

What Happened to the Hours AI Gave Back

Most firms say their people got faster. Few can say where the time went.

Micah Laughlin, Chief Strategy Officer · September 25, 2026

In short. 71% of AI-using Texas firms told the Dallas Fed in May 2026 that AI raised productivity for the people using it, and 89% of nearly 6,000 executives told the Atlanta Fed and NBER it had no effect on their firm. Stratos Edge Research finds both are true: the gain is real at the desk and invisible on the P&L until a company sets a baseline before go-live and chooses where the freed hours go.

Three things that are all true.

The AI line in the budget is growing. The people using the tools say they are faster. The P&L says nothing changed. Most executives reading this are holding all three at once, and the board has started asking which one is wrong.

All three are true. Most companies are getting a gain from AI and losing track of it. The Federal Reserve's own surveys show where.

The gain is real at the desk. In the Dallas Fed's May 2026 survey of Texas firms using AI, 71% said it had raised productivity for the employees using it. Three-quarters said it had not changed how many workers they need. Two wholesalers in the same release described the mechanism this issue is about, in their own words: growing on current staffing, so hiring is delayed as the business grows. Those two did the thing. The rest reported the productivity and left the hours where they fell.

The gain is invisible at the firm. In the Atlanta Fed and NBER survey of nearly 6,000 executives, 89% said AI had no effect on their firm's productivity over the past three

years. Most of those firms use AI. They are not wrong, and neither are the Texas firms. Dallas asked whether the people using AI are faster than the people who are not. Atlanta asked whether the firm's sales per employee moved. Both answers are true at once because of arithmetic. The St. Louis Fed puts time saved at about 2.2 hours a week for the people using generative AI. In the New York Fed's panel, the median AI-using firm had fewer than one worker in five using it. Run those together and the firm-level effect is about one percent, which no sales-per-employee line will show, however well anyone measures it. Shallow deployment makes the gain small at the firm. The absence of a baseline makes it unprovable at any level.

This issue's position follows from that arithmetic: the team is the unit. Firm-wide, one in five people use the tool. On the team where the tool is the workflow, everyone does. Forty people each freeing two hours a week is two people's worth of capacity, enough to take on more volume, open a program, or change a requisition. The gain exists at the team level, dilutes to nothing at the firm level, and goes unmeasured at the team level.

Almost nobody converts it to a decision. The New York Fed's August 2026 survey is the only public U.S. figure that turns capacity into a staffing decision. Among service firms using AI, 15% had hired fewer people because of it. Thirteen percent had hired more. The rest made no staffing decision with the capacity. The survey cannot see the firm that moved the hours into volume or a backlog without touching headcount, and neither can any other U.S. survey. Note what the 13% did: the capacity made growth possible and they staffed it. That is as much a return as an avoided hire, and only visible as one if someone can show the growth would not have happened otherwise.

The executives selling the story have not banked it either. The St. Louis Fed classified every AI-and-productivity mention on U.S. earnings calls by tense. About 95% describe gains still to come. That share has not moved since 2023. The people with the strongest incentive to claim a realized gain are, on the record, describing a forecast.

Meanwhile the spend is visible. In KPMG's Q2 2026 U.S. pulse, 26% of executives can see their AI costs in real time and two-thirds have a cost dashboard. Cost controls are arriving before any measurement of return, so the organization can say what it spent long before it can say what it got. The requisition gets approved for the same reason: the CFO can see the AI line and cannot see the capacity line, so headcount looks like the only lever that still moves.

Same technology, two outcomes.

Stratos Edge Research reads the three findings as one: most companies that deployed AI got faster at something, and very few of them realized a return from it. The ones who did are not running better models. They measured the gain, decided what it would become, and moved it there on purpose.

That sounds too simple to carry an article. It carries this one because the pattern holds in every dataset that splits companies by whether they measured.

The productivity is the raw material. AI frees hours almost everywhere it is used, and freed hours on their own do nothing for the business; they get absorbed into the workday and nobody can say where they went. The return exists only when the gain has been measured and reallocated with intention to work the business chose. Take on 15% more volume with the same team. Start the program that has been waiting for people. Put the freed hours into accuracy, turnaround, or the customers who were getting the short version. Staff next year's growth without the five hires in the budget. Each of those is a decision, and a decision needs a baseline to be made with any confidence.

