

The Productivity Paradox.

AI, the Measurement Gap, and the Coming Productivity Boom

DATE	June 2026
AUTHOR	Joseph Logan
PUBLISHER	Expansion Effect Research
WEBSITE	expansioneffect.com
LICENSE	CC BY 4.0

“You can see the computer age everywhere except in the productivity statistics.”

ROBERT SOLOW, 1987

“We’re in 1905. The electric motor works, and everyone’s bought one. Almost nobody has redesigned the factory.”

JOE REIS, APRIL 2026

Contents.

1.	The Solow Lineage: Why Paradoxes Are Normal	7
1.1	<i>The Original Paradox</i>	
1.2	<i>The Electrification Precedent</i>	
1.3	<i>The Steam Precedent</i>	
1.4	<i>The Compressed Timeline Hypothesis</i>	
2.	The Anatomy of the Micro-to-Macro Gap	10
2.1	<i>The Evidence Base</i>	
2.2	<i>Amdahl's Law Applied to Labor</i>	
2.3	<i>The Coordination Cost Problem</i>	
2.4	<i>The Perception Gap</i>	
3.	The Five Countervailing Forces	14
3.1	<i>The Composition Effect</i>	
3.2	<i>The Pseudo-Work Problem</i>	
3.3	<i>The Saturation Effect</i>	
3.4	<i>The Cognitive Erosion Effect</i>	
3.5	<i>The Organizational Inertia Effect</i>	
4.	The Intangible Accumulation: Inside the J-Curve	17
4.1	<i>The Investment Phase</i>	
4.2	<i>The Intangible Capital Stock</i>	
4.3	<i>The Taylorist Transition</i>	
5.	The Measurement Blind Spot: Deflationary Expansion	20
5.1	<i>The GDP Accounting Problem</i>	
5.2	<i>Digital Dark Matter</i>	
5.3	<i>The Right Instruments</i>	
6.	The International Dimension: The US-EU Adoption Divergence	23
6.1	<i>The Adoption Gap</i>	
6.2	<i>The Management Practices Explanation</i>	
6.3	<i>The ICT Precedent and the Coming Divergence</i>	

7.	Sector-Level Exposure: Where the Gains Will Appear First	25
7.1	<i>The Concentration Problem</i>	
7.2	<i>Sector-Level Exposure Matrix</i>	
7.3	<i>The Financial Services Case</i>	
7.4	<i>The Healthcare Paradox</i>	
8.	Actor Mapping: Who Controls the Redesign	28
8.1	<i>The Four Actor Classes</i>	
8.2	<i>The Bottleneck Actor</i>	
8.3	<i>The New Entrant Threat</i>	
9.	The Scenario Matrix: Three Paths to 2028	30
9.1	<i>Scenario Architecture</i>	
9.2	<i>Scenario A: The Fast Path (~25%)</i>	
9.3	<i>Scenario B: The Base Path (~40%)</i>	
9.4	<i>Scenario C: The Slow Path (~35%)</i>	
10.	The Hari Causal Chain: Five Links to the Inflection	32
10.1	<i>Framework</i>	
10.2	<i>The Five Links</i>	
10.3	<i>Joint Probability</i>	
11.	The Leading Indicators Dashboard	34
11.1	<i>Primary Indicators</i>	
11.2	<i>Secondary Indicators</i>	
12.	Open Questions and Intellectual Honesty	35
12.1	<i>The Orchestration Layer May Be Too Small</i>	
12.2	<i>The Measurement Problem May Be Permanent</i>	
12.3	<i>The Redesign Bottleneck May Be More Durable</i>	
12.4	<i>Model Capability Plateaus</i>	
12.5	<i>The Geopolitical Variable</i>	
13.	Methodological Appendix	37
13.1	<i>On the Hari Causal Chain Framework</i>	
13.2	<i>On the Historical Base Rates</i>	
13.3	<i>On the Deflationary Expansion Paradox</i>	

EXECUTIVE SUMMARY

Three years after the public release of ChatGPT, the most widely adopted general-purpose technology in history has produced no measurable acceleration in aggregate total factor productivity. Task-level experiments show gains of 15 to 55 percent on bounded cognitive work. Enterprise pilots report efficiency improvements across dozens of workflows. C.H. Robinson, the world's largest freight broker, reports 40 percent productivity gains since 2022 and 600 hours saved per day from a single AI agent. And yet the Bureau of Labor Statistics' multifactor productivity data for the US private business sector shows no structural break from its post-2005 trend.

This is not a contradiction. It is a pattern. Every general-purpose technology in the modern era has produced the same apparent paradox: the technology works at the task level long before it moves the macroeconomic needle. Steam, electricity, and computing each followed the same trajectory. The lag between adoption and aggregate productivity acceleration has ranged from 15 to 80 years, with the most recent precedent (computing and ICT) resolving in approximately 24 years.

This report makes five arguments.

First, the current productivity stagnation is structurally normal and historically expected. The technology is not failing. The measurement instruments are looking in the wrong place.

Second, the micro-to-macro gap is not a mystery. It is the predictable consequence of Amdahl's Law applied to labor: task-level speedups cannot aggregate into job-level or firm-level productivity gains until the serial bottlenecks (coordination, review, accountability, and organizational design) are redesigned around the new capability.

Third, the \$700 billion in AI capital expenditure currently flowing through the US economy is not productivity. It is the investment phase of the J-curve. Firms are accumulating intangible capital (workflows, training, agent infrastructure, and organizational knowledge) that will not appear in output statistics until the investment phase ends and the deployment phase begins.

Fourth, when the productivity boom arrives, traditional GDP accounting will structurally understate it. The collapse of cognitive task costs toward near-zero is a deflationary expansion: more value delivered at lower measured price. The instruments we use to measure productivity were designed for an economy where cognitive inputs had stable, positive costs. They were not designed for an economy where a legal department can process ten thousand contracts for the same cost as one hundred.

Fifth, the productivity inflection is not a single event. It is a cascade. The first sectors to complete organizational redesign will show visible TFP acceleration. The aggregate will follow with a lag of 12 to 18 months as the redesign diffuses across the economy. The leading indicators that will signal the cascade are not the ones currently being watched.

The joint probability that this causal chain completes on the projected timeline, with a visible TFP inflection in US aggregate data by Q2 2028, is approximately 40 percent. The remaining 60 percent is distributed across scenarios in which the timeline extends, the measurement instruments fail to capture the gains, or the organizational redesign bottleneck proves more durable than historical analogues suggest.

This is not a pessimistic projection. It is a precise one.

1. The Solow Lineage. Why Paradoxes Are Normal.

1.1 The Original Paradox

In 1987, Robert Solow wrote a book review for the *New York Times* in which he observed, almost in passing, that "you can see the computer age everywhere except in the productivity statistics." The observation became one of the most cited sentences in modern economics. It described a genuine puzzle: the United States had been investing heavily in information technology since the early 1970s, and yet measured total factor productivity had been decelerating since 1973. The computers were everywhere. The productivity was not.

The Solow Paradox resolved in the late 1990s. Between 1995 and 2004, US labor productivity growth averaged 2.7 percent per year, compared to 1.4 percent in the preceding two decades. The acceleration was real, sustained, and eventually visible in every major productivity measure. It took approximately 24 years from the mass adoption of the personal computer (roughly 1971) to the visible productivity acceleration (roughly 1995).

The resolution of the paradox was not a mystery in retrospect. Erik Brynjolfsson and Lorin Hitt documented in 1996 that firms with higher IT investment showed higher productivity, but only with a lag of five to seven years, and only when the IT investment was accompanied by organizational change. The technology alone was not sufficient. The organizational redesign was the bottleneck.

1.2 The Electrification Precedent

The computing precedent is the most recent and therefore the most useful for calibration. But the electrification precedent is the most structurally illuminating, because it shows the mechanism most clearly.

Electric motors became commercially viable in the late 1880s. By 1900, approximately 5 percent of US factory mechanical drive capacity was electrified. By 1910, the figure was 25 percent. By 1920, it exceeded 50 percent. The adoption curve was steep and sustained. And yet measured total factor productivity in US manufacturing showed no acceleration during the 1890s or 1900s or 1910s.

The acceleration came in the 1920s. Paul David's landmark 1990 paper, "The Dynamo and the Computer," quantified the discontinuity: trend productivity growth in US manufacturing jumped from 1.5 percent per year during 1899–1914 to 5.1 percent during 1919–1929. The same technology across both periods. Radically different results.

