



**BY WILLIAM P. STERLING, PH.D.**  
*Global Strategist*

# GLOBAL PERSPECTIVES

MARCH 2026

## THE COMING INTELLIGENCE EXPLOSION: When AI Builds AI, Who Wins On Wall Street?

HIGHLIGHTS

- ▶ The AI intelligence explosion is likely to become an economy-wide cost story, not just a tech story. BlackRock recently estimated that a modest shift in labor's share of costs could generate \$110 trillion in combined present value and lift corporate earnings by 31% — with beneficiaries extending far beyond technology.
- ▶ Recursive self-improvement compresses the timeline. Frontier labs are automating their own R&D, meaning the AI cost revolution may arrive sooner and cut deeper than most investors assume — and legacy software companies may have half the adaptation window they expect.
- ▶ The market may be underestimating three things at once: the broad earnings uplift from AI-driven margin expansion, the speed of disruption for many industries, including legacy software, and the emerging secular growth category of AI governance and compliance.

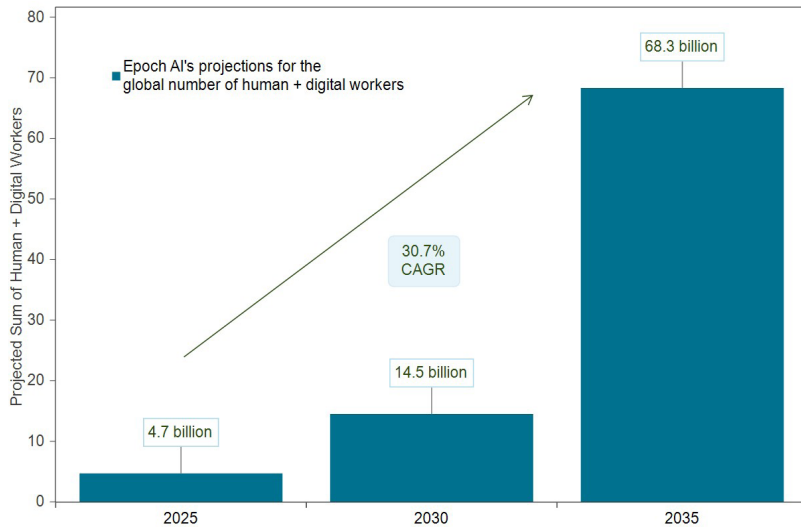
### CALIBRATING THE INTELLIGENCE EXPLOSION

Start with two numbers. The first: 68.3 billion. That is Epoch AI's projection for the global effective workforce — humans plus AI equivalents — by 2035, roughly 15 times today's human labor force of 4.7 billion (**Figure 1**). The second: nine months. That is roughly how

long OpenAI says it will take to deploy “hundreds of thousands of automated AI research interns,” with a fully automated research workforce targeted within two years.<sup>1</sup>

**FIGURE 1**

**The Coming Intelligence Explosion:  
Number of Effective Global Workers to Surge**



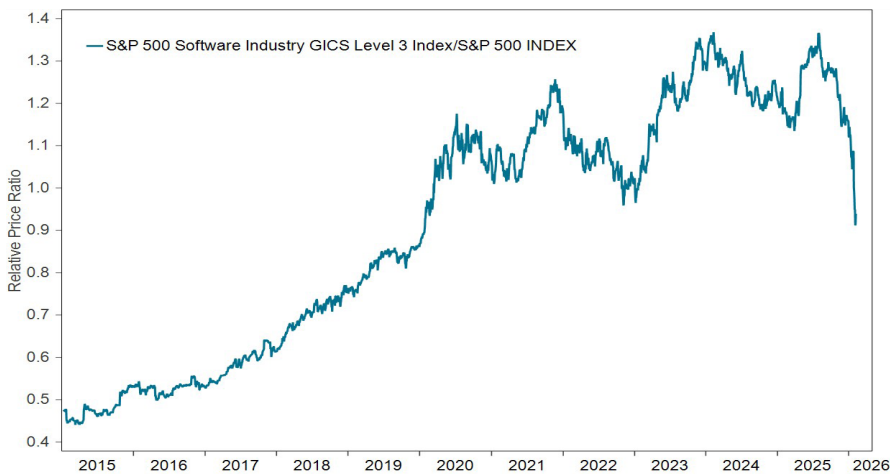
Sources: GW&K Investment Management, Epoch AI, and Macrobond.