The companies that set a baseline can make the decision, name what the capacity became, and show the board the return. The companies that did not are getting the

same productivity gain and no return from it, because nothing was measured and nothing was moved.

That is the whole finding.

The finding, in four numbers.

71%

of AI-using Texas firms say AI raised productivity for the employees using it. The gain is real.

Source: DALLAS FED, MAY 2026

+1.8%

output per worker is what 748 U.S. financial executives reported from AI in 2025. The gain implied by the same firms' own revenue and headcount numbers was "much smaller" in every industry. Real at the desk, missing on the P&L.

Source: ATLANTA FED, MARCH 2026

15% / 13%

of AI-using service firms hired fewer people because of AI. **13%** hired more. About a third retrained. The rest made no staffing decision with the capacity, and no survey asks what else they did with it.

Source: NEW YORK FED, SEPTEMBER 2026

15% vs 3%

is the share reporting established ROI with full visibility into AI costs, against the share without it. The cohort that measured is the cohort that realized a return.

Source: KPMG GLOBAL AI PULSE, Q2 2026;
GLOBAL SAMPLE

The gain shows up in the first number. The return shows up in the last.

The cohort that can see it.

Among U.S. healthcare organizations furthest along in scaling AI, 18% measure AI's effect on revenue or cost against a defined baseline with a named owner, in Deloitte's spring 2026 survey of 64 healthcare CFOs. Among the organizations earlier in scaling, 31% do. The organizations moving fastest are the ones least able to say what they got, and two-thirds of the scalers rely on before-and-after comparisons rather than a defined baseline.

Read the inversion as a mechanism rather than a lapse and it explains itself. Attribution needs a before. The before exists once, in the weeks ahead of the first deployment on a process, and it is cheap then: the numbers are already in the workflow system and nobody has touched the process yet. Scale through that window and it closes. You cannot go back and measure hours per claim in the quarter before the tool went in. The scalers did not skip measurement so much as spend the only window they had. The same mechanism explains why 95% of earnings-call productivity claims have stayed in the future tense for three years: with no before, the claim can never move to the past.

So there are two perishable things in this story, and they expire in opposite directions. The baseline expires at go-live. The freed hours expire within the quarter, absorbed into the workday. The return lives in the window between them. Most organizations reach the end of a deployment holding neither end.

The asset, then, is every process on the roadmap that has not gone live yet. Each one still has a before.

The rest of the evidence points the same way. Wharton found that 72% of U.S. enterprise leaders now track metrics for generative AI tied to profitability, throughput or productivity, and three in four report positive returns. MIT NANDA found that 95% of enterprise gen AI pilots show no measurable P&L impact. Different populations, different definitions of return, and we do not average them.

Read together they say what this issue says: the people who report a return are the people who built the instrument to see one. The cleanest split is global: in KPMG's Q2 2026 pulse of 2,145 senior leaders, organizations with full visibility into AI costs report established ROI at 15%, against 3% for those without it. McKinsey's State of AI found that of twelve adoption practices, tracking well-defined KPIs was the one most correlated with bottom-line impact, and fewer than one in five organizations were doing it.

All of these are self-reported and correlational; the organizations that measure may be further along for other reasons, and Method says so at length. The direction does not reverse, and no survey we found points the other way.

What the capacity became.

The return is realized at the moment the measured gain is reallocated, and the destination is a choice. The requisition is the one boards see first, so it runs through this issue. It is not the only destination, and the evidence says it is not the main one.

The Atlanta Fed survey asked 748 U.S. financial executives why they invested in AI and what they got. Production efficiency and labor productivity ranked highest as motives. Reducing labor cost ranked lower. The benefits firms said they realized in 2025 matched the motives: efficiency, decision speed, output. Headcount effect: negligible. What companies wanted from AI was more done, and that is what the ones who can see it report getting.

The named examples say the same thing. Lemonade's metric is premium per employee, and the story is premium more than doubling while headcount drifted down 6%, not a cut. C.H. Robinson reports shipments per person per day up more than 40%. Oscar Health says it can serve more members without adding headcount

at the same rate. Two Dallas Fed wholesalers described growing on current staffing. Every one of those is capacity turned into volume.

