

THE point.

Living Through Your First Industrial Revolution

March 8, 2026

By:

CIO

THE CIO GROUP



Living Through Your First Industrial Revolution

By Steven Wieting and David Bailin

With no reference points, AI will usher all of us into a period of technological, social and psychological changes so dramatic that it will alter how we live, work and play. Welcome to your first industrial revolution.

A few weeks ago, a speculative essay by Citrini Research titled “The 2028 Global Intelligence Crisis” imagined a future in which AI’s success becomes the economy’s undoing - mass white-collar displacement, a 38% drawdown in the S&P 500, and a consumer economy in freefall. Investor reaction to the piece moved the Dow down 821 points in a single session. Michael Burry, the hedge fund manager who profited hugely from the mortgage crisis, shared it with the comment, “And you think I’m bearish.”

Within days, Citadel Securities published a systematic rebuttal, arguing that Citrini had ignored the natural brakes that slow any technological transition. Linas Beliūnas wrote a careful dissection concluding that the two-year crisis timeline was implausible - but that a slower, decade-long grind might ensue and never trigger the necessary government response. Nate Silver warned that the political impact of AI is “probably understated” and that Silicon Valley would play its hand “overconfidently and badly.” Meanwhile, Michael Spencer argued that most people underestimated the sheer scale of AI infrastructure being built and therefore underestimated its impact.

Every author agrees on the enormity of the impact of AI. But they all view the rate of change and the range of impacts very differently. By applying a framework organized along six dimensions we can understand the most likely impacts and develop a reasonable sense of how fast these may occur. We have labeled the six as follows: the Rate of Change, the Competitive Divide, Revolutionary New Business Models, Economic Consequences, the Software Question and Politics. Taken together, they describe something we have no living memory of - a major industrial revolution happening at 10x the speed of steel.

I. The Rate of Change: What Is Real and What Is Not

Every business in America is confronting the same question: what happens when the cost of cognitive work drops by an order of magnitude or more?

The efficiency gains are real. AI agents do not take sick days, negotiate raises, or require benefits. They draft, summarize, code, route customer inquiries, and manage workflows at a fraction of the cost of the humans who did it before. For companies that deploy these tools well, the impact on margins is immediate. These tools allow for layoffs.

Agentic coding tools have reached the point where a competent developer working alongside AI can replicate core functionality of a mid-market SaaS product in weeks. A Fortune 500 procurement manager told Citrini that he had been in conversations with AI companies about



replacing a vendor entirely. The vendor renewed the contract at a 30% discount. As Beliūnas noted, that anecdote rings true not because AI replacement is imminent, but because the threat of it may shift bargaining power, especially for followers in major markets.

Citadel noted that productivity shocks are positive supply shocks in economic terms: they lower marginal costs, expand potential output, and increase real income at the macro level. That is the textbook answer. But Beliūnas identifies a critical nuance: the rate of AI capability improvement will not uniformly match the rate of economic and institutional adoption.

Enterprise software adoption moves at the speed of budget cycles, legal review, integration timelines, and institutional risk tolerance. E-commerce took from 1995 to 2010 to reach 5% of US retail sales. The smartphone, perhaps the fastest consumer technology adoption in history, took five years to reach majority penetration.

So, the likely scenario is not Citrini's two-year cataclysm. It is a rolling wave of efficiency gains that arrive unevenly across industries, accelerating over three to seven years, with some sectors transformed quickly and others absorbing the change more slowly. That is still enormous. It is just not instant.

II. Adopt Fast or Degrade Faster

The businesses that actively employ AI - fast and well - will see expanded margins, accelerated product cycles, and a structural cost advantage over those that do not. Code that took a team three months to write ships in three weeks. A legal review that consumed a paralegal's week takes an afternoon. Customer support queries that once required humans resolve autonomously.

We believe that businesses that do not adopt will see fairly rapid competitive degradation, not a slow erosion - a step function. A competitor that is 20% more efficient is annoying. A competitor that is 80% more efficient and willing to price accordingly is an existential risk. The competitive gap between AI-adopters and AI-avoiders will widen faster than any gap created by prior technology waves, because the tool itself improves with use while the laggards who do not adopt it stand still.