Epoch AI projects that by 2035 the global effective workforce — humans plus AI equivalents — will be roughly 15 times today’s human labor force of 4.7 billion.

Investors have recently been dissecting the “SaaSocalypse” — the near-\$1 trillion rout in software stocks triggered by Anthropic’s impressive product launches and the market’s belated reckoning with agentic AI (Figure 2). That conversation matters, but it is focused on the tools. The far more consequential story is the factory that builds the tools: frontier AI labs are automating their own research at a pace that could reshape how investors think about the entire technology stack — and the timeline on which every sector faces disruption.

**FIGURE 2**

**“SaaSocalypse”: Software vs the S&P 500 Index**



Sources: GW&K Investment Management, Bloomberg, S&P Global, and Macrobond.

Nearly six years of the software industry’s relative outperformance has been given up since last August when tools like Anthropic’s Claude Code began to revolutionize the role of AI in speeding up the development of complex software projects.

<sup>1</sup> These figures were quoted in Dean W. Ball, “On Recursive Self-Improvement,” Hyperdimensional Substack, February 5, 2026.

There is, however, a third number that reframes the entire debate — and it comes not from a technologist but from one of the most influential investment minds on Wall Street.

### THE \$110 TRILLION COST REVOLUTION

In his 2026 outlook, BlackRock’s Rick Rieder — the firm’s Chief Investment Officer of Global Fixed Income and a top-contender for the recently filled US Federal Reserve Chair role — frames AI not as a revenue story but as a cost revolution.<sup>2</sup> The distinction is critical. While Wall Street has been arguing about which software companies lose revenue to AI agents, Rieder is focused on the far larger prize: What happens to corporate profitability across every sector when AI compresses labor’s share of total costs?

The arithmetic is startling. Labor currently accounts for roughly 55% of total US business-sector costs. Rieder’s illustrative scenario asks what happens if AI and related automation reduce that share by just five percentage points — from 55% to 50% — and if three-quarters of those savings accrue to the companies deploying AI while the remaining quarter flows to AI service providers.

The result: approximately \$1.2 trillion in annual labor cost savings, translating into roughly \$878 billion in incremental after-tax corporate profits per year. Discounted at reasonable rates, the present value of the corporate piece alone is on the order of \$82 trillion, with another \$27 trillion accruing to AI providers. The combined figure: \$110 trillion (**Figure 3**).

**FIGURE 3**

#### The AI Cost Revolution — Present Value of Labor Savings

Metric	Value
Current labor costs as % of business-sector output	~55%
Pro-forma labor costs (after AI adoption)	~50%
Annual labor cost savings	~\$1.2 trillion
Share retained by corporates	75%
Incremental annual after-tax corporate profit	~\$878 billion
Implied increase in corporate earnings	~31%
Present value of corporate profit increase (5% growth, 3% discount)	~\$82 trillion
Present value of AI provider revenue stream	~\$27 trillion
Combined present value of AI labor savings	~\$110 trillion

Source: Rick Rieder, “The Odds Are Changing: Investing in 2026,” BlackRock Advisor Center, January 2, 2026. Based on Bureau of Economic Analysis data as of 6/30/2025. Illustrative scenario; not a forecast.

A recent illustration by BlackRock suggested that the present value of bringing down labor costs from 55% of total costs to 50% of total costs would amount to about \$110 trillion. For reference, the total market capitalization of US listed equities is about \$70 trillion.

<sup>2</sup> Rick Rieder, “The Odds Are Changing: Investing in 2026,” Blackrock Advisor Center, January 2, 2026.

Put differently: a 9% reduction in labor costs translates mechanically into roughly 31% higher corporate earnings.<sup>3</sup> And that estimate assumes a modest five-percentage-point shift in cost structure — not the wholesale transformation that Epoch AI’s projections imply. If the intelligence explosion delivers even a fraction of what the frontier labs are targeting, the earnings uplift could be substantially larger.

This reframes the SaaSocalypse entirely. Yes, some software companies lose. But the corporate sector as a whole wins — and wins big. The value that migrates away from labor and from certain software vendors does not disappear. It shows up as margin expansion across every industry that successfully deploys AI agents, from financial services and healthcare to logistics and manufacturing. The stock market, in aggregate, is a beneficiary of the intelligence explosion, even if specific sectors within it are disrupted.

### INSIDE THE RECURSIVE LOOP

What makes Rieder’s cost-revolution framework so powerful is that it may be conservative, because it may not fully account for the self-reinforcing dynamics inside the labs that are building these AI systems.