Deloitte's coding of health-system newsrooms found that where value was demonstrated at all, it was more often operational or clinical than financial: turnaround, accuracy, a clinical outcome. Those are destinations too, and they are the ones a CFO cannot see from the finance system.

So the destinations for measured capacity are, in rough order of how often the evidence shows them: more volume on the same team; growth staffed from inside; a program or backlog that had been waiting for people; speed, accuracy or quality on the existing work; and an avoided hire. The order a board asks about them runs the other way. Any of the five is a return if it was chosen and measured. None of them is a return if it happened by drift. An organization that measured can name which destination its capacity went to and put a number on it. An organization that did not has capacity flowing to "other tasks," which is where the Danish register study found 80% of saved time going, unnamed.

What we could not find is independent evidence on the quality and customer-satisfaction destinations. Those figures exist only in vendor case studies, which this issue does not use. The absence is worth stating: the outcomes most companies say they want from AI are the ones nobody outside a vendor has measured.

Two companies, one loop closed.

Named examples in this issue are company-published only: filings, shareholder letters, an executive on the record. We wanted companies under 5,000 employees. We got one closed loop, and it is the contrast with the second company that makes the point.

Lemonade closed the loop. A digital insurer of roughly 1,300 employees, Lemonade defined in-force premium per employee as a named operating metric and has reported it every quarter. It went from about \$0.40 million in Q2 2022 to \$1.07 million in Q2 2026. In-force premium more than doubled since late 2022 with headcount down about 6%. The company set a dated checkpoint the capacity gain is meant to deliver: its first adjusted-EBITDA-positive quarter in Q4 2026. A number was set, the capacity was redeployed instead of hired, and there is a checkpoint on the calendar. That is what a baseline looks like when it is doing its job. Lemonade was built on automation and the metric is its own, so this is not a retrofit onto a legacy operation. Whether it makes the checkpoint is a question with a date on it, which is more than most companies can say.

Duolingo set the rule and skipped the number. In April 2025, Duolingo chief executive Luis von Ahn published a memo: headcount would be added only if a team could not automate more of its work, and AI use would count in performance reviews. No metric was published for either. A year later, on the record, von Ahn said Duolingo had dropped the AI-usage evaluation rule after employees asked whether they were meant to use AI for its own sake. Full-time headcount grew through the period. The rule was the right instinct. Without a baseline, there was nothing to hold it to, and it reversed.

Both companies are AI-native consumer businesses, which is not the reader's situation. The contrast is still the cleanest one on the public record: same conviction, same year, one wrote the number down.

The other failure is the same mistake in the other direction. Commonwealth Bank of Australia announced 45 customer-service redundancies after introducing a voice bot, then reversed and apologized in August 2025 when call volumes rose instead of falling. A headcount decision made without the number, and the number arrived anyway.

SIDEBAR: ENTERPRISE NAMES, FOR FORM NOT SCALE

Walmart said in September 2025 it expects its 2.1 million global headcount to stay flat for three years while revenue grows, with roles changing as AI rolls into customer service, supply chain and back-of-store work. A named line and a date. Not checkable until late 2028, and no interim reporting found.

C.H. Robinson (about 13,000 employees) reported a more than 40% increase in shipments per person per day since 2022 across its North American surface transportation unit, and told investors that automation of the quote-to-cash cycle lets it "decouple headcount growth from volume growth," with average headcount down 12.9% year over year. Over our size line, but it is the plainest statement of the mechanism from a U.S. executive on an earnings call.

Oscar Health (about 2,400 employees) told investors in August 2026 that AI across benefits, billing, claims and member support means it can "serve more members without adding headcount at the same rate," with SG&A ratio at a record low 14.2%. A dollar target exists. A published baseline for the capacity does not. Closest mid-size healthcare example found, and half a loop.

The decision, worked.

The decision as it lands on a desk, with illustrative numbers. A composite, no client.

A specialty services company. About 600 people. A revenue cycle and documentation team of 40. The growth plan for next year needs that team to handle 15% more volume, and the operating budget carries five new positions to do it. AI tooling for claims preparation and documentation went live in Q1.