One of Citrini's sharpest observations is that the historical disruption model - where incumbents resist new technology, lose share to nimble entrants, and die slowly - will not apply here. The incumbents cannot afford to resist. The most threatened companies are likely to become AI's most aggressive adopters. Cut headcount, redeploy savings into AI, maintain output at lower cost. Each company's individual response is rational. The collective result is an acceleration of change across every sector with a highly paid, white-collar cost structure and the use of robotics for many logistical roles among blue collar workers.

This creates three categories of business. AI-native insurgents that have no legacy costs and build from scratch with AI at the core. AI-adaptive incumbents that restructure aggressively



enough to compete - and many will, precisely because the tool that threatens them is the same tool they can wield. And AI-indifferent companies that become the Kodaks and Blockbusters of this cycle. The middle category defies the standard disruption narrative and is the most important: incumbents with deep customer relationships, proprietary data, and regulatory moats can fight back effectively if they move fast enough.

This scenario applies directly to the US economy because we have few employment protections. Economies that protect employment with very high costs of change may discourage AI adoption, limit some macro and micro disruption, but ultimately place their firms at a growing competitive disadvantage.

III. New Business Models, Old Ones on Fire

AI does not optimize existing businesses by itself. Its strategic deployment can either generate creative disruption or structural failure.

A highly vulnerable category is what Citrini labelled “habitual intermediation” - businesses built on monetizing friction, inertia and human limitations. Over fifty years, the American economy constructed businesses that profit from busy, impatient, and willing clients who did not understand the degree to which they paid dearly for convenience. Subscriptions that passively renew despite months of disuse. Insurance policies that count on policyholder inertia or that charge for unnecessary coverage. Financial advice premised on navigating complexity that AI finds neither complex nor tedious.

Any business relying on human to human interaction where the “relationship” mattered will be subject to AI’s brutal transparency. As AI agents begin handling consumer decisions, the businesses that depended on those frictions will see their moats drain. An AI agent does not have a favorite app. It does not feel the pull of a well-designed checkout experience. It checks every option and picks the cheapest, fastest one. Advice can displace relationships when performance measurement and fee transparency become spotlighted.

But that is only half the story. New AI-native businesses are emerging that are scalable and profitable in ways their predecessors never were. They are born capital-light, with marginal costs approaching zero. A company that once needed fifty engineers and a two-year runway can launch with five people and a few months.

When the cost of a factor of production collapses, demand for what it enables expands. Cheaper software means more software gets built. Cheaper analysis means more decisions get informed by analysis. This is the historical pattern that Citadel identified and it is real. Keynes predicted the 15-hour workweek because he underestimated the elasticity of human wants. We did not work less. We consumed more and worked more.

There are some limitations to this argument. Visa and Mastercard are not going to face fee pressure immediately because the value of their networks is high, structural and initially hard to



replicate. But one would be naïve to think that other networks and peer to peer systems will not take real market shares as stablecoins evolve.

Disruption is not just the replication of code and a business model. DoorDash's moat is not the code - it is the two-sided marketplace of driver, restaurant, and consumer density in every geography. You can replicate the app in weeks; you cannot replicate the network. Similarly, the stablecoin payments scenario - AI agents routing around credit card interchange via Solana - requires regulatory frameworks, merchant adoption, and consumer trust infrastructure that simply do not exist yet.

The disruption of intermediation businesses will be uneven, contested, and slower than the most dramatic scenarios suggest. Some incumbents will die. Some will adapt. The new models will be scalable and profitable. But the transition will be messy, not clean.

IV. The Economic Consequences

A GPU cluster producing the output previously attributed to 10,000 workers generates GDP without generating the restaurant visits, apartment leases, and income tax revenue those workers produced. That type of output shows up in national accounts but never circulates through households. This is a rapid step function accelerating a trend already underway. Labor's share of income has been declining for decades (see Figure 1). Productivity generates profits and only a small portion benefits workers, white or blue collar.

Bank of America's analysts noted in February 2026 that recent productivity gains have powered corporate profits, with labor income steadily falling as a share of GDP. The economist Robert Allen told the Financial Times that this dynamic - what he calls an "Engels pause," where GDP rises while wages stagnate - may have been underway in the US since the early 1970s. AI is likely to accelerate that dramatically.

