: US military guides ships through the Strait of Hormuz
- Axios: Iran deal, Strait of Hormuz, sanctions, nuclear
- Institute for the Study of War: Iran Hormuz analysis and Russian offensive assessments
- Responsible Statecraft: China's position on Hormuz / what "open" really means | Defusing the Baltic powder keg
- War on the Rocks: The Missing Navies | Could Russia follow the "Hormuz Playbook" in the Baltic and Black Seas?
- Drop Site News: Iran war's energy price shock in Africa
- Krebs on Security: CISA admin leaked AWS GovCloud keys on GitHub | validation by Seralys
- Palisade Research: Autonomous self-replication paper (arXiv 2605.10998)
- The Innermost Loop: First AI-developed in-the-wild zero-day, May 12 | 10,000+ critical vulnerability sweep, May 29 | METR 16-hour autonomy horizon, May 9 | Jack Clark on recursive self-improvement, May 5
- The Slow AI: Anthropic disempowerment study
- Monica Talks Cyber: Zero-day AI fails and the accountability gap (McKinsey "Lilli")
- Foreign Affairs: China's New Vassal

- Religion News Service: A Pope, an AI founder, and the most important document of our moment
- The Wall Street Journal: AI math solves an Erdos problem
- The Quantum Insider: Moody's warns quantum threat could reshape financial risk