COMPANY A

Measured before go-live. Hours per claim, claims per person per week, rework rate, denial rate, two months of baseline. Six months in, the team is handling 13% more claims per person at the same rework rate. Thirteen percent of 40 people is about five people of capacity, and the director has a choice about where it goes. She puts three people's worth against the growth plan and two against the denial backlog that has been sitting since last year, with a target on the denial rate. The five requisitions are on the CFO's desk. The answer is now a comparison, not an argument: the capacity the plan needs already exists, measured against a number the team agreed to before anyone had an opinion, and the rest of it has a named destination. The CFO closes three of the five and holds two against a re-measure in Q4. The board sees the AI spend on one line, the avoided cost on the next, and a denial-rate target with a date. The team sees growth staffed from inside and a backlog worked, which is a different message than a headcount cut, and it is true.

COMPANY B

Deployed the same tools on the same date and did not measure. The team is also handling more. Ask the director and she will tell you the tools help. Ask her by how much and she will say it depends. Six months in, the same five requisitions arrive. There is no number to hold them against, so the request is judged on the growth plan alone, and the growth plan says five. The CFO approves five. The AI line and the headcount line both go up. In the next board meeting someone asks what the AI spend returned, and the answer is a productivity anecdote.

Same technology, same team size, same date. The difference is three salaries this year, two more pending a re-measure, a denial-rate target with a date, and a return the board can see. It was decided in the two months before go-live.

Two things to notice. The first is that Company A did not need a sophisticated instrument. Hours per claim and a rework rate are on any operations dashboard. The

second is that nobody at Company A left. The growth got staffed from capacity that was already paid for, and a backlog got worked. The CEO wanted that outcome from AI in the first place, and only the company that can prove the capacity exists and say where it went gets it.

By seat.

SEAT	WHAT CHANGES	FIRST MOVE
CEO	The board question becomes "what did we do with what AI did," and it has no answer without an owner.	Name the owner of the capacity decision before the next deployment, make it the person who owns the baseline, and ask for the destination, not the productivity number.
CFO	Cost visibility is arriving; capacity visibility is not. An approval against a process with no baseline is approved blind.	Every AI approval carries two numbers, the cost and the baseline of the process it touches; pull the baselines for undeployed processes this quarter, while they are free.
IT and operations	The baseline is an operations artifact: hours per unit, units per person, error or rework rate, measured before go-live.	If the vendor's business case does not name the process and the number, it is a forecast. Ask for the number before the contract.

CEO.

The board question is not "what did AI do." It is "what did we do with what AI did." The first question has no good answer without the second. Name the owner of the capacity decision before the next deployment, and make it the same person who owns the baseline. Then ask them for the destination, not the productivity number.

CFO.

Cost visibility is arriving. Capacity visibility is not. Every AI approval should carry two numbers: the cost, and the baseline of the process it touches. A requisition against a process with an AI deployment and no baseline is a requisition being approved blind. Treat it as one. The same test applies to a growth plan, a new program, or a quality target that assumes the capacity is there. And for the processes not yet deployed, the baseline is free this quarter and unavailable next. Pull it now.

IT and operations.

The baseline is an operations artifact, not a finance one. Hours per unit, units per person, error or rework rate, measured for a period before go-live, on the process AI will touch. If the vendor's business case does not name the process and the number, the business case is a forecast. Ask for the number before the contract.

SEGMENT	WHAT THE EVIDENCE SAYS	THE UNIT TO MEASURE
Health systems and specialty providers	Deloitte's spring 2026 survey: the fastest scalers are the least able to attribute a return.	Revenue cycle, prior authorization, documentation and scheduling all have a countable unit already in the workflow system.
Manufacturers	In the New York Fed's panel no manufacturer reported an AI layoff, and one in five had retrained workers.	A quote turned or a defect caught, not headcount.
Professional and financial services	The Atlanta Fed found realized gains concentrated in high-skill services, especially finance, where the reported gain sits furthest from the revenue line.	Billable or processed units per person, before and after.

SIDEBAR: WHERE THIS LANDS BY SEGMENT

Health systems and specialty providers. The Deloitte finding is the segment's own: the fastest scalers are the least able to attribute. Revenue cycle, prior auth, documentation and scheduling all have a countable unit. The baseline exists in the workflow system already. Nobody has pulled it as a baseline.

Manufacturers. In the New York Fed's panel no manufacturer reported an AI layoff and one in five had retrained workers. Capacity is showing up in planning, quality and quoting, where the unit is a quote turned or a defect caught. Measure that, not the headcount.

