

The AI Market Panic Explained: Why Running Local Models Puts You on the Right Side of the Gap

February 27, 2026 · by Mark Bartlett

[Download this guide as PDF](#)

Quick Answer: The AI market panics of February 2026 — Citrini's viral '2028 Global Intelligence Crisis' memo, IBM's 13% single-day crash, \$2 trillion wiped from software stocks — are driven by a real fear: AI replaces white-collar jobs faster than the economy absorbs the shock. The doom case is plausible. The bull case (services deflation returning \$4,000-7,000 per household annually, 532,000 new business applications in January alone) is also plausible. What neither side talks about is the capability-dissipation gap — the distance between what AI can do right now and how fast that capability actually reaches the economy. That gap is wide and stays wide because of regulatory, organizational, cultural, and trust inertia. If you're already running local models, testing new releases, and building AI into your workflows, you're operating at the capability frontier while most of the economy moves at the dissipation rate. That advantage compounds.

Related: [VRAM Requirements for Local LLMs](#) · [OpenClaw Setup Guide](#) · [Best Local LLMs for Mac](#) · [Running LLMs on Mac M-Series](#)

On February 23, 2026, IBM stock dropped 13.2%. Its worst day in 26 years. Over \$31 billion in market cap gone. The cause: Anthropic published a blog post about COBOL modernization. Not a product launch. Not an earnings miss. A blog post. Claude Code can now map dependencies across thousands of lines of COBOL and document workflows that would take human analysts months. The market read that sentence and sold.

The day before, a speculative fiction piece from Citrini Research called "The 2028 Global Intelligence Crisis" went viral. Written as a fictional macro memo from June 2028, it modeled a scenario where AI-driven white-collar displacement cascades through the economy: layoffs lead to reduced spending, reduced spending triggers mortgage defaults among high-income earners, private credit collapses, and the S&P 500 drops to 3,500. A Substack post crashed the market. Bloomberg covered the author's surprise that it worked.

In the first three weeks of February, over \$2 trillion was wiped from software market caps. Microsoft, Nvidia, Oracle, Meta, Amazon, and Alphabet collectively lost \$1.35 trillion in a single week. Workday dropped 10% despite beating earnings. A logistics startup's demo crashed C.H.

Robinson 24%. Every few days, another AI capability announcement triggers another sell-off. The pattern keeps repeating.

I don't think either the doomers or the bulls have the full picture. There's a concept that explains what's actually happening, and it has practical implications for anyone running AI on their own hardware.

The doom case, taken seriously

The Citrini memo deserves honest engagement because it's well-constructed. James Van Geelen and Alap Shah weren't writing clickbait. They were modeling a specific, plausible failure mode.

The mechanism: AI gets better. Companies cut white-collar headcount because the math says to. Displaced workers spend less. Since white-collar workers make up roughly 50% of US employment and drive about 75% of discretionary spending, and the top 10% of earners account for over half of consumer spending, even modest displacement creates outsized hits to consumption. Mortgages written against stable high incomes start defaulting when those incomes disappear. Private credit, which grew from \$1 trillion in 2015 to over \$2.5 trillion by 2026, has been backing SaaS companies at valuations that assumed perpetual growth. When AI agents can do the work those SaaS products automate, the valuations collapse and the loans go bad.

The memo's most quoted line: "In 2008, the loans were bad on day one. In 2028, the loans were good on day one. The world just changed after."

They also introduced "Ghost GDP" — output that appears in national accounts but never circulates through paychecks and spending. AI produces more, but fewer humans touch the money on the way through. In their scenario, labor's share of GDP drops from 56% (2024) to 46%, the sharpest decline on record. Unemployment hits 10.2%. The S&P falls 38%.

This went viral because it has emotional coherence. Every white-collar worker who's watched Claude do their job in 30 seconds can feel the plausibility. That doesn't make it right, but dismissing it as "just fiction" misses why it moved markets.

The bull case (two arguments)

Argument 1: policy response is inevitable

The Citrini scenario requires a specific assumption that's easy to miss – no meaningful policy response as things deteriorate. But politicians respond to mass economic pain because they want to keep their jobs. The 2008 financial crisis proved this. A divided Congress that couldn't agree on anything still passed TARP because the alternative was unthinkable. When unemployment visibly spikes, governments act. They act badly, slowly, and with enormous waste, but they act.