If unemployment rises meaningfully - and some credible observers expect it will, even if the timeline is unclear at present - the consequences may cascade through the consumer economy in predictable ways. The top 20% of earners account for roughly 65% of all consumer spending. These are the workers exposed to AI-driven displacement in professional services, finance, consulting, and technology.

If the result were linear, unemployment would see consumption decline faster relative to the number of jobs lost. A 2% decline in white-collar would equate to a 3–4% hit in discretionary spending. But this assumes that there are no new businesses formed, that Stripe's 4,000 layoffs do not equal 25 or more new businesses. It is likely that technologically displaced workers will have a strong motivation to rise as challengers in their industries.

It should not be dismissed lightly that every technological advance in history has ultimately made the economy bigger. While there was no vision or appetite to take macroeconomic steps to stabilize employment during the first industrial revolution of the late 18th century, Keynes



provided the rationale in the 20th century. The pandemic of the 2020s shows the scale at which governments can act to intervene in economies (see Figure 2).

With this said, if the government is unable or unwilling to assist those displaced - through retraining programs, direct transfers, or expanded safety nets - and if the Federal Reserve's lowering of interest rates proves insufficient to stimulate renewed hiring (because the displacement is structural, not cyclical), the negative economic impacts can accumulate and become self-reinforcing:

For consumers: Consumer confidence declines not because of a single shock, but because the ground feels unstable. The employed spend cautiously because they fear being next.

For spending and demand: Discretionary categories contract first - dining, travel, entertainment, home improvement. But the contraction bleeds into housing if real incomes fall and stay down. If mortgage income assumptions prove to be suspect, real estate prices could see declines.

All of this depends on the rate of change and on the speed with which white collar workers and executives adapt. Slow increases in unemployment may not trigger an emergency government response. If GDP keeps growing at 2–3% while the median worker quietly loses ground for a decade, the political system has no crisis to react to. The metrics everyone watches keep looking fine while households lose purchasing power so gradually it never registers as an emergency.

In many respects, this is what has unfolded in the past four decades, giving rise to populist political solutions now. Extreme, populist solutions to counter AI may undermine the economy in entirely different ways. This is the scenario we are least able to anticipate.

V. The Software Question

A common misconception is that AI disruption is primarily a software industry story - that SaaS companies are uniquely vulnerable while other sectors watch from a safe distance. The reality is the opposite. AI touches every industry with a white-collar cost structure. But software companies face a distinctive double challenge.

First, they must use AI to code faster, embed AI functionalities into their products, and anticipate entirely new use cases. They must become the “AI inside” that enterprise customers want to buy rather than build. Every renewal conversation now includes an implicit question: could we use our own AI tools for all or part of the functions? The software company's answer must be a convincing “no” - either because its product is deeply integrated into the customer's workflow, or its data moat is irreplaceable, or its embedded AI capabilities are so far ahead that building in-house would be foolish.

Some software companies have moats and time. Systems of record with deep regulatory integration - think Epic in healthcare or Bloomberg in financial data - are not easily replicated



regardless of how good agentic coding becomes. Their switching costs are measured in years and regulatory certifications, not lines of code.

Other software companies have limited barriers to disruption. If your SaaS product is priced per seat and your customers are cutting headcount, your revenue declines even if no customers cancel. You don't need wholesale replacement of the product to get margin compression. You just need your customers to need fewer humans.

The broader principle extends beyond software: capital-light businesses are more vulnerable in this environment because their competition will be even more capital-light. When the primary input to your business is human intelligence, and AI makes human intelligence less scarce, your cost advantage disappears. A consulting firm, a design agency, a legal services provider, a financial advisory practice - each faces competitors who can deliver comparable output with a fraction of the headcount. The barrier to entry drops not by 20% but by 80%. The incumbents with deep client relationships, proprietary methodologies, and brand trust will survive. Those coasting on credentialism and billable hours will not.

Michael Spencer's work on AI scaling reminds us that the infrastructure behind all of this is growing at a pace that defies conventional planning. Trillions of dollars of business investment are flowing into AI compute. Energy constraints are real and might limit how rapid and disruptive AI can be. , Regardless, the investment community and the technology industry are building for a future where AI is not a feature but the foundation of the economy.