Professional and financial services. The Atlanta Fed survey found realized gains concentrated in high-skill services, especially finance. That is also where the productivity paradox is largest: the reported gain is furthest from the revenue line. Billable or processed unit per person, before and after, is the instrument.

The capacity test.

The Stratos Edge Research capacity test is the operational close: five questions for any decision that lands on a process where AI has been deployed in the last eighteen months: a requisition, a growth target, a new program, a service-level commitment.

FIVE QUESTIONS

- Step 1.** What was the unit rate before go-live? Hours per unit, units per person, whatever the process counts in. If the answer is "we didn't capture that," the rest of the test is a forecast, and there is no way back to the before. For the next process, this question is the one to answer first.
- Step 2.** What is it now? Same unit, same measurement, at least a quarter of data.
- Step 3.** Where did the difference go? Named work, not "other tasks." More volume, a program, a backlog, a quality target, a hire not made. If nobody assigned it, it is not capacity, it is slack, and it has already been absorbed.
- Step 4.** Does the plan need more of this unit than the measured gain provides? If the gain covers it, the requisition is a choice, not a need. If the gain was sent somewhere else on purpose, the requisition is a need and the board should see both lines.
- Step 5.** When do we re-measure? A date. The Lemonade checkpoint, scaled to the process.

A company that can answer all five has realized a return on its AI spend and can put it in front of a board, whichever destination it chose. A company that cannot has the same productivity gain and no return, and the requisition goes through because it is the only decision left that looks like one. The technology was never the variable. Measurement and intention were.

Where the hours went.

Where did the hours AI gave back go? The closest answer the evidence allows comes from Denmark, the one place anyone has looked with administrative data: 80% of the time AI saved went to "other tasks," unnamed, on no line anyone could point to.

One survey has asked employees what they were told to do with the time. BCG put the question to close to 12,000 employees and managers in more than a dozen countries in 2026. Two-thirds said they received little or no guidance on what to do with the hours AI saved them. More than half said they did not redirect the time to anything more strategic. Global sample, so it sits here as counterweight. We found no U.S. survey that asks the question at all.

So the hours went into the workday, unassigned, at the companies that did not measure and did not say, and into a destination with a name and a number at the companies that did. The difference was whether anyone measured, whether they measured before the before was gone, and whether anyone told the team where the hours were supposed to go. The first two take a quarter of work. The last one takes a sentence.

Method.

What this issue is.

A synthesis of published evidence, with two named company contrasts drawn from company-published materials. Nothing here is a Stratos Edge measurement.

The thesis.

Stratos Edge Research's thesis is that the productivity gain from AI is the input, not the return. The return is realized only when the gain has been measured and reallocated with intention to work the business chose. Organizations that set a baseline before deployment can make that reallocation, prove it, and re-measure. Organizations that did not are getting the same productivity gain and no return from it, so the board sees spend with nothing to show.

What would disprove it.

A U.S. dataset in which organizations with a pre-deployment baseline report booked return at the same rate as organizations without one. Or company-published cases in which capacity was assigned and proven with no baseline in place. If either shows up, it gets the same page count as the finding that confirms us.

What this evidence proves.

Across every dataset we found that splits organizations by a measurement practice, the measured group reports a return at a multiple of the unmeasured group. The Federal Reserve panels show why: capacity gains are real at the individual level, invisible at the firm level, and converted into a staffing decision by about a quarter of firms. The named contrast (Lemonade and Duolingo) shows what a baseline does and what its absence does, in two companies with the same conviction.

What it does not prove.

No U.S. survey splits reported return by whether a baseline existed before deployment. Every cross-tab is contemporaneous; organizations that measure may be further along for other reasons. Direction is inferable, causation is not shown.

The two gaps we hunted.

First, a McKinsey KPI-tracking cross-tab against EBIT impact. It does not exist as a published table. The March 2025 edition reports KPI tracking as the practice most

correlated with EBIT impact in a regression (R^2 of 0.20), and the August 2026 edition reports high performers at twice the rate of defined measurement processes. We use both, labeled global. Second, filings or shareholder letters from U.S. companies under 10,000 employees stating hires avoided against plan, with a number. We found no closed loop in healthcare or manufacturing at that size. Oscar Health and Privia Health have the language without the baseline. C.H. Robinson has the loop and is over the size line. Two sub-10,000 companies (PSQ Holdings, Cloudflare) attribute headcount changes to AI with numbers, but both are workforce reductions, not avoided hires, and are excluded from the body for that reason. Lemonade remains the only sub-5,000 U.S. company with a named metric, a redeployment, and a dated checkpoint.