The explanation was architectural. Factory owners in the 1880s and 1890s had replaced their steam engines with electric motors but kept the same factory layout. Multi-story buildings. Central power shafts. Belt-and-pulley systems distributing energy to every floor. The new motor was bolted onto the old architecture. The gains from electrification were real but modest, because the factory was still organized around the constraints of steam.

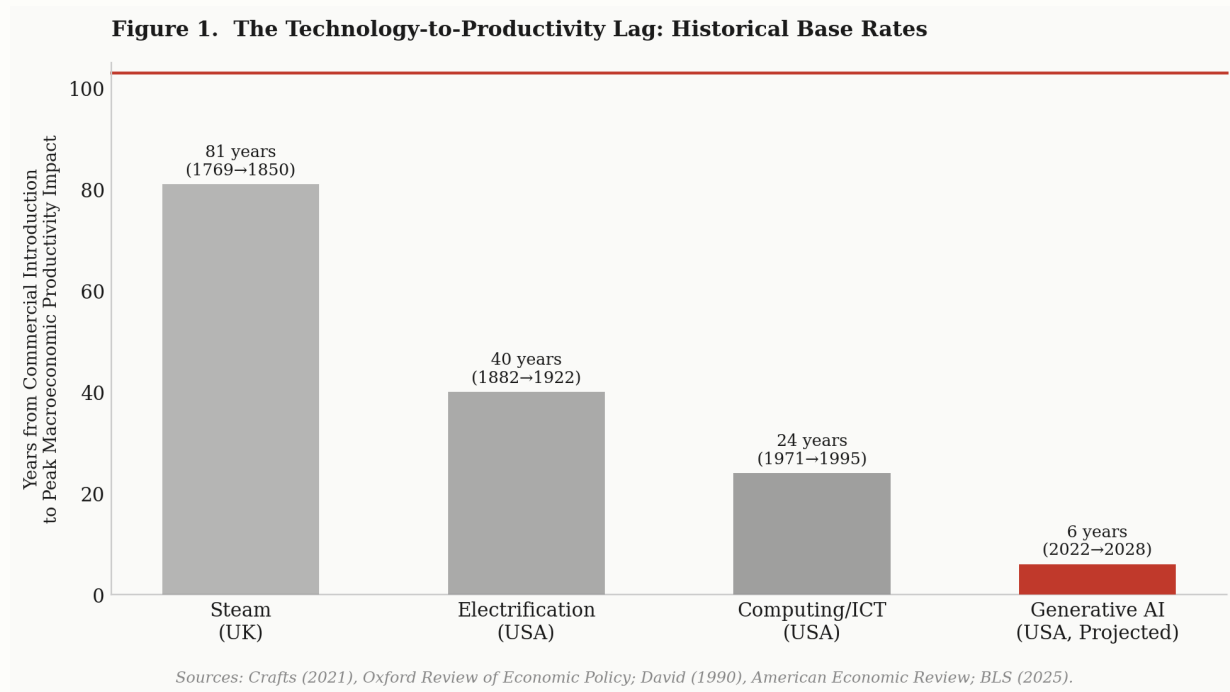
The productivity acceleration came when factories were torn down and rebuilt. Single-story layouts. Individual motors per machine, what engineers called "unit drive." Workflows redesigned around the capabilities of electric power rather than the constraints of steam. The architecture changed, not just the energy source. And when the architecture changed, the productivity gains were not incremental. They were transformative.

The lag between the commercial viability of the electric motor (approximately 1882) and the visible productivity acceleration (approximately 1920) was roughly 38 years. The lag between 50 percent factory electrification and the productivity acceleration was approximately 10 years.

1.3 The Steam Precedent

The steam engine is the longest-lag case in the modern record. James Watt's improved steam engine became commercially available in 1776. Measurable aggregate productivity acceleration in the British economy did not appear until the 1830s and 1840s, a lag of approximately 60 years. The reasons are now well understood: the steam engine required not just adoption but the construction of an entirely new physical infrastructure (railways, canals, urban factories) and an entirely new organizational form (the industrial firm, the factory system, the wage labor contract). The technology was not the bottleneck. The institutional and organizational redesign was.

1.4 The Compressed Timeline Hypothesis



HISTORICAL LAG FROM GENERAL-PURPOSE TECHNOLOGY ADOPTION TO AGGREGATE PRODUCTIVITY ACCELERATION

Each successive general-purpose technology has shown a shorter lag between adoption and aggregate productivity acceleration. Steam: approximately 60 years. Electricity: approximately 38 years. Computing/ICT: approximately 24 years. The compression is not accidental. Each successive technology arrived in an economy with more organizational flexibility, more capital mobility, more educated workers, and more institutional knowledge about how to manage technological transitions.

AI arrives in an economy with all of these advantages plus one additional factor: we know the pattern. Paul David's paper exists. The electrification analogy is widely understood. The organizational redesign bottleneck has been named. This does not eliminate the lag, since organizational inertia is not dissolved by intellectual awareness, but it may compress it further. The working hypothesis of this report is that the AI productivity lag will be in the range of 8 to 15 years from mass adoption (approximately 2022–2023), placing the visible aggregate productivity acceleration between 2030 and 2038 in the base case.

The more optimistic scenario, which is the scenario this report's causal chain models, is that the acceleration begins to appear in sector-level data by 2027 and in aggregate TFP data by 2028. This is not the base case. It is the fast-path scenario, and it requires specific conditions to hold.

2. The Anatomy of the Micro-to-Macro Gap.

2.1 The Evidence Base

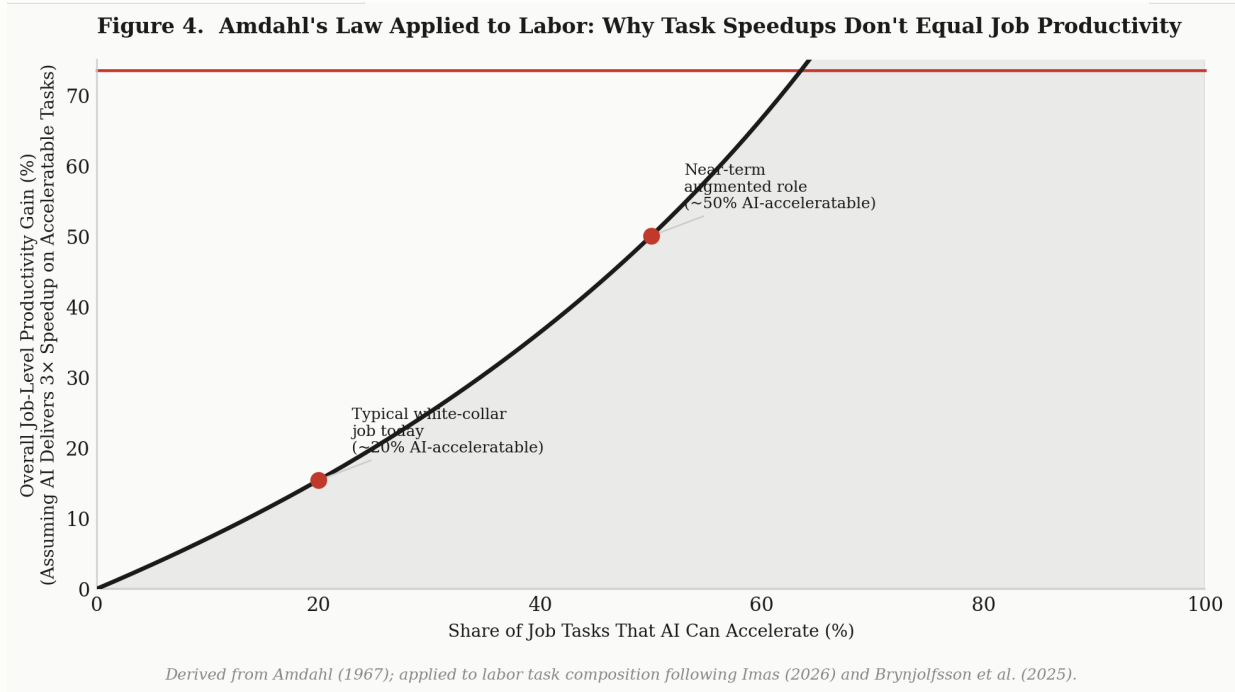
The micro-level evidence for AI productivity gains is substantial, varied, and growing. A selective but representative sample:

STUDY	SETTING	FINDING
Brynjolfsson, Li & Raymond (2023)	Customer service agents, Fortune 500	14% more issues resolved per hour; largest gains for novice workers
Noy & Zhang (2023)	Writing tasks, online experiment	37% faster completion; 18% quality improvement
GitHub / Peng et al. (2023)	Software developers, greenfield tasks	55% faster task completion
BCG / Dell'Acqua et al. (2023)	Management consultants, bounded tasks	25% faster; 40% higher quality on AI-suitable tasks
METR (2025)	Experienced engineers, mature codebases	19% <i>slower</i> with AI; 39-point perception gap
Faros AI (2025)	10,000+ developers, 1,255 teams	21% more tasks; 98% more PRs; no company-level improvement; PR review time +91%
C.H. Robinson (2022–2026)	Freight logistics, full organizational redesign	40% productivity gain; 3M+ shipping tasks automated; 600 hours/day saved

The pattern in this table is not random. The studies showing the largest gains are either (a) greenfield tasks with low institutional complexity, (b) settings where the AI tool was integrated into a redesigned workflow rather than bolted onto an existing one, or (c) both. The studies showing flat or negative results are those where experienced workers used AI tools within existing organizational structures without workflow redesign.