VI. The Political Dimension: Bigger Than Anyone Thinks

AI will have both positive and negative impacts on the American economy. Those impacts will be weaponized politically - associated with policies, fear tactics, and blame - and they will bleed into the November 2026 midterm elections with more force than most political observers currently expect.

Nate Silver has been among the clearest voices on this point. In his February 2026 essay "The Singularity Won't Be Gentle," he argued that the political impact of AI is probably understated for several interconnected reasons.

First, Silicon Valley is bad at politics. It is insular, extraordinarily wealthy, and plausibly stands to benefit from changes that would be undesirable to a large and bipartisan fraction of the public. Silver expects the tech industry to play its hand the way any rich gambler on a winning streak would: overconfidently and badly.

Second, cluelessness about AI on the political left means the blowback will be greater once it arrives. Silver has noted that AI leaders like Sam Altman claim their technology will profoundly reshape society in ways nobody was asking for, while accumulating wealth at a pace that makes the Gilded Age look modest. The public's trust in large technology companies has already fallen sharply - Gallup found it dropped from 32% to 24% between 2020 and 2025.



Third, disruption to the creative and professional classes produces outsized political impact. Writers, editors, consultants, financial advisors, and the broader knowledge-worker class are not just economically significant - they are culturally vocal, politically engaged, and concentrated in the swing districts and media markets that determine elections. Silver has observed that AI was not a significant political issue in 2024, but that was probably the last election for which that was true.

The 2026 midterms are shaping up to be a referendum on economic anxiety as much as on any single policy. Marist polling already shows Democrats with a 14-point lead on the generic congressional ballot. Trump's approval has fallen below 40% in Silver's average, dragged down by cost-of-living concerns. AI-related job displacement - even if still in its early stages - will amplify the sense that the economy is working for the few and not the many. Both parties will attempt to claim the issue. Republicans will warn that regulation hands the advantage to China. Democrats will push for worker protections and transfers. The actual policy substance will matter less than the narratives, and the narratives will be shaped by whichever real-world AI impacts are most visible by autumn.

Silver wrote in late February 2026 that the 42% of people who said AI would be only "a marginal issue" or a "non-issue" in the next election cycle are probably going to be wrong. AI is becoming the highest-stakes intersection of economics and politics since globalization, and its effects will be felt at the ballot box sooner than the policy establishment is prepared for.

Our Industrial Revolution

The original Industrial Revolution took roughly eighty years to transform the British and US economies. It produced staggering gains in output, unprecedented concentrations of wealth, decades of worker immiseration before wages finally caught up, political upheaval that reshaped the party system, and social dislocations that took a generation to resolve. It was the defining economic event of the 19th century, and nobody who lived through its early decades understood what it would become.

We are now in the early days of our own industrial revolution. The AI revolution will be bigger - because it affects cognitive work, not just physical labor, and cognitive work is the majority of our modern economy. It will be bolder - because the pace of capability improvement is measured in months, not decades, and the capital committed to its acceleration already exceeds anything in industrial history. And it will be more dangerous - because the political and institutional structures that eventually tamed the first Industrial Revolution's worst excesses took fifty years to build, and we may not have fifty years.

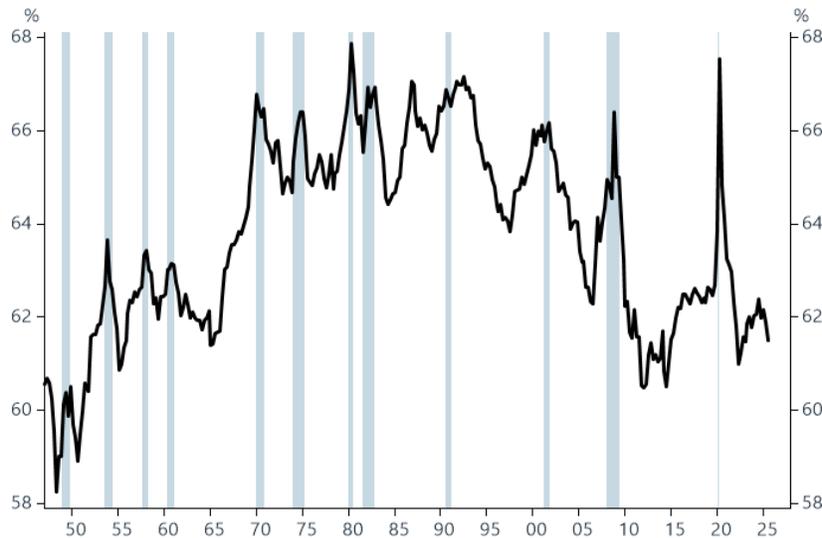
Citrini's two-year crisis probably will not materialize as written. Citadel is right about the S-curve of adoption and the natural brakes that exist. Beliūnas is right that a slow grind is harder to respond to than a crisis. Spencer is right that the infrastructure being built will be the foundation