The destinations we could not source.

Independent evidence that measured AI capacity was redirected to quality, accuracy, or customer-satisfaction outcomes does not exist at the source tiers this issue uses. The Library holds several such cases, all vendor-published. They are excluded. The Deloitte newsroom coding (operational and clinical value more often demonstrated than financial) is the closest independent signal and is used as such.

Source discipline.

U.S. sources carry the body. Three global surveys (KPMG Global, McKinsey, BCG) appear labeled as counterweight where they are used. Vendor-commissioned research is not used as an anchor. Where sources contradict (Wharton against MIT NANDA; Atlanta Fed reported against implied), both are cited beside each other and not averaged. KPMG's U.S. Pulse publishes no established-ROI percentage, so the 7% and the 15%-versus-3% split are global figures.

Non-U.S. evidence held in method only.

Humlum and Vestergaard (NBER, May 2025) linked about 25,000 Danish workers in eleven AI-exposed occupations to administrative earnings and hours records. Users reported saving 2.8% of work hours. Eighty percent put the time into "other tasks,"

unnamed. Three to seven percent of the savings passed through to earnings. Employment and wage bill at high-adoption workplaces showed no differential change. It is the strongest available evidence that saved hours do not become a booked outcome on their own. It is Danish, so it is held out of the U.S. evidence sections and appears once, labeled, in the closing section.

Sample notes and supporting figures.

Atlanta Fed CFO survey (Nov 2025 to Jan 2026): 603 CFOs in The CFO Survey panel plus 145 senior finance executives; reported +1.8% output per worker from AI in 2025 against a "much smaller" gain implied by the same firms' AI-attributed revenue and employment changes. Dallas Fed (May 2026): 313 Texas firms, 203 AI users, skewed to smaller and mid-sized service firms; 76% no change in need for workers, 77% no change in wages for AI-using employees. Atlanta Fed and NBER (Nov 2025 to Jan 2026): firms with 10 or more employees, median about 100, executive self-assessment in five categories; 89% no productivity effect is across all firms, and 78% of U.S. firms in the sample use AI, so at least 86% of users also reported no effect; realized U.S. productivity effect +0.24% over three years, expected +2.25% over the next three; more than 90% no employment effect. St. Louis Fed (Bick, Blandin, Deming): 5.4% of work hours saved among gen AI users, 1.4% across all workers. Census HTOPS (March 2026): among last-week users, 31% saved one to two hours, 15% three to four, 15% more than four, 10% none; the "three hours a week" figure vendors quote is real for the top third and is not the median. New York Fed (Aug 2026): New York and northern New Jersey; median 17% of workers using AI at service adopters; 4% laid off because of AI, about a third retrained, no manufacturer reported an AI layoff. Census BTOS: AI-related employment decreases at 2% of firms; 66% of AI-using firms use it only to augment. St. Louis Fed earnings calls: about 490,000 transcripts from 5,198 U.S. public companies; non-AI productivity talk runs about 75% future tense. KPMG U.S. Q2 2026: 66% cost dashboards, 61% require approvals for AI spend, \$202 million average planned investment, employee resistance 20% from 5%. KPMG Global Q2 2026: CEO clearly accountable 14% established ROI vs 4%; 76% report meaningful value, 7% established ROI. Deloitte: 64 CFOs at health systems over \$1

billion and plans over 500,000 members, mid-market not represented; 68% of scalers rely mainly on before-and-after comparisons; separately, 17,622 newsroom articles from 62 large health systems and plans, demonstrated-value articles 9% (2023) to 19% (May 2026), expected-value framing 35 to 45% throughout. BCG AI at Work 2026: 42% of regular frontline users report saving about eight hours a week; 66% little or no guidance; more than half not redirected. All survey figures are self-reported; none are audited financials.

Corrections.

No corrections have been logged to this issue since publication. A correction is recorded here with the date, the original text, the corrected text, and the reason.

Sources.