C.H. Robinson is the outlier in the table, and the most important data point. It is the only entry that represents a completed organizational redesign rather than a controlled experiment or a pilot. The 40 percent productivity gain is not a task-level speedup. It is a system-level transformation. CEO Dave Bozeman's description of the mechanism is precise: "We compressed tasks that used to take hours or days into seconds. That's fundamentally changed how fast we can respond to customers." The order-to-cash process went from 15–30 minutes to 32 seconds. That is not a task speedup. That is a workflow redesign.

2.2 Amdahl's Law Applied to Labor



AMDAHL'S LAW APPLIED TO LABOR: MAXIMUM JOB-LEVEL PRODUCTIVITY GAIN VS. AI-AUTOMATABLE TASK SHARE

Gene Amdahl formulated his law in 1967 in the context of parallel computing. The law states that the maximum speedup achievable by parallelizing a portion of a program is limited by the fraction of the program that must remain serial. If 90 percent of a program can be parallelized and 10 percent must remain serial, the maximum theoretical speedup is 10x, regardless of how many processors are added.

The law applies directly to labor productivity. A job is not a single task. It is a collection of tasks, some of which can be accelerated by AI and some of which cannot, at least not yet. The serial portion of most knowledge-work jobs includes: coordination with colleagues, review and approval processes, accountability for decisions, client communication, and the exercise of judgment in ambiguous situations. These are not tasks that current AI systems can fully automate. They are the bottlenecks.

Consider a software developer whose coding speed increases by 55 percent with AI assistance. If coding represents 40 percent of the developer's total work time, and the remaining 60 percent (code review, meetings, documentation, stakeholder communication, debugging in production) is unchanged, the overall job-level productivity gain is approximately 22 percent. This is a real gain. But it is not 55 percent. And it will not appear in firm-level output statistics until the firm's overall throughput increases, which requires that the non-coding bottlenecks also be addressed.

The Faros AI finding (98 percent more PRs merged, no company-level improvement, PR review time up 91 percent) is Amdahl's Law in action. The coding bottleneck was removed. The review bottleneck was not. The work did not disappear. It migrated downstream.

2.3 The Coordination Cost Problem

The MIT Sloan / Microsoft paper "Chaining Tasks, Redefining Work" (April 2026) formalizes a related mechanism. The paper demonstrates mathematically that the productivity gains from AI are bounded not by the speed of individual task completion but by the cost of coordination between tasks. Each handoff from an AI system to a human, for review, approval, or judgment, introduces a coordination cost that partially offsets the task-level speedup.

The implication is significant. The firms that will capture the largest productivity gains from AI are not those that deploy AI tools most aggressively. They are those that redesign their workflows to minimize the number of human-AI handoffs. This requires not just technology adoption but organizational redesign, the same bottleneck that delayed the productivity gains from electrification by 38 years.

2.4 The Perception Gap

The METR (2025) finding deserves extended treatment, because it reveals a mechanism that is not captured by the standard productivity literature.

In a randomized controlled trial, 16 experienced open-source developers completed 246 tasks on mature repositories (1 million-plus lines of code). Developers using AI tools were 19 percent *slower* than the control group. They believed they were 20 percent *faster*. The perception gap was 39 percentage points.

The mechanism is not mysterious. AI tools accelerate the most visible portion of the work (generating code, suggesting completions, drafting boilerplate) while the time spent prompting, reviewing, and correcting near-miss output quietly absorbs the savings. The developer's subjective experience is one of reduced effort, which registers as speed. The clock says otherwise.

This perception gap has a structural implication for the macro data. If workers consistently overestimate their AI-driven productivity gains, then survey-based productivity measures will overstate the gains, while output-based measures will show no improvement. This is precisely the pattern observed in the current data: enterprise surveys report widespread productivity gains; BLS output-per-hour data shows no structural break.

The gap will close when AI tools are integrated into workflows in ways that make the output measurable: when the AI is not a tool the worker uses but a component of a system whose throughput is tracked. C.H. Robinson tracks its AI agents' output in real time. Most enterprises do not.

3. The Five Countervailing Forces.

The Berkeley California Management Review paper by Gruda and Aeon (2026) identifies five structural forces that can neutralize or reverse expected AI productivity gains. These forces are not theoretical. They are observable in current enterprise deployments. Understanding them is essential for projecting when and how the productivity gains will materialize at scale.

3.1 The Composition Effect

The first force is the composition effect: the assumption that any unit task performed by AI results in a unit productivity gain, and that the sum of these gains necessarily produces a net productivity improvement. This assumption fails because it ignores the behavioral, contextual, and systemic dynamics that task-level gains trigger.

The Jevons Paradox, named for the 19th-century economist William Stanley Jevons who observed that improvements in steam engine efficiency led to *increased* total coal consumption, applies directly here. When AI makes cognitive tasks cheaper and faster, the demand for those tasks increases. A marketing team that previously produced four campaigns per quarter may now produce twelve. The per-campaign cost has fallen. The total cognitive load on the team has risen. The net productivity effect depends on whether the twelve campaigns produce proportionally more value than the four, which is an empirical question most enterprises are not yet equipped to answer.

3.2 The Pseudo-Work Problem

The second force is the generation of pseudo-work: AI-generated outputs that mimic productive activity but have no measurable impact on organizational performance. An AI can generate redundant summaries, contextless reports, and responses to queries that never needed to be answered. The low marginal cost of AI generation leads to an inflation of outputs that overwhelm attention flows, validation circuits, and collaboration channels.

The Upwork Research Institute (2024) survey captures this dynamic precisely: 96 percent of executives anticipated productivity gains from AI; 77 percent of employees reported that AI had *increased* their workload. Thirty-nine percent now spend more time reviewing or moderating AI-generated content. The AI is generating more. The humans are filtering more. The net output is not obviously higher.

3.3 The Saturation Effect

The third force is the saturation effect: the phenomenon in which a task becomes counterproductive when its productivity exceeds what the rest of the system can absorb. A legal team that can now review contracts ten times faster creates a bottleneck at the business development team, which must now process ten times more contract redlines. The gain at one node becomes a burden at the next.

This is the organizational equivalent of Amdahl's Law. The system-level throughput is bounded by the slowest node, regardless of how fast the other nodes operate. Until the entire workflow is redesigned to accommodate the new speed of the fastest node, the gains from that node will be partially or fully offset by the bottlenecks it creates.

3.4 The Cognitive Erosion Effect

The fourth force is the most subtle and the most consequential for long-run productivity. As AI takes over functions like writing, summarizing, and proposing ideas, workers tend to withdraw effort from these tasks, not because they are lazy, but because the cognitive load has been reduced. Over time, this withdrawal erodes the human faculties associated with productive effort: analyzing problems, confronting ideas in teams, forming judgments without algorithmic assistance.

The evidence for this effect is early but consistent. Studies show that increased AI use is associated with cognitive offloading and reduced critical thinking. Automation bias, the tendency to accept automated suggestions even when they are incorrect, is well-documented in aviation, medicine, and now knowledge work. The long-run productivity implication is that the workforce that emerges from a decade of heavy AI use may be less capable of the high-judgment work that AI cannot yet perform. This is not a reason to avoid AI adoption. It is a reason to be deliberate about which cognitive functions are offloaded and which are preserved.

