

AI Capital Displacement & Its Consequences

By Robert E. Tucker

Grey Matter Partners

January, 2026

AI Capital Displacement & its Consequences

Executive Summary

In 2025, job creation slowed sharply even as GDP remained resilient. The popular explanation, AI-driven labor replacement, fits surprisingly little of the macro data. The more plausible mechanism is capital displacement: trillions redirected into low-labor-intensity infrastructure, pulling financial oxygen away from job-dense sectors. The result is stagnation without layoffs, and an economy growing around a narrow, capital-heavy core.

1. Soft Hiring, Strong GDP: The Disconnect

The labor market is weakening without the usual signals of recession or mass layoffs.

While there are many indicators that the current economy remains healthy, new job creation in 2025 was sluggish. In the last year, the US economy created only about one quarter as many net nonfarm payroll jobs as it did in 2024 which was itself a weaker year than the post-pandemic boom. While government job losses played a significant role in the scale of the shortfall, even excluding government employment reductions, new job creation has lagged persistently and significantly. The result is a labor market that looks stagnant by recent historical standards, even as headline measures continue to suggest ongoing economic expansion.

A story that has gained increasing traction in media outlets and in conventional wisdom has been that we're seeing the first waves of AI job replacement, but there is limited evidence so far of broad, economy-wide AI-driven job displacement. Instead, recent research has tended to be relatively pessimistic about AI enterprise capabilities and impact. Recent reporting on MIT-linked research, for example, found that roughly 95% of enterprise GenAI pilots are not yet delivering measurable P&L impact despite billions in enterprise spending. Companies are seeing productivity gains, but these gains are more consistent with cost cutting and expense containment, and not on the scale or of the nature we'd expect from massive AI-related capital commitments. According to research from Gartner, 2025 AI spending (including data centers, infrastructure and corporate implementation) was in the range of 1.5 trillion dollars. Commitments and contracts over the next few years suggest that cumulative AI spend could easily reach 6-7 trillion dollars by the end of 2027, half of which (~3T) is US-based spend.

2. Scaling Hits Constraints

AI spending is exploding even as marginal capability gains get slower, costlier, and harder to monetize.

To put this number in perspective, AI spending in 2025 was a line item that's larger than half of annual global military spending - half of the military spending on planet earth. But it's not creating very many jobs. AI isn't taking your job. It's taking the line item for jobs and forcing it to become smaller. AI is absorbing free cash that would otherwise have gone to payroll growth, hiring, and labor-dense expansion. A trillion dollars in AI investment is a trillion dollars that otherwise could have been invested in sectors or businesses with much higher ratios of capital to job creation. In short, AI isn't yet working at scale, but it is costing at a scale that has inevitable implications for the labor market.

We thus have demonstrably massive AI spend, equally demonstrable lack of AI contribution to corporate earnings outside the infrastructure winners, and multi-year capital commitments that consume trillions of dollars of financial oxygen from the labor market. The emperor might not be naked, but he is thinly dressed.

We also appear to be moving out of the era of "fast and easy progress" leading to predictable jumps in capability. Core scaling approaches (add more context, add more compute, add more training data) are showing diminishing returns, and progress is getting dramatically more expensive in capital and energy, even as data center electricity demand accelerates.

Hallucinations remain a persistent structural problem under current training and evaluation incentives, not a simple bug to patch. Security issues such as prompt injection are also deeply rooted, because these systems tend to mix instructions and data in ways that are hard to fully separate. And value alignment is not just an engineering task. It runs into epistemology and real-world human priorities, where "correct" is not a purely technical target. More compute may buy incremental capability, but it does not, by itself, resolve these structural constraints. AI has stopped obeying Moore's law and is starting to look much more like Boeing.

3. The Danger of Historical Precedents

History suggests that foundational infrastructure booms rarely crown early winners and often destroy capital before stabilizing.

Those who argue that additional trillions in AI investment tend to argue that productivity gains are lagged and that today's disappointing ROI is exactly what you should expect in the early stages of a general-purpose technology. On this view, we're not seeing returns because

companies haven't reorganized workflows, retrained staff, and rebuilt systems around new capabilities. Strong proponents of AI spend liken this moment to electrification and railroads: AI may look inefficient in its first phase, but pulling back now risks repeating the classic mistake of abandoning a transformative infrastructure build just before the payoff arrives.

This argument should make everyone nervous. The first half-century of railroad development was plagued by serial collapses: speculative "paper railroads," many of which never laid track, followed by massive capital destruction and a real banking-and-business shakeout in the Panic of 1893, the worst U.S. depression until the 1930s. The path to electrification was no smoother. Even foundational architectural choices were contested for years, and while AC ultimately prevailed, DC systems lingered for decades in pockets of the economy. Electricity later became central to modern life, but only once demand caught up to the infrastructure and the system stopped behaving like a speculative experiment.

4. AI Capital Spend & Missing Jobs

The jobs slowdown looks less like automation and more like trillions being redirected into low-employment infrastructure.