1. [KPMG US, "AI Quarterly Pulse Survey," Q2 2026 edition \(last modified June 24, 2026\).](#)
2. [KPMG International, "Global AI Pulse Q2 2026: From deployment to value realization," June 2026, n=2,145.](#)
3. [Yotzov, Barrero, Bloom et al., "Firm Data on AI," NBER Working Paper 34836 / FRB Atlanta WP 2026-3, revised March 2026.](#)
4. [Baslandze, Edwards, Graham et al., "How Might AI Change the Workplace? Evidence from Corporate Executives," FRB Atlanta Policy Hub: Macroblog, March 25, 2026 \(FRB Atlanta WP 2026-04 / NBER 34984\).](#)
5. [Federal Reserve Bank of Dallas, Texas Business Outlook Surveys, Special Questions on AI, May 26, 2026.](#)
6. [Abel, Deitz, Emanuel, Montalbano, "Businesses Are Using AI to Transform Work, Not Cut Jobs," FRB New York Liberty Street Economics, September 1, 2026.](#)

7. [Ozkan, Kalyani, Sullivan, "AI and Productivity: What Firms Are Saying on Earnings Calls," FRB St. Louis On the Economy, July 31, 2026.](#)
8. [U.S. Census Bureau, "The Microstructure of AI Diffusion," CES WP 26-25 \(BTOS AI supplement\), 2026.](#)
9. [Bick, Blandin, Deming, "The Rapid Adoption of Generative AI," FRB St. Louis Working Paper 2024-027, revised 2025.](#)
10. [Stepler, "About a Third of Workers Who Used AI in the Last Week Said They Completed Tasks One to Two Hours Faster," U.S. Census Bureau \(Household Trends and Outlook Pulse Survey, March 2026\), August 11, 2026.](#)
11. [Beauchene, Duranton, Martin, Lyon, Walters, "AI at Work: Strategy Matters More Than Tools," BCG, June 3, 2026 \(~12,000 respondents, global\).](#)
12. [Janisch, Fera, Gosu, Bhatt, Shukla, "Momentum isn't a metric," Deloitte Center for Health Solutions, August 12, 2026.](#)
13. [Wharton Human-AI Research and GBK Collective, "Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise," October 2025 \(more than 800 U.S. enterprise decision-makers, companies over \\$50M revenue\).](#)
14. [MIT NANDA, "The GenAI Divide: State of AI in Business 2025."](#)
15. [McKinsey, "The state of AI: How organizations are rewiring to capture value," March 12, 2025.](#)
16. [McKinsey, "The state of AI in 2026: On the road to ROI," August 25, 2026.](#)
17. [Lemonade, Inc., shareholder letters Q4 2025, Q1 2026, Q2 2026 \(8-K exhibits\) and Q2 2026 investor presentation.](#)
18. [Luis von Ahn, "AI-first" memo, LinkedIn, April 28, 2025; Fortune, April 13, 2026; Rapid Response podcast, May 13, 2026.](#)
19. [Walmart workforce conference, September 26, 2025, as reported by WSJ, Fortune, CNBC, Fox Business.](#)
20. [C.H. Robinson Q4 2025 earnings call, January 28, 2026.](#)
21. [Oscar Health Q2 2026 earnings call, August 6, 2026.](#)
22. [ABC News \(Australia\), "CBA backtracks on AI job cuts as chatbot lifts call volumes," August 21, 2025.](#)

23. Humlum and Vestergaard, "[Large Language Models, Small Labor Market Effects](#)," NBER Working Paper 33777, May 2025.

Cite this issue.

Stratos Edge Research. What Happened to the Hours AI Gave Back. Issue 1, September 2026.
Editor: Micah Laughlin. <https://www.stratosedge.ai/insights/hours-ai-gave-back>

Data appendix: [issue-1-data.csv](#) ([Download the PDF](#)), one row per number in this issue with metric, value, population, sample, source, URL, date and evidence tier. Series methodology and citation forms at [stratosedge.ai/research](https://www.stratosedge.ai/research).

Confidence **Before Technology**

STRATOS EDGE RESEARCH · THE AI ADOPTION INTELLIGENCE COMPANY · STRATOSEDGE.AI

© 2026 Stratos Edge LLC. All rights reserved. This article may be quoted with attribution and a link to the original. It may not be reproduced in full or used to train models without written permission.