3.5 The Organizational Inertia Effect

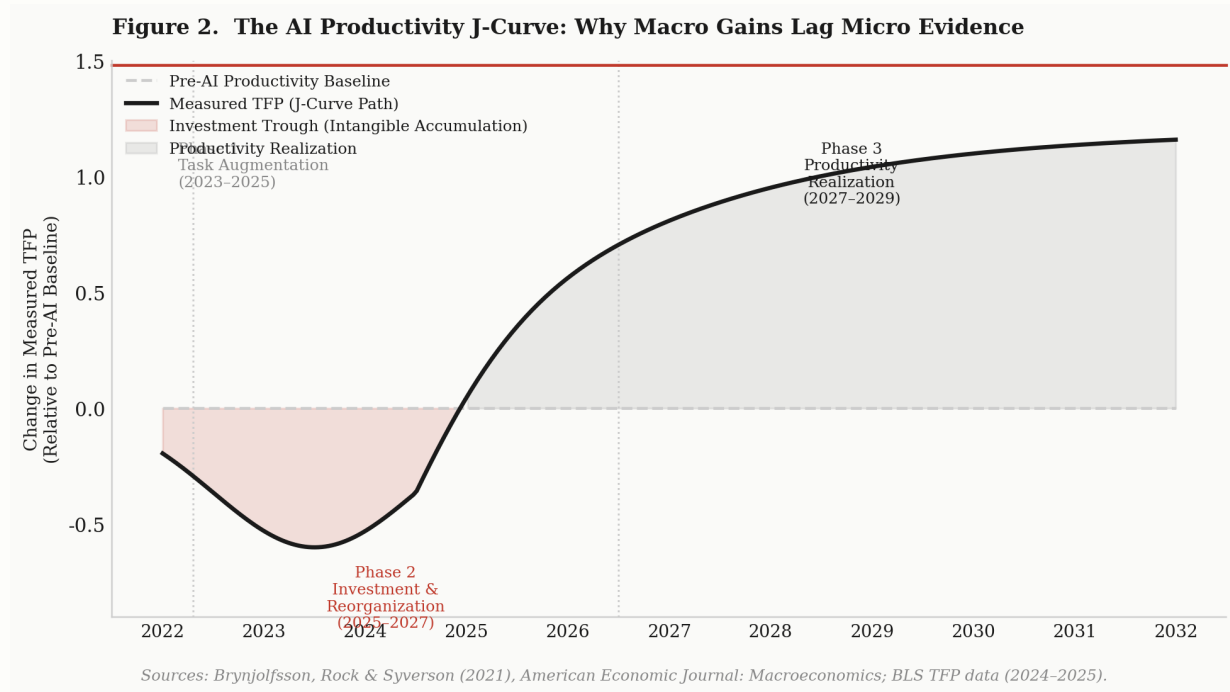
The fifth force is the most familiar and the most powerful: organizational inertia. AI productivity gains are not immediately accessible without profound transformation of organizational settings and underlying norms. When new tools are introduced into organizations with entrenched legacy systems, those systems absorb the technology without altering the underlying processes. The technology becomes grafted onto existing routines, neutralizing much of its transformative potential.

The Berkeley paper identifies two shifts that must occur before the gains become accessible. The first is what the authors call the "Technological Stockholm Syndrome": the paradoxical attachment that arises after the initial resistance to change. Workers who initially resist AI tools eventually become dependent on them, but in ways that replicate the old workflow rather than redesigning it. The second shift is the fundamental redesign of work processes so that AI is not an appendage grafted onto legacy systems but a fully integrated actor within a rebalanced operational ecosystem.

Both shifts take time. The first typically takes 12 to 24 months. The second typically takes 3 to 7 years. This is not a pessimistic estimate. It is consistent with the historical record for every prior general-purpose technology.

4. The Intangible Accumulation. Inside the J-Curve.

4.1 The Investment Phase



THE AI INVESTMENT J-CURVE: INTANGIBLE CAPITAL ACCUMULATION VS. MEASURED PRODUCTIVITY OUTPUT

The J-curve is the most useful framework for understanding the current moment. In the investment phase of a J-curve, firms allocate resources to a new capability (technology, training, organizational change) and output temporarily *declines* relative to what it would have been without the investment. The resources consumed by the investment are not producing output. They are building the capacity to produce output in the future.

The current AI investment cycle is the largest in history. US technology companies have committed to approximately \$700 billion in AI-related capital expenditure in 2026 alone. This capital is flowing into data centers, GPU clusters, model training, software development, and, critically, the organizational infrastructure required to deploy AI at scale: workflow redesign, employee training, agent development, and integration with existing systems.

None of this investment is currently showing up as productivity. It is showing up as capital accumulation, specifically as Intellectual Property Products (software and R&D) in the BLS national accounts. The BLS data shows that capital input growth in the US private business sector accelerated to 2.9 percent in 2024, driven primarily by IPP investment. This is the investment phase of the J-curve. The output phase has not yet begun.

4.2 The Intangible Capital Stock

The most important insight from the J-curve framework is that the investment phase is not wasted. It is building an intangible capital stock that will generate returns when the deployment phase begins. This stock consists of several components:

Workflow capital: The redesigned processes, agent configurations, and system integrations that allow AI to operate within an organization's specific context. This is the organizational equivalent of the single-story factory layout. It cannot be purchased off the shelf. It must be built.

Training capital: The human knowledge of how to work effectively with AI systems: how to prompt, how to evaluate outputs, how to identify when AI is wrong, and how to escalate appropriately. This knowledge is accumulated through practice and cannot be transferred by instruction alone.

Data capital: The proprietary datasets, feedback loops, and fine-tuning investments that make AI systems more effective in a specific organizational context. A model fine-tuned on a company's historical contracts is more valuable than a general-purpose model. The fine-tuning investment is intangible capital.

Agent infrastructure: The network of AI agents, tools, and integrations that allow AI to operate across an organization's systems. C.H. Robinson's 30-plus AI agents represent years of infrastructure investment. The value of that infrastructure is not the agents themselves but the network effects between them.

The accumulation of this intangible capital stock is not visible in current productivity statistics. It will become visible when the stock is large enough to generate measurable output gains, which is the definition of the J-curve inflection point.

4.3 The Taylorist Transition

The on-the-job leisure phenomenon, workers using AI to complete their assigned work faster and then filling the remaining time with non-work activities, is real and observable. The Atlassian (2025) survey found that 99 percent of developers report time savings from AI, and 68 percent report saving 10 or more hours per week. These hours are not showing up as additional output. They are being absorbed by organizational inefficiency, leisure, or both.


This is not a new phenomenon. Frederick Winslow Taylor documented an identical pattern in early 20th-century factories, which he called "soldiering": workers deliberately working below their capacity to avoid having their piece-rates cut. The solution was not moral exhortation. It was the assembly line, which made individual speed-hiding impossible by setting the pace of work at the system level rather than the individual level.

The AI equivalent of the assembly line is the agentic workflow: a system in which AI agents route tasks, monitor execution, and set the pace of work, making it impossible for individual workers to absorb AI-generated time savings without the savings appearing in system-level output metrics. This transition from passive AI tools (where the worker controls the prompt and the speed) to active orchestration layers (where the system routes tasks and monitors execution) is the organizational redesign that will unlock the macro productivity gains.

The transition is underway. It is not complete. Most enterprises are still in the passive-tool phase. The firms that complete the transition to active orchestration first will show the productivity gains first, and the sector-level data will begin to reflect this divergence before the aggregate data does.

5. The Measurement Blind Spot. Deflationary Expansion and Digital Dark Matter.

5.1 The GDP Accounting Problem



THE MICRO-TO-MACRO PRODUCTIVITY GAP: TASK-LEVEL AI GAINS VS. AGGREGATE TFP TREND

Gross Domestic Product measures the market value of all goods and services produced in an economy. When the price of a good or service falls, GDP falls, even if the quantity and quality of the good or service are unchanged or improved. This is not a flaw in GDP accounting. It is a design feature. GDP is a measure of market value, not of welfare or utility.

The problem arises when the price of a good or service falls toward zero. A legal department that previously processed 100 contracts per month at a cost of \$500 per contract now processes 10,000 contracts per month at a cost of \$5 per contract. The total cost has fallen from \$50,000 to \$50,000, unchanged. But the output has increased by a factor of 100. In GDP accounting, this looks like flat productivity. In any meaningful sense, it is a 100x productivity gain.

This is the deflationary expansion paradox. AI is collapsing the cost of cognitive tasks toward near-zero. The value delivered to organizations and individuals is increasing dramatically. The measured price of that value is falling. GDP accounting will capture the falling price and miss the rising value.