Stacking every worst-case assumption simultaneously – rapid displacement, no consumption recovery, no policy intervention, no new industry formation – creates a scenario that's internally consistent but historically unlikely. Real economies are messier. Negative feedback loops exist, but so do circuit breakers.

Citadel Securities published a detailed rebuttal making a related point: productivity shocks are positive supply shocks. They lower costs, expand output, and increase real income. Citadel also pointed to Indeed data showing software engineer demand up 11% year-over-year in early 2026 and Fed data showing generative AI daily workplace usage remaining “unexpectedly stable.” If displacement were accelerating the way the doom narrative assumes, those numbers would look different.

Argument 2: services deflation

Michael Bloch published a direct counterpoint called “The 2028 Global Intelligence Boom.” His argument: most consumer spending goes to services that are essentially complexity navigation. Tax prep. Insurance comparison. Travel booking. Real estate transactions. Legal review. AI agents could compress these costs by 40-70%.

His numbers: average household spending on complexity-navigation services runs \$8,000-12,000 annually. If AI deflates that by half, it returns \$4,000-7,000 per household per year. That's a tax-free raise that requires no legislation.

That money doesn't evaporate. It goes back into the economy. Bloch points to the Census Bureau's January 2026 business formation data: 532,319 new applications, up 7.2% from December. The cost of launching a business – software, legal, accounting, marketing, design – has dropped 70-80%. People who lose white-collar jobs don't all sit at home. Some of them start companies.

In Bloch's scenario, real median household purchasing power rises 18% from 2025 to 2027, the largest three-year gain since the postwar boom. Not because wages skyrocket, but because services get cheaper. A household that needed \$100K in 2025 needs only \$85K in 2027 for the same standard of living.

Is this optimistic? Yes. Is it as internally consistent as the doom case? Also yes. Which is the whole problem.

The capability-dissipation gap

Both narratives sound compelling because both describe real forces. What's missing is the variable that determines which force dominates and when.

There are two curves on the same chart.

The first is AI capability. It goes up fast. Gemini doubled its reasoning benchmark scores in three months. Claude went from "interesting chatbot" to "autonomously modernizes COBOL codebases" in eighteen months. METR's autonomy evaluations show capabilities doubling every few months on unsaturated benchmarks. This curve is steep and accelerating.

The second curve is societal dissipation, the rate at which AI capabilities actually permeate the real economy. How fast companies reorganize workflows. How fast regulators approve new processes. How fast the accountant down the street stops doing things the way she learned in 2015. This curve is much, much flatter.

The gap between these two curves is where we all live right now. It explains everything that seems contradictory about this moment. Why AI demos are stunning but most offices haven't changed. Why the stock market swings violently between euphoria and panic. Why both Citrini and Bloch can be right about the direction but wrong about the timing.

The doom scenario assumes dissipation accelerates to match capability – that the economy absorbs AI at something close to the rate AI improves. The boom scenario assumes the gap stays wide enough for adjustment. Neither assumption is tested because we've never had a technology that improved this fast.

Four forces that keep the gap wide

The gap doesn't persist by accident. There are structural reasons AI capability doesn't translate instantly into economic disruption.

Regulatory inertia

Financial services need regulator approval to change processes. Healthcare needs HIPAA and FDA clearance. Government procurement cycles run three to seven years. 95% of ATM transactions still run on COBOL – not because banks don't know better but because migrating a system that processes trillions of dollars annually is a decade-long risk management exercise. Nobody's doing that because of a blog post, no matter how good Claude is at reading the code.

Organizational inertia

Companies aren't rational actors that optimize instantly. Headcount decisions filter through HR policies, employment law, union agreements, and management politics. Someone in legal has concerns. Someone in HR has a process. The gap between "Claude can technically do this task" and "we've reorganized our department around AI" is measured in years, not months. Pilot programs at large enterprises take so long to evaluate that the AI capability they were piloting becomes obsolete before the report is filed.

Cultural inertia

Most people still don't use AI daily. When Toby Lutke, the CEO of Shopify, a tech company, had to issue a company-wide memo in April 2025 making AI usage a "baseline expectation" and built it into performance reviews – at a company full of engineers who presumably understand the technology – that tells you how slowly habits change. Now multiply that by every non-tech company in America. The law firms. The accounting practices. The hospital that still faxes referrals.