Dean Ball, a researcher at the Foundation for American Innovation, has laid out the mechanics with unusual clarity. The major frontier labs — OpenAI, Anthropic, and Google DeepMind — have already begun automating large fractions of their research and engineering operations. This is not a future aspiration. It is happening now, behind closed doors, and Ball argues it may be the single most important development in AI this year.

The logic is recursive. Today’s most capable AI models — Claude Opus 4.6, GPT-5 and its successors, and Gemini — are already strong enough to contribute meaningfully to AI research. They can write and debug code, run experiments, tune hyperparameters, generate and test hypotheses, and in some cases propose novel model architectures. The labs are scaling this involvement rapidly. Within a year or two, Ball projects, the effective “workforces” of each frontier lab could grow from the low thousands of human researchers to tens of thousands, then hundreds of thousands — digital researchers that, as Ball puts it, “will neither sleep nor eat nor take breaks. And their only objective is to make themselves, the AI systems, smarter” (Figure 4).

**FIGURE 4**  
The Scaling of Frontier Lab Workforces

Timeframe	Human Researchers	AI Agent Equivalents	Total Effective Staff
Early 2026	2,000 – 5,000	Hundreds	~5,000
Late 2026	3,000 – 7,000	Tens of thousands	~30,000 – 50,000
2027 – 2028	5,000 – 10,000	Hundreds of thousands	~200,000 – 500,000

Source: Estimates based on Dean W. Ball analysis and OpenAI public statements. Per-lab illustrative figures.

AI researcher Dean Ball argues that as major AI labs automate their research, within a year or two the effective “workforces” of each frontier lab will grow from the single-digit thousands to tens of thousands, and then hundreds of thousands.

<sup>3</sup> In highly competitive sectors, AI-driven efficiencies may primarily manifest as disinflation and quality gains rather than sustained margin expansion.

If each generation of AI helps build a better successor, the result is a positive feedback loop — not a mythical overnight singularity, but a concerted, compounding acceleration in the pace of improvement. The historical rate of AI capability growth has roughly doubled every four months (Figure 5). Recursive automation could steepen that curve considerably. The physical constraints are real — compute availability, training time, energy infrastructure — but within those limits, the trajectory points toward an exponential increase in effective AI research capacity.

**FIGURE 5**

**Software Task Time Horizon at 50% Success Is Rising Exponentially**



Source: Measuring AI Ability to Complete Long Tasks, by Thomas Kwa, Ben West, Joel Becker, and 21 others, METR.org, <https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/>

Based on analysis from METR (Model Evaluation and Threat Research) and Epoch AI, the “time horizon” of AI agents — defined as the length of tasks a model can complete with 50% probability — has been doubling every 4 months since 2024, an acceleration from a previous 7-month doubling rate.

This 4-month doubling trend indicates that AI systems are becoming capable of handling increasingly longer, more complex, and multi-step tasks at a rapid, accelerating pace.

**WHAT EXPONENTIAL PROGRESS MEANS FOR MARKETS**

The convergence of Rieder’s cost-revolution math and Ball’s recursive-improvement thesis produces an investment framework that many investors may not fully appreciate. If AI capability is accelerating faster than consensus assumes, the cost revolution arrives sooner and cuts deeper than Rieder’s already-impressive base case. And the beneficiaries are not confined to the technology sector.

The bull case for AI infrastructure gets steeper and longer. Companies like Nvidia, TSMC, and the hyperscale cloud providers are already benefiting from a massive compute buildout. If recursive self-improvement works, compute demand doesn’t plateau after a few model generations — it compounds. Energy infrastructure plays, including natural gas, nuclear restarts, and grid equipment, are similarly exposed to upside. The US power grid was not built for an economy running hundreds of thousands of simultaneous AI research processes.

The bull case for broad equities is Rieder’s core insight. A 31% increase in corporate earnings from labor-cost compression alone would represent one of the largest margin expansions in modern market history. And it accrues to any company that deploys AI effectively, not just the companies that build it. Rieder is explicit: “The key is to own the beneficiaries of that transition — businesses with scalable models, durable cash flows, and credible plans for harnessing AI to lift productivity.”

Financial services firms, healthcare companies, industrial conglomerates with complex operations, and professional services firms all stand to capture a share of the \$82 trillion in present value that flows from labor-cost reduction.