The paradox is not hypothetical. It is already observable in sectors where AI has been most deeply integrated. The cost of generating a first draft of a legal document has fallen by approximately 80 percent since 2023. The cost of generating a software function has fallen by approximately 60 percent. The cost of generating a customer service response has fallen by approximately 70 percent. These cost reductions are real and measurable. They are not showing up as productivity gains in the BLS data. They are showing up as price deflation in the relevant service categories.

5.2 Digital Dark Matter

The deflationary expansion paradox is a specific instance of a broader measurement problem that economists have called "digital dark matter": the economic value generated by digital goods and services that is not captured in GDP because it is provided at zero or near-zero marginal cost.

Erik Brynjolfsson and colleagues estimated in 2019 that the consumer surplus generated by free digital services (search, social media, navigation) amounted to approximately \$2.5 trillion per year in the United States, equivalent to roughly 11 percent of GDP, and was entirely invisible to national accounts. The AI equivalent of this dark matter is larger, because AI is not just generating consumer surplus. It is generating producer surplus: reducing the cost of production in ways that increase output without increasing measured price.

The measurement problem is not merely academic. It has direct implications for monetary policy, fiscal policy, and investment allocation. If the productivity gains from AI are real but unmeasured, then policymakers who rely on measured productivity to set interest rates, project tax revenues, and allocate public investment will systematically underestimate the economy's productive capacity. The policy errors that result from this underestimation could be substantial.

5.3 The Right Instruments

The instruments currently used to measure productivity were designed for an economy where cognitive inputs had stable, positive costs. They are not adequate for an economy where those costs are collapsing. The following instruments are better suited to tracking AI-driven productivity gains in the current environment:

Revenue per employee at AI-adopting firms: This measure captures the output gains from AI adoption at the firm level, net of the cost reductions. Firms that have completed organizational redesign should show rising revenue per employee even as their headcount is flat or declining.

Output quality-adjusted price indices: Standard price indices measure the cost of a unit of output without adjusting for quality improvements. A legal document generated in 30 minutes by an AI-assisted lawyer is not the same product as a legal document generated in 8 hours by an unassisted lawyer, even if the market price is the same. Quality-adjusted indices would capture the productivity gain that standard indices miss.

Task completion rates per worker-hour: Rather than measuring output value (which is subject to price deflation), this measure tracks the number of tasks completed per unit of labor input. It is a direct measure of the task-level productivity gains that the micro-level studies document.

Intangible capital investment as a leading indicator: The current level of AI-related intangible capital investment is a leading indicator of future productivity gains, just as electrification investment was a leading indicator of the 1920s productivity boom. Tracking the accumulation of this capital stock provides advance warning of the coming productivity acceleration.

6. The International Dimension. The US-EU Adoption Divergence.

6.1 The Adoption Gap

The Bick et al. (2026) paper, published as a Brookings Papers on Economic Activity study and simultaneously as a CEPR VoxEU column, provides the most comprehensive current data on the US-EU AI adoption divergence.

In January-February 2026, 43 percent of US workers reported using generative AI for work. The comparable figure across six European countries (UK, Germany, France, Italy, Sweden, Netherlands) was 32 percent, ranging from 26 percent in Italy to 36 percent in the UK. The intensive margin gap is even larger: US workers spend 5.2 percent of their work hours using AI, compared to less than one-third of that figure in Germany, France, and Italy.

At the firm level, the gap is similarly pronounced. Approximately 34 percent of US firms use AI for producing goods and services, compared to 20 percent on average across 32 European countries. The range within Europe is wide: from over 35 percent in Scandinavian countries to less than 10 percent in Eastern European countries.

COUNTRY/REGION	WORKER AI ADOPTION (2026)	FIRM AI ADOPTION (2025)
United States	43%	~34%
United Kingdom	36%	~28%
Sweden	~35%	~38%
Netherlands	~33%	~36%
Germany	~30%	~22%
France	~29%	~18%
Italy	26%	~15%
EU Average (32 countries)	~32%	20%

Sources: Bick et al. (2026); EU-ICT-Firm Survey (2025)

6.2 The Management Practices Explanation

The most striking finding in the Bick et al. paper is that compositional differences (workforce age, education, industry mix) explain only about half of the US-EU adoption gap. The other half is explained by management practices, specifically whether employers actively encourage AI use and provide AI tools and training.

A one standard deviation increase in a worker's management index (measuring performance incentives and employer encouragement) is associated with a 9.6 percentage point higher AI adoption rate. Countries where firms actively encourage AI use show adoption rates 20 to 30 percentage points higher than countries where firms are passive or resistant.

The implication for productivity is direct. The US productivity advantage from AI adoption is not primarily a function of the US workforce being smarter, younger, or more technically sophisticated. It is a function of US management practices being more aggressive in deploying AI and more willing to redesign workflows around it. This is a management problem, not a technology problem. And management problems are, in principle, solvable.

6.3 The ICT Precedent and the Coming Divergence

The US-EU AI adoption gap closely mirrors the US-EU ICT adoption gap of the 1990s. From 1995 to 2025, US productivity grew by approximately 90 percent, compared to approximately 30 percent in the euro area. A substantial portion of this divergence has been attributed to faster ICT diffusion in the US.

If the AI adoption gap follows the same trajectory as the ICT gap, the productivity divergence between the US and Europe will widen significantly over the next decade. The Bick et al. paper finds that a 10-percentage-point increase in firm-level AI adoption is associated with approximately 0.6 percentage points of additional annualized productivity growth per year. If the current US-EU adoption gap of approximately 14 percentage points persists, it implies an annual productivity growth differential of approximately 0.8 to 0.9 percentage points, compounding to a substantial divergence over a decade.

The silver lining identified in the paper is that the gap is not inevitable. It is primarily a product of management practices, which can change. European firms that actively encourage AI adoption and redesign workflows around it will close the gap. The question is whether they will do so before the divergence becomes self-reinforcing.

7. Sector-Level Exposure. Where the Gains Will Appear First.

7.1 The Concentration Problem

The current US productivity data shows a pattern that is consistent with the early stages of a technology-driven productivity boom: the gains are real but highly concentrated. The San Francisco Fed's May 2026 analysis of BLS industry-level data identifies three sectors that account for the bulk of the current productivity outperformance relative to pre-AI trends: Information, Retail Trade, and Professional and Technical Services (PSTS).

This concentration is not surprising. These are the sectors with the highest AI adoption rates, the most digitized workflows, and the most organizational flexibility to redesign around AI capabilities. They are also the sectors where the measurement instruments are most adequate, where output is most easily quantified and where price deflation is most easily tracked.

The sectors not yet showing productivity gains (Healthcare, Education, Government, Construction) are characterized by high regulatory burden, strong labor protections, complex accountability structures, and workflows that are difficult to digitize. These are not sectors where AI cannot eventually drive productivity gains. They are sectors where the organizational redesign bottleneck is most severe.

7.2 Sector-Level Exposure Matrix

The following matrix assesses the AI productivity exposure of major US economic sectors across four dimensions: current AI adoption rate, workflow redesign feasibility, measurement adequacy, and estimated time to visible TFP impact.

SECTOR	AI ADOPTION	REDESIGN FEASIBILITY	MEASUREMENT	EST. TFP IMPACT
Information / Tech	Very High	High	Good	2025–2026 (underway)
Financial Services	High	High	Good	2026–2027
Professional Services	High	Medium	Moderate	2026–2028
Retail Trade	Medium-High	High	Good	2026–2027
Logistics / Transport	Medium-High	High	Good	2026–2027 (C.H. Robinson model)
Healthcare	Medium	Low	Poor	2029–2032
Legal Services	Medium	Medium	Poor	2027–2029
Education	Low-Medium	Low	Poor	2030+
Construction	Low	Very Low	Poor	2030+
Government	Low	Very Low	Poor	2032+

The sectors at the top of this matrix are where the productivity gains will appear first in the data. The sectors at the bottom are where the gains will be largest in absolute terms, because they have the most inefficiency to eliminate, but where the organizational redesign bottleneck is most severe.

7.3 The Financial Services Case

Financial services is the sector most likely to show the next major productivity inflection after Information. The sector combines high AI adoption rates, highly digitized workflows, strong measurement infrastructure, and significant organizational flexibility.

JPMorgan Chase has publicly committed to deploying AI across 1,000 use cases by 2026, with a particular focus on in-house tool development tailored to employee workflows. The firm's AI-assisted contract review system (COIN) processes in seconds what previously took lawyers 360,000 hours per year. This is not a task speedup. It is a workflow redesign that has eliminated a category of work entirely.