of the new economy. Silver is right that the politics will be uglier than anyone in Silicon Valley expects.

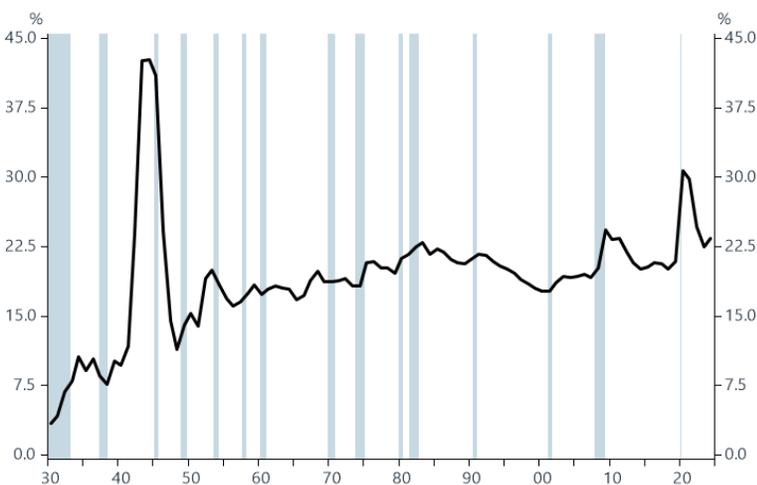
For businesses, the imperative is clear: adopt AI aggressively, or watch your competitive position erode in real time. For investors, the imperative is to distinguish between companies that will ride this wave and those whose business models depend on frictions that are evaporating. For policymakers, the imperative is to start building the institutional architecture for an economy where human intelligence is no longer the scarce input - before the impacts become so significant that a major shift in the role of government forces political change.

Figure 1 US Employee Compensation as a Share of National Income



Source: CIO Group, Haver Analytics

Figure 2: How Large Can Government Intervention be in the Economy? See WWII and 2020 Federal Outlays as % of GDP



Source: CIO Group, Haver Analytics



Sources

The Citrini Research essay referenced throughout, “The 2028 Global Intelligence Crisis,” is available at: <https://www.citriniresearch.com/p/2028gic>

Citadel Securities’ rebuttal, arguing that Citrini confused recursive technology with recursive adoption, was published in February 2026.

Linas Beliūnas’ analysis of the Citrini thesis, concluding that a slower decade-long transition may be more dangerous than a two-year crisis, was published in February 2026.

Nate Silver’s essay “The Singularity Won’t Be Gentle” on the political implications of AI was published in February 2026.

Michael Spencer’s Why Scaling AI is Underestimated was published on Substack on February 2026.

Bank of America’s analyst note on productivity gains powering corporate profits while labor income falls as a share of GDP was published in February 2026.

Robert Allen’s discussion of the “Engels pause”—where GDP rises while wages stagnate—was reported in his research in 2009.

Gallup polling data on public trust in large technology companies (declining from 32% to 24% between 2020 and 2025) is cited in Section VI. Marist polling data on the generic congressional ballot is cited in Section VI.



Contact us for more information

info@ciogroup.com

Stay subscribed

New York

41 Madison Ave, Floor
31, New York, NY

New Jersey

45 W. River Road, Suite 202
Rumson, New Jersey, NJ

Phone

+1 (212) 516-4000