Trust inertia

Even organizations that want to adopt AI can't trust output by default. They shouldn't. Building verification systems, evaluation harnesses, and human-in-the-loop processes is expensive and slow. The shift from "I do this work myself" to "I verify AI doing this work at scale" is an enormous organizational transformation. Nate Jones has been writing about this specific bottleneck. The generation problem is solved. The review problem is just beginning. The companies that have built systems where AI reviews AI, with humans handling exceptions, have a structural advantage that compounds over time. Most companies haven't started.

Why this gap is the opportunity

The people operating at the capability frontier while the rest of the economy moves at the dissipation rate capture outsized returns. This isn't a prediction. It's arithmetic. If you can do something that AI makes possible, and the people competing with you haven't adopted that capability yet because inertia is keeping them in 2024, you have an advantage that persists until they catch up.

Because the four forces of inertia are structural, not temporary, catching up takes years. The organizational architecture required (testable specifications, eval harnesses, rollback processes) can't be bought off the shelf. It has to be built. And each new model release makes existing AI fluency more valuable, not less, because new capabilities land on a foundation of practical understanding that took real time to develop.

The solo consultant who integrated AI into their workflow last month is doing work that their competitors will spend the next two years learning to replicate. By then, the consultant will have moved again.

Where you sit on the chart

If you're reading InsiderLLM, you're probably on the capability curve, not the dissipation curve. Think about what that actually means.

You're running [Qwen 3.5 locally](#) while your coworkers haven't heard of it. You set up [OpenClaw](#) last month while your company is still in a quarterly AI review meeting. You understand [quantization](#), [VRAM budgets](#), prompt engineering, model routing – practical AI fluency that most professionals won't develop for years.

Toby Lutke runs structured evaluations against every new model release. You do the same thing every time a new GGUF drops on HuggingFace, even if you don't call it that. You're pulling the model, testing it against your use cases, comparing it to last week's release, and building intuition about what works. That intuition compounds.

The people panicking about the Citrini memo and the people celebrating the Bloch response are both watching from the sidelines. You're not watching. You're running the models.

What to do with this

This isn't "learn AI" – that was 2024's advice and it was too vague to be useful.

The career move is specific: become the person who can walk into a room of executives reading Citrini's memo and say "I've tested this. Here's what AI can actually do in our workflow. Here's what it can't. Here's the 90-day plan." Technical people understand models. Business people understand workflows. Almost nobody bridges both. That's the gap within the gap. It pays well and it will pay well for years because the dissipation curve is slow.

The practical move is even simpler: keep doing what you're doing. Every week you spend running models on your own hardware, testing new releases, figuring out what's real versus what's hype, that's building the kind of grounded AI fluency that can't be faked and can't be acquired in a weekend bootcamp. Every month the capability-dissipation gap stays wide is a month you're compounding an advantage.

Stop doom-scrolling market sell-offs. The stocks will do whatever they do. What you control is how fast you close the gap between AI capability and AI integration in your own work.

A note on negativity bias

The Citrini memo hit harder than the Bloch response for a reason that has nothing to do with which is more accurate. We're wired to pay disproportionate attention to threats. "AI will crash the economy" gets 10-50x more engagement than "AI-driven deflation could increase real purchasing power for median households." The doom narrative isn't wrong because it went viral. But the fact that the equally rigorous counterargument barely registered should make you question whether you're getting the full picture.

Both sides have real data. Both have internally consistent models. The honest answer is that nobody knows which curve steepens faster. But the honest response to uncertainty isn't paralysis. It's preparation. And preparation, in this case, looks like exactly what you're already doing: [running AI on your own hardware](#), building practical fluency, and closing your personal capability gap while the rest of the world argues about whether the gap exists.

Related guides

- [VRAM Requirements for Local LLMs](#)

- [OpenClaw Setup Guide](#)
- [Best Local LLMs for Mac in 2026](#)
- [Running LLMs on Mac M-Series](#)
- [Qwen 3.5 Local AI Guide](#)
- [Best Models for OpenClaw](#)

Get notified when we publish new guides.

[Subscribe](#) — free, no spam

Source: <https://insiderllm.com/guides/ai-market-panic-capability-dissipation-gap/>

Free guides for running AI locally