The financial services sector also benefits from a measurement advantage: its output is primarily financial, and financial output is well-measured. When a bank's AI systems allow it to process more transactions, underwrite more loans, or generate more trading revenue with the same headcount, the productivity gain appears directly in revenue-per-employee statistics. This is why financial services is likely to be one of the first sectors to show visible TFP acceleration in the BLS data.

7.4 The Healthcare Paradox

Healthcare is the sector with the most to gain from AI productivity improvements and the least likely to show those gains in the near term. The reasons are structural.

AI diagnostic tools have demonstrated accuracy rates that match or exceed specialist physicians in specific domains: dermatology, radiology, ophthalmology, and pathology. The productivity implications are enormous: a radiologist who can review twice as many scans per day, with higher accuracy, is a genuine productivity gain. But the productivity gain does not appear in measured output unless the healthcare system is redesigned to capture it, which requires changes to billing codes, liability frameworks, regulatory approval processes, and clinical workflow design that are measured in years, not months.

The healthcare sector also faces the most severe version of the accountability bottleneck. AI systems cannot legally own liability for diagnostic errors. Every AI-assisted diagnosis requires a human physician to review and sign off. Until the liability framework changes, which requires legislative action in most jurisdictions, the human review step cannot be eliminated, and Amdahl's Law limits the productivity gain to the fraction of the physician's time that is not consumed by the review step.

8. Actor Mapping. Who Controls the Redesign.

8.1 The Four Actor Classes

The organizational redesign that will unlock AI's macro productivity gains is not a single event. It is a distributed process involving four distinct actor classes, each with different incentives, capabilities, and timelines.

Frontier AI Labs: OpenAI, Anthropic, Google DeepMind, and Meta AI control the foundational models that underpin most enterprise AI deployments. Their incentive is to demonstrate productivity gains that justify continued investment and validate their valuations ahead of IPO. Their capability is to improve model performance and reduce inference costs. Their timeline is measured in months.

Enterprise Software Platforms: Salesforce, ServiceNow, Workday, SAP, and Microsoft are the intermediaries through which AI capabilities reach most enterprise workflows. Their incentive is to embed AI deeply enough in existing workflows that switching costs prevent customer churn. Their capability is to integrate AI into the systems of record that govern enterprise operations. Their timeline is measured in years.

Consulting and Systems Integration Firms: McKinsey, Accenture, Deloitte, and their peers are the organizational redesign specialists. Their incentive is to generate billable hours from the transformation process. Their capability is to redesign workflows, manage change, and train workforces. Their timeline is measured in years.

Enterprise Leadership: CEOs, CFOs, and COOs are the ultimate decision-makers about whether and how to redesign organizational workflows around AI. Their incentive is to improve financial performance. Their capability is to allocate capital and direct organizational change. Their timeline is measured in quarters and years.

8.2 The Bottleneck Actor

Of these four actor classes, enterprise leadership is the bottleneck. Frontier labs can improve models. Enterprise software platforms can embed AI in workflows. Consulting firms can design the redesign. None of this matters unless enterprise leadership decides to commit to the organizational transformation required to capture the gains.

The evidence suggests that this commitment is unevenly distributed. C.H. Robinson's Dave Bozeman represents one end of the spectrum: a CEO who has spent four years systematically redesigning his organization's workflows around AI, tracking real-time KPIs, and holding the organization accountable for measurable output gains. The result is 40 percent productivity improvement and eight consecutive quarters of market outperformance.

At the other end of the spectrum are the organizations described in the Joe Reis analysis: enterprises that have purchased Copilot subscriptions, mandated their use, and declared themselves AI companies, without redesigning a single workflow. These organizations are bolting electric motors onto steam-engine factories. The motors work. The factory layout is unchanged. The productivity gains are not coming.

The distribution of enterprise leadership commitment is the primary variable in the productivity timeline. If the C.H. Robinson model diffuses rapidly across industries, the productivity inflection will arrive in the 2027–2028 window. If the Copilot-mandate model predominates, the inflection will be delayed by several years.

8.3 The New Entrant Threat

The historical record is clear on one point: the firms that capture the largest productivity gains from general-purpose technologies are often not the incumbents. They are the new entrants who build the single-story factory from scratch, without the organizational inertia, legacy systems, and institutional assumptions of the incumbents.

The AI equivalent of the single-story factory is the AI-native firm: a company built from the ground up with AI as a core component of its operational architecture, rather than as a tool bolted onto legacy processes. These firms do not need to redesign their workflows because their workflows were designed for AI from the beginning.

The productivity implications are significant. AI-native firms in financial services, legal services, healthcare administration, and logistics are already operating at productivity levels that legacy incumbents cannot match without fundamental organizational transformation. The competitive pressure from these new entrants is one of the forces that will eventually compel legacy incumbents to complete the organizational redesign, or be displaced.

9. The Scenario Matrix. Three Paths to 2028.

9.1 Scenario Architecture

The following three scenarios represent the range of plausible outcomes for AI-driven productivity growth through 2028. They are not equally likely. The probabilities assigned reflect the current state of the evidence and the historical base rates for technology-driven productivity transitions.

9.2 Scenario A: The Fast Path (Probability: ~25%)

Conditions: Enterprise leadership commitment to organizational redesign accelerates rapidly, driven by competitive pressure from AI-native entrants. Enterprise software platforms successfully embed agentic workflows into existing systems of record. The on-the-job leisure leak is resolved by the transition to active orchestration layers. Measurement instruments are updated to capture quality-adjusted output gains.

Timeline: Sector-level TFP acceleration visible in Information and Financial Services by Q4 2026. Aggregate TFP acceleration visible in BLS data by Q2 2027. GDP growth revised upward by 0.5 to 1.0 percentage points annually by 2028.

Markers: C.H. Robinson-style productivity announcements from major financial services and professional services firms. Revenue-per-employee ratios at AI-adopting firms diverging sharply from industry averages. BLS sector-level productivity data showing sustained above-trend growth in Information and Financial Activities.

Risks: Measurement instruments may not capture the gains quickly enough to appear in official statistics. Regulatory responses to AI-driven job displacement may slow adoption. Model capability plateaus may limit the scope of workflow re-design.

9.3 Scenario B: The Base Path (Probability: ~40%)

Conditions: Enterprise leadership commitment is uneven. Some sectors (Information, Financial Services, Logistics) complete organizational redesign on schedule. Others (Healthcare, Education, Government) are delayed by regulatory and institutional bottlenecks. The US-EU adoption gap persists, creating a divergence in national productivity trajectories.

Timeline: Sector-level TFP acceleration visible in leading sectors by Q2 2027. Aggregate TFP acceleration visible in BLS data by Q4 2027 to Q2 2028. GDP growth revised upward by 0.3 to 0.6 percentage points annually by 2029.

Markers: Widening divergence in revenue-per-employee between AI-adopting and non-adopting firms within the same industry. Challenger Gray & Christmas AI-cited layoff data continuing to accelerate through 2026. New business formation in AI-native categories accelerating.

Risks: The J-curve trough may be deeper and longer than modeled if the organizational redesign bottleneck proves more durable. The deflationary expansion paradox may cause GDP accounting to understate the gains even when they arrive. A recession or geopolitical shock could disrupt the investment cycle.

9.4 Scenario C: The Slow Path (Probability: ~35%)

Conditions: Organizational inertia proves more durable than historical analogues suggest. Enterprise software platforms fail to embed AI deeply enough to drive workflow redesign. The on-the-job leisure leak persists, absorbing productivity gains before they appear in output statistics. Regulatory responses to AI-driven displacement slow adoption in key sectors.


Timeline: No aggregate TFP acceleration visible in BLS data before 2030. Sector-level gains in Information and Financial Services visible by 2028, but insufficient to move the aggregate. GDP growth remains near trend through 2028.

Markers: Continued flat or declining TFP in BLS multifactor productivity data. Widening gap between enterprise survey productivity reports and output-based measures. Increasing regulatory scrutiny of AI deployment in healthcare, legal, and financial services.

Risks: The slow path is not a failure scenario. It is the scenario in which the productivity gains are real but arrive later than the fast or base paths. The gains do not disappear; they are delayed. The cost of the slow path is the opportunity cost of the delay: the economic value that would have been generated in the years between the fast path and the slow path.