Which brings us back to jobs. Every decision-maker reading these words knows, roughly, what their company is spending on AI and AI-adjacent implementation. Our clients increasingly describe an "arms race" mentality: spending driven less by demonstrated advantage than by fear of falling behind, even as the commercial "weapons" themselves remain unreliable and uneven in practice. Yet it is almost impossible to walk into a board meeting and say you want to pause a multi-year AI program because headcount now looks like the more dependable path to growth. That story sounds like retreat in the face of media certainty that AI is the next electricity.

But the money has to come from somewhere, and payroll is often the most tempting place to find it. Flat headcount and cost compression can mechanically raise productivity in the technical sense, but there is a point where rational corporate incentives, applied at scale, weaken the collective economic environment. One useful leading indicator is the University of Michigan's Consumer Sentiment Index. It remains anomalously low by historical standards a full 2-sigma below the mean at the time this article was written. When consumers stay this pessimistic, spending becomes fragile. Stress tends to surface first in the weakest parts of household balance sheets, then work its way up the value chain.

And the opportunity cost for this level of AI investment? The scale speaks for itself. A \$3T housing push (6 - 10 million units) plausibly supports on the order of ~10 to 30 million job-years of employment during the buildout, depending on the mix of single-family vs multifamily construction. At healthcare's labor intensity, \$3 trillion is not an abstract number. It plausibly

funds 20 to 30 million job-years of direct healthcare employment, meaning 10 to 15 million jobs supported per year if deployed over two years. That's what job-dense capital looks like. Compare that to AI: A reasonable estimate is that ~\$3T of AI infrastructure and implementation spending supports only a few million direct job-years, not tens of millions, because the spend is concentrated in capital-intensive buildout rather than labor-dense sectors. That's an order of magnitude fewer jobs created per dollar of spend. That's what leads to a healthy economy only creating half a million jobs in a year. It's not the only reason, but it's hard to argue that it's not a significant part of the reason.

5. The Economics of a Crowded Table

Even if a few firms win big, crowded economics and weak exit paths make average AI returns structurally uncertain.

It's true that housing and healthcare don't carry the same valuation multiples. Fine. But neither can AI, at least not on today's evidence. Outside the infrastructure winners, most companies still can't point to a clean, measurable P&L impact that matches the scale of spending. If this were almost any other investment category, boards and markets would be demanding clearer proof of return by now. Instead, we have a strange inversion: valuation increases quickly while value grows slowly, if at all.

AI also has an exit-structure problem. The ecosystem competes with itself in ways that make large acquisitions hard to justify. There is no obvious accretive logic for one frontier model company to buy another when the customer bases overlap and profitability is still thin. And the infrastructure layer is even less buyable: NVIDIA isn't going to acquire TSMC, and the hyperscalers aren't going to merge into each other. Once companies are passing through \$3 to 4 trillion in market cap, the normal M&A logic breaks. Who is going to buy a \$4T company and credibly explain to shareholders how that's accretive?

All of which means, logically, that the average rate of return on AI investment cannot resemble what the infrastructure winners will earn, and that's a sector problem. With housing and healthcare, we understand the business models, we understand the infrastructure, and we have decades of pattern recognition around when consolidation works and when it doesn't. With AI, those questions are still unresolved: where durable margins live, which layers become commoditized, and what "normal" returns will look like once the buildout phase ends.

We are effectively running a multi-trillion-dollar global experiment in value creation, and one of the few things we can say with confidence is that AI infrastructure is among the least job-dense places to deploy capital at scale. If even a portion of that spending were redirected toward labor-intensive sectors, it would function as one of the largest employment impulses in modern history.

6. Cautions for Decision-Makers

The biggest risk is mistaking narrative compliance and short-term tightening for durable strategic advantage.

What then to do? Three cautionary notes offer a practical first step:

The first is to exercise caution **mistaking narrative safety for strategic safety**. Executives approve AI spending not because it clearly advances the business, but because it protects them from future blame. Saying yes aligns with conventional wisdom. Saying no requires being right later, with evidence. In that environment, leaders optimize for cover, not advantage. The result is capital committed for reputational reasons rather than operational ones.

The second caution is not to **confuse constraint with progress**. Hiring freezes, tighter workflows, and higher output per employee are taken as signs that AI is working. In reality, many of these gains come from squeezing existing teams harder and stripping out slack. It looks like efficiency, but it is often just strain. Leaders mistake short-term tightening for long-term leverage, and by the time the difference becomes obvious, capacity has already eroded.

The third caution is not to **treat growth as reversible and capital as permanent**. Large AI investments are approved as long-term, irreversible commitments, while hiring, expansion, and talent development are treated as flexible and deferrable. This flips the traditional logic of adaptability. Firms lock themselves into rigid infrastructure while starving the very systems that allow them to respond to change. What feels prudent in the moment quietly reduces future options.

These mistakes compound. Capital pours in under the cover of efficiency. Damage is deferred, not denied. Growth that never happens is treated as harmless. From the outside, the system looks disciplined. Inside, it is quietly eating its future. Until that pattern is named, it will continue to be mistaken for prudence and repeated as policy.