The bear case for legacy software could also arrive faster than anyone expects. This is the less obvious implication of recursive improvement. A company that looks safe today because current AI agents “aren’t good enough yet” to replace its specialized workflow might find that gap closing in 12 to 18 months, not the five years its management is banking on.

The SaaSocalypse of February 2026 was triggered by tools that barely existed six months earlier. If recursive improvement is real, the next wave of disruption may hit sooner and harder. Boris Cherny, the head of Anthropic’s Claude Code, shipped 22 and 27 code commits on consecutive days, all produced by Claude. An OpenAI engineer posted simply: “100% — I don’t write code anymore.” The progression from 10% AI-written code in 2023 to 100% at the frontier in early 2026 is the pace investors should be calibrating against (**Figure 6**).

**FIGURE 6**

**The Share of AI-Written Code (2023 – 2026)**

Period	AI Share of New Code	Source / Context
2023	~10% – 15%	Early GitHub Copilot adoption
Mid-2025	~29%	Science study: new Python on GitHub (US)
Late 2025	~30%	Microsoft internal figure (Nadella)
Late 2025 (frontier)	70% – 90%+	AI-native organizations
Early 2026 (frontier)	~100%	Anthropic, OpenAI insiders

Sources: Science; Fortune; Anthropic; Microsoft. Frontier figures represent leading-edge organizations, not broad enterprise adoption.

The share of AI-written code exploded in late 2025, contributing to rapid new product launches and a sudden reassessment of the business models of legacy software companies.

**THE GOVERNANCE PREMIUM**

There is a category of investment implication that the market has barely begun to price: the value of AI governance, compliance, and oversight infrastructure. As AI builds AI and corporate deployment accelerates to capture Rieder’s labor savings, the demand for guardrails scales at least as fast as the capability it is guarding.

Europe’s AI Act is already imposing requirements on high-risk systems. The US is moving toward disclosure and audit standards for AI in finance, healthcare, and critical infrastructure. Every enterprise deploying AI agents in regulated industries will need tools to monitor, log, and verify what those agents do. Audit trails, explainability layers, and compliance dashboards are not glamorous, but they are becoming as essential as cybersecurity was a decade ago — and cybersecurity, once it became mandatory enterprise spending, produced a generation of high-growth, high-margin companies. It is still early days, but AI governance is likely on the same trajectory.

## PRICING THE FEEDBACK LOOP

The market's current positioning tells a familiar story: overweight the hardware, underweight the disruption risk to legacy software, and largely neutral on both the broad corporate margin story and the governance opportunity. That allocation made sense when AI progress felt roughly linear and adoption curves were slow and predictable.

It makes less sense if the labs are right about recursive self-improvement and if Rieder is right that the primary economic impact of AI is cost compression rather than new revenue creation. The Epoch AI projections — a 30% compound annual growth rate in effective cognitive labor through 2035 — are built on exactly this acceleration. Even discounted heavily, they imply a pace of change that stress tests virtually every sector's forward estimates.

The practical investment framework comes down to four key points.

1. AI infrastructure demand is real but cyclical. Nvidia and the hyperscalers benefit from recursive AI development, but capex cycles in compute buildouts have historically been volatile.
2. The disruption timeline for legacy software is shorter than consensus: companies relying on a five-year adaptation window may have half that, or less.
3. Broad corporate earnings stand to benefit massively from Rieder's cost revolution, and the market has not yet priced a 31% earnings uplift into the non-tech sectors that stand to capture the largest share of \$82 trillion in present value.
4. Governance and compliance are the next secular growth category. As AI capability compounds, so does the institutional demand for oversight infrastructure.

The SaaSocalypse discussion asked the right question: what happens when AI agents can do the work? The recursive loop asks the harder one: what happens when they get better at it every single month? And Rieder's framework supplies the answer that matters most for portfolios: the value does not vanish when labor's share declines. It migrates — to corporate margins, to AI providers, and to the investors positioned on the right side of the largest cost revolution in modern economic history.

*William P. Sterling*

William P. Sterling, Ph.D.,  
Global Strategist

### DISCLOSURES:

This represents the views and opinions of GW&K Investment Management. It does not constitute investment advice or an offer or solicitation to purchase or sell any security and is subject to change at any time due to changes in market or economic conditions. The comments should not be construed as a recommendation of individual holdings or market sectors, but as an illustration of broader themes. Data is from what we believe to be reliable sources, but it cannot be guaranteed. GW&K assumes no responsibility for the accuracy of the data provided by outside sources.

© GW&K Investment Management, LLC. All rights reserved.