10. The Hari Causal Chain. Five Links to the Inflection.

10.1 Framework



US TOTAL FACTOR PRODUCTIVITY BY BUSINESS CYCLE: THE POST-2005 STAGNATION AND PROJECTED INFLECTION

The Hari v5 Causal Chain Projection maps the sequence of observable events that must occur for the projected outcome to materialize. Each link in the chain has an assigned probability based on historical base rates, current evidence, and institutional clock-speeds. The joint probability is the product of the individual link probabilities, adjusted for conditional dependencies.

Projected Outcome: Visible acceleration in US aggregate total factor productivity (BLS multifactor productivity data) by Q2 2028, representing a structural break from the post-2005 trend of approximately 0.4 percent annual TFP growth.

10.2 The Five Links

Link 1: The Token Reckoning (Q3 2026) Probability: 80%

Enterprise AI spending reaches an inflection point at which the cost of unstructured, tool-level AI use becomes visible to CFOs. The "tokenmaxxing" phenomenon, in which companies burn their entire annual AI budget in the first months of the year through uncontrolled model usage, forces a transition from ad hoc AI tool deployment to structured agentic workflow design. This is the organizational equivalent of the moment when factory owners realized that bolting electric motors onto steam-engine layouts was not generating the expected returns.

Observable markers: Enterprise AI budget governance frameworks becoming standard practice. Major enterprise software platforms releasing agentic workflow products with measurable ROI tracking. CFO commentary on AI ROI shifting from aspirational to accountability-focused.

Link 2: The Workflow Redesign Wave (Q4 2026 – Q2 2027) Probability: 65%

The Token Reckoning triggers a wave of organizational workflow redesign in the sectors with the highest AI adoption rates and the most organizational flexibility: Information, Financial Services, and Logistics. This redesign is not incremental. It involves the elimination of entire categories of work (the accounts payable room, the contract review team, the first-level customer service queue) and the reallocation of the humans who performed that work to higher-judgment functions.

Observable markers: Challenger Gray & Christmas AI-cited layoff data accelerating through Q4 2026. Revenue-per-employee ratios at AI-adopting firms in Information and Financial Services diverging sharply from industry averages. C.H. Robinson-style productivity announcements from major firms in these sectors.

Link 3: The Measurement Lag Resolution (Q1 – Q3 2027) Probability: 55%

The BLS begins to capture the productivity gains from the Workflow Redesign Wave in sector-level output data. This requires two things: that the gains are large enough to appear in output statistics despite the deflationary expansion effect, and that the BLS measurement instruments are adequate to capture quality-adjusted output improvements. The probability here is lower than for Links 1 and 2 because measurement lag resolution is partly outside the control of any single actor.

Observable markers: BLS sector-level productivity data for Information and Financial Activities showing sustained above-trend growth. Academic papers using firm-level data to document the productivity gains from organizational redesign. BEA consideration of quality-adjusted price indices for AI-intensive service categories.

Link 4: The Aggregate Signal (Q4 2027 – Q1 2028) Probability: 50%

The sector-level productivity gains in leading sectors become large enough to produce a visible signal in aggregate BLS multifactor productivity data. This requires that the leading sectors account for a sufficient share of aggregate output to move the needle, and that the lagging sectors (Healthcare, Education, Government) do not offset the gains with continued stagnation.

Observable markers: BLS multifactor productivity data for the private business sector showing above-trend growth for two or more consecutive quarters. Federal Reserve commentary acknowledging AI-driven productivity acceleration. Upward revisions to potential GDP estimates by CBO and Fed.

Link 5: The Deflationary Expansion Recognition (Q2 2028) Probability: 40%

The economic and policy community recognizes that the measured productivity gains are understating the true gains due to the deflationary expansion effect. New measurement frameworks are proposed or adopted that capture quality-adjusted output improvements in AI-intensive service categories. This recognition is the final link because it is the most dependent on intellectual and institutional change, which is the slowest-moving variable in the chain.

Observable markers: BEA or BLS announcement of new measurement frameworks for AI-intensive services. Academic consensus forming around the deflationary expansion paradox. Policy discussions about the implications of unmeasured productivity for monetary and fiscal policy.

10.3 Joint Probability

LINK	INDIVIDUAL PROBABILITY	CUMULATIVE JOINT PROBABILITY
Link 1: Token Reckoning	80%	80%
Link 2: Workflow Redesign Wave	65%	52%
Link 3: Measurement Lag Resolution	55%	29%
Link 4: Aggregate Signal	50%	14%
Link 5: Deflationary Expansion Recognition	40%	6%

The joint probability of the full causal chain completing on schedule, with all five links materializing in the projected timeline, is approximately 6 percent. This is not the probability that the productivity gains will materialize. It is the probability that they will materialize *and* be recognized *and* be measured *and* be attributed correctly *and* generate a policy response, all within the projected timeline.

The probability that the productivity gains materialize in some form by 2028, even if they are not fully measured or recognized, is significantly higher: approximately 40 percent. The remaining 60 percent is distributed across scenarios in which the timeline extends (most likely), the gains are real but unmeasured (possible), or the organizational redesign bottleneck proves more durable than modeled (less likely but non-trivial).

11. The Leading Indicators Dashboard.

The following indicators are proposed as a practical monitoring framework for tracking progress through the causal chain. They are designed to bypass the measurement blind spots of the BLS and provide advance warning of the coming productivity acceleration.

11.1 Primary Indicators

Revenue-per-Employee at AI-Adopting Firms *Rationale:* Captures firm-level productivity gains net of cost reductions. A rising ratio at AI-adopting firms, diverging from industry averages, is the clearest signal that organizational redesign is generating measurable output gains. *Data source:* SEC filings, earnings calls, Bloomberg financial data. *Threshold:* Sustained divergence of 10+ percentage points between AI-adopting and non-adopting firms within the same industry over four or more consecutive quarters.

Challenger Gray & Christmas AI-Cited Layoff Data *Rationale:* The most direct current measure of AI-driven workforce displacement. Three consecutive months as the leading layoff attribution category (achieved as of May 2026) indicates that the workflow redesign wave has begun. Acceleration through Q4 2026 would confirm Link 2. *Data source:* Challenger Gray & Christmas monthly reports. *Threshold:* AI-cited cuts exceeding 50 percent of total monthly cuts for three or more consecutive months.

Enterprise Software Platform Agentic Workflow Adoption *Rationale:* The transition from passive AI tools to active orchestration layers is the organizational equivalent of the unit-drive transition in electrification. Adoption of agentic workflow products by major enterprise software platforms is a leading indicator of the workflow redesign wave. *Data source:* Salesforce, ServiceNow, Workday, SAP earnings calls and product announcements. *Threshold:* Agentic workflow products generating more than 10 percent of revenue for two or more major enterprise software platforms.

New Business Formation in AI-Native Categories *Rationale:* AI-native firms are the new entrants who build the single-story factory from scratch. Accelerating new business formation in AI-native categories is a leading indicator of competitive pressure on incumbents to complete organizational redesign. *Data source:* Census Bureau Business Formation Statistics; Crunchbase; PitchBook. *Threshold:* New business formation in AI-native categories (AI services, AI-assisted professional services, AI-native logistics) growing at 20+ percent annually for two or more consecutive years.

11.2 Secondary Indicators

BLS Sector-Level Productivity Data for Information and Financial Activities *Rationale:* These are the sectors most likely to show visible TFP acceleration first. Sustained above-trend growth in these sectors is a leading indicator of the aggregate signal. *Threshold:* Annual TFP growth of 2+ percent in both sectors for two or more consecutive years.

Upwork / Freelance Platform High-Value Contract Ratio *Rationale:* The ratio of high-value, long-engagement contracts to low-value, task-based contracts on freelance platforms is a proxy for the emergence of Orchestration Work, the new layer of human contribution that AI cannot automate. A rising ratio indicates that the new work is forming. *Threshold:* High-value contracts (>\$10,000) as a share of total platform value rising above 60 percent.

LinkedIn Skill Re-Titling Curves *Rationale:* The emergence of new job titles incorporating AI-orchestration language ("AI Operations Manager," "Workflow Architect," "Agent Systems Lead") is a leading indicator of the new work layer forming in the labor market. *Threshold:* AI-orchestration job titles appearing in more than 5 percent of new knowledge-work job postings.

12. Open Questions and Intellectual Honesty.

This report has made a series of confident projections. Confidence is not certainty. The following are the most important ways this analysis could be wrong.

12.1 The Orchestration Layer May Be Too Small

The central argument of this report, and of the companion new-work analysis, is that AI will create a new layer of human contribution (Orchestration Work) that absorbs the workers displaced from the tasks AI automates. This argument has a mathematical vulnerability: the new layer may be too small to absorb the displaced workers.

If AI automates 30 percent of the tasks currently performed by knowledge workers, and the new orchestration layer requires only 10 percent as many workers as the automated tasks employed, the net employment effect is negative. The productivity gains are real. The new work is real. But the gap between the two is a structural unemployment problem, not a temporary transition.

The historical record does not resolve this question definitively. The electrification transition created more jobs than it destroyed, but the new jobs were in different sectors, in different locations, and required different skills than the old jobs. The transition was painful for the workers caught in the gap, even if the aggregate outcome was positive. The same may be true for the AI transition.

12.2 The Measurement Problem May Be Permanent

The deflationary expansion paradox may not be resolved by better measurement instruments. It may be a permanent feature of an economy in which cognitive tasks have near-zero marginal cost. If this is the case, GDP will permanently understate the economy's productive capacity, and the productivity gains from AI will never appear in official statistics. Not because they are not real, but because the instruments are structurally inadequate.

This is not a catastrophic outcome. The gains are real regardless of whether they are measured. But it has significant implications for policy: if policymakers rely on measured productivity to set interest rates, project tax revenues, and allocate public investment, they will systematically underestimate the economy's capacity and make policy errors as a result.

12.3 The Organizational Redesign Bottleneck May Be More Durable Than Modeled

The electrification analogy suggests a lag of approximately 38 years from commercial viability to aggregate productivity acceleration. The ICT analogy suggests approximately 24 years. This report's base case assumes a lag of 8 to 15 years for AI, reflecting the compression of each successive technology transition.

But the compression may not continue. The organizational redesign bottleneck, the need to tear down the old factory and build a new one, is not obviously faster for AI than for electricity. It may be slower, because AI requires changes not just to physical infrastructure and organizational structure but to the cognitive habits, professional identities, and institutional assumptions of the workforce. These are harder to change than factory layouts.

If the organizational redesign bottleneck is as durable for AI as it was for electricity, the aggregate productivity acceleration will not arrive until the late 2030s or early 2040s. This is not a failure of the technology. It is a failure of the organizational transition.

12.4 Model Capability Plateaus May Limit the Scope of Redesign

The entire argument of this report assumes that AI model capabilities will continue to improve at a rate sufficient to drive the workflow redesign wave. If model capabilities plateau, if the current generation of large language models represents the ceiling of what is achievable without a fundamental architectural breakthrough, then the scope of workflow redesign will be limited to the tasks that current models can perform reliably.

Current models are highly capable at bounded, well-defined cognitive tasks: writing, coding, summarizing, classifying, and generating. They are less capable at open-ended reasoning, causal inference, and tasks that require deep domain expertise in novel situations. If this capability gap does not close, the sectors most dependent on open-ended reasoning (healthcare, legal, scientific research) will not be able to complete the organizational redesign that the productivity argument requires.

12.5 The Geopolitical Variable

The AI productivity transition is occurring in a geopolitical environment of increasing fragmentation. Export controls on advanced semiconductors, data localization requirements, and regulatory divergence between the US, EU, and China are all potential constraints on the diffusion of AI capabilities across borders and across sectors.

The US-EU adoption gap documented in the Bick et al. paper is partly a product of regulatory environment: European firms face stricter data protection requirements, more aggressive labor regulations, and a more cautious regulatory posture toward AI deployment. If these constraints tighten, the EU may fall further behind the US in AI-driven productivity growth, widening the divergence that the ICT transition began in the 1990s.

13. Methodological Appendix.

13.1 On the Hari Causal Chain Framework

The Hari v5 Causal Chain Projection framework is designed to produce "roughly right" projections rather than "precisely wrong" ones. The individual link probabilities in Chapter 10 are not derived from formal statistical models. They are informed judgments based on historical base rates, current evidence, and institutional clock-speeds. The joint probability calculation assumes conditional independence between links, which is a simplification: in reality, the links are positively correlated (if Link 1 materializes, Link 2 is more likely), which means the joint probability is somewhat higher than the product of the individual probabilities.

The framework is designed to be updated as new evidence arrives. The leading indicators in Chapter 11 are the primary inputs to the update process. If the Token Reckoning materializes on schedule (Q3 2026), the probability of Link 2 should be revised upward. If it does not, the probability of Link 2 should be revised downward.

13.2 On the Historical Base Rates

The lag estimates for steam (60 years), electricity (38 years), and computing (24 years) are drawn from the academic literature on general-purpose technologies, primarily David (1990), Crafts (2002), and Brynjolfsson and McAfee (2014). These estimates are subject to significant uncertainty: the starting point (commercial viability of the technology) and the ending point (visible aggregate productivity acceleration) are both subject to measurement and interpretation.

The compression hypothesis, that each successive GPT shows a shorter lag, is an empirical regularity, not a law. It is possible that the AI transition will not follow this pattern. The reasons to expect compression (organizational flexibility, capital mobility, intellectual awareness of the pattern) are real but not decisive.

13.3 On the Deflationary Expansion Paradox

The deflationary expansion argument is the most speculative element of this report. The claim that GDP accounting will structurally understate AI-driven productivity gains is defensible in theory but difficult to quantify in practice. The legal department example (100x output at the same measured cost) is illustrative rather than representative: most AI productivity gains are not this dramatic, and many are accompanied by price reductions that do appear in GDP statistics.

The more modest version of the argument, that GDP will *partially* understate the gains because quality improvements in AI-intensive services are not fully captured by standard price indices, is well-supported by the existing literature on digital dark matter (Brynjolfsson et al., 2019) and is the version that should be treated as the working hypothesis.

14. References.

- Acemoglu, D. (2025). "The Simple Macroeconomics of AI." *Economic Policy*.
- Atlassian (2025). *State of Teams Report: AI and Developer Experience*.
- Bick, A., Blandin, A., Deming, D., Fuchs-Schündeln, N., and Jessen, J. (2026). "Mind the Gap: Diverging AI Adoption in Europe and the United States." *Brookings Papers on Economic Activity*. NBER Working Paper 34995.
- Bozeman, D. (2026). Interview with Asutosh Padhi, *The Exchange*, McKinsey & Company, February 19, 2026.
- Brynjolfsson, E. and Hitt, L. (1996). "Paradox Lost? Firm-Level Evidence on the Returns to Information Systems Spending." *Management Science* 42(4): 541–558.
- Brynjolfsson, E., Collis, A., Diewert, W.E., Eggers, F., and Fox, K.J. (2019). "GDP-B: Accounting for the Value of New and Free Goods in the Digital Economy." *NBER Working Paper* 25695.
- Brynjolfsson, E., Li, D., and Raymond, L. (2023). "Generative AI at Work." *NBER Working Paper* 31161.
- Crafts, N. (2002). "The Solow Productivity Paradox in Historical Perspective." *CEPR Discussion Paper* 3142.
- David, P. (1990). "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80(2): 355–361.
- Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. (2023). "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *Harvard Business School Working Paper* 24-013.
- DORA (2025). *State of DevOps Report 2025*.
- Faros AI (2025). *The AI Productivity Paradox Report*.
- GitClear (2026). *Coding on Copilot: 2025 Data Suggests Downward Pressure on Code Quality*.
- Gruda, D. and Aeon, B. (2026). "Seven Myths about AI and Productivity: What the Evidence Really Says." *California Management Review Insight*, October 2025; extended analysis January 2026.
- METR (2025). "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity." *METR Research Report*.
- Noy, S. and Zhang, W. (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381(6654): 187–192.
- Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." *arXiv* 2302.06590.
- Reis, J. (2026). "We're in 1905: Why Electricity (Not Dot-Com) Is the Right AI Analogy." *The Weekend Windup*, April 19, 2026.
- Solow, R. (1987). "We'd Better Watch Out." *New York Times Book Review*, July 12, 1987.
- Stack Overflow (2025). *Developer Survey 2025*.

Upwork Research Institute (2024). *AI at Work: The Great Paradox*.

This report is part of the Expansion Effect Research Series. The companion analyses, "AI Job Displacement: June 2026" and "The Abstraction Layer: AI-Driven Creation of New Work," are available at expansioneffect.com.

The Hari v5 Causal Chain Projection Framework is a proprietary analytical tool developed by the Expansion Effect research team. Probability estimates reflect informed judgment based on historical base rates and current evidence. They are not statistical forecasts and should not be treated as investment advice.