The Implications of AI for the Labor Market

Nate Persily

James B. McClatchy Professor of Law, Stanford Law School

Co-Director, Stanford Law AI Initiative

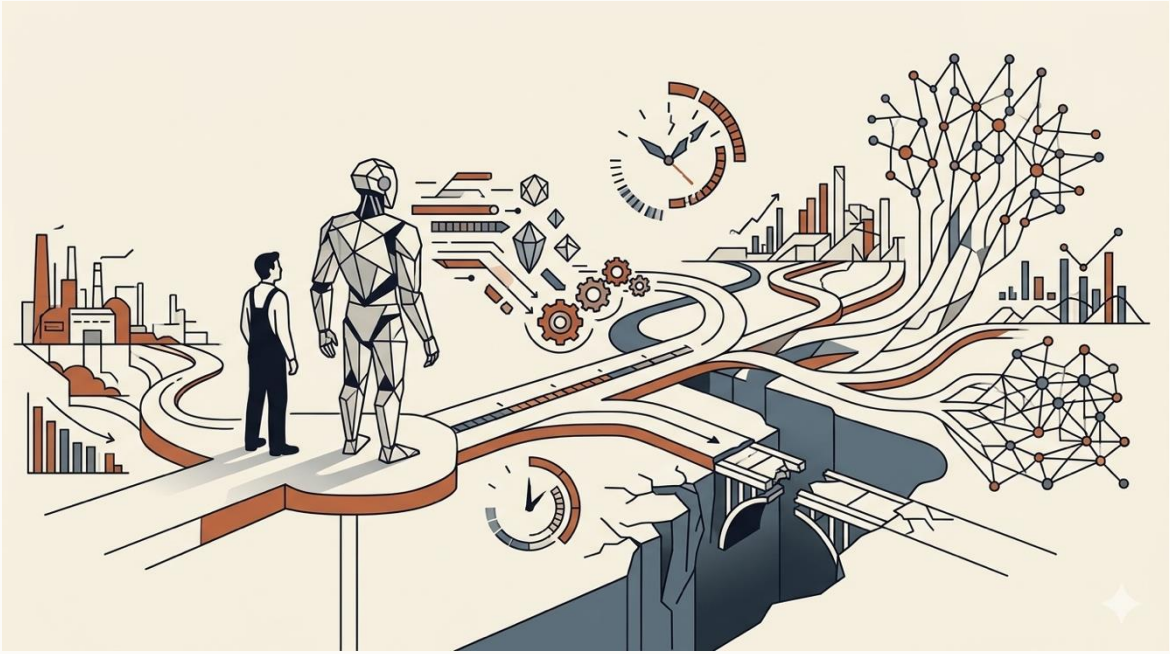
Stanford AI Governance Lab

@Persily · npersily@stanford.edu



Four questions

- Which tasks can AI do, and which is it actually doing?
- Which jobs will shrink, which will grow, and which will change?
- Who bears the costs of the transition?
- How fast, and how would we know?



AI and work: utopia, dystopia, or overhyped?



OPTIMISTIC VIEW · HAL VARIAN

If “displace more jobs” means “eliminate dull, repetitive, and unpleasant work,” the answer would be yes. … **The work week has fallen from 70 hours a week to about 37 hours now, and I expect that it will continue to fall. This is a good thing.** Everyone wants more jobs and less work. Robots of various forms will result in less work, but the conventional work week will decrease, so there will be the same number of jobs (adjusted for demographics, of course).

— Hal Varian, Chief Economist, Google



DYSTOPIAN VIEW · DARIO AMODEI

Anthropic CEO warns AI could eliminate half of all entry-level white-collar jobs

Chris Morris · May 28, 2025





SKEPTICAL VIEW · DARON ACEMOGLU

- **~5%** of U.S. tasks profitably performed by AI within 10 years
- **≤0.66%** TFP gain over a decade; GDP boost closer to **1%**
- **“nontrivial, but modest effect, and certainly much less than … Goldman Sachs and the McKinsey Global Institute”**

Applying AI to Extend Worker Expertise

ECONOMIC POLICY

The Simple Macroeconomics of AI

Daron Acemoglu · 2025

Sources: Hal Varian; Chris Morris (May 28, 2025); Acemoglu, The Simple Macroeconomics of AI (2025).

Known unknowns

The capability frontier
Generative AI, other AI, robotics

The geographic diffusion of AI technology

The speed of development and adoption

The future policy landscape

The likelihood of AGI or an intelligence explosion

Some general themes

Automation v. augmentation

Jobs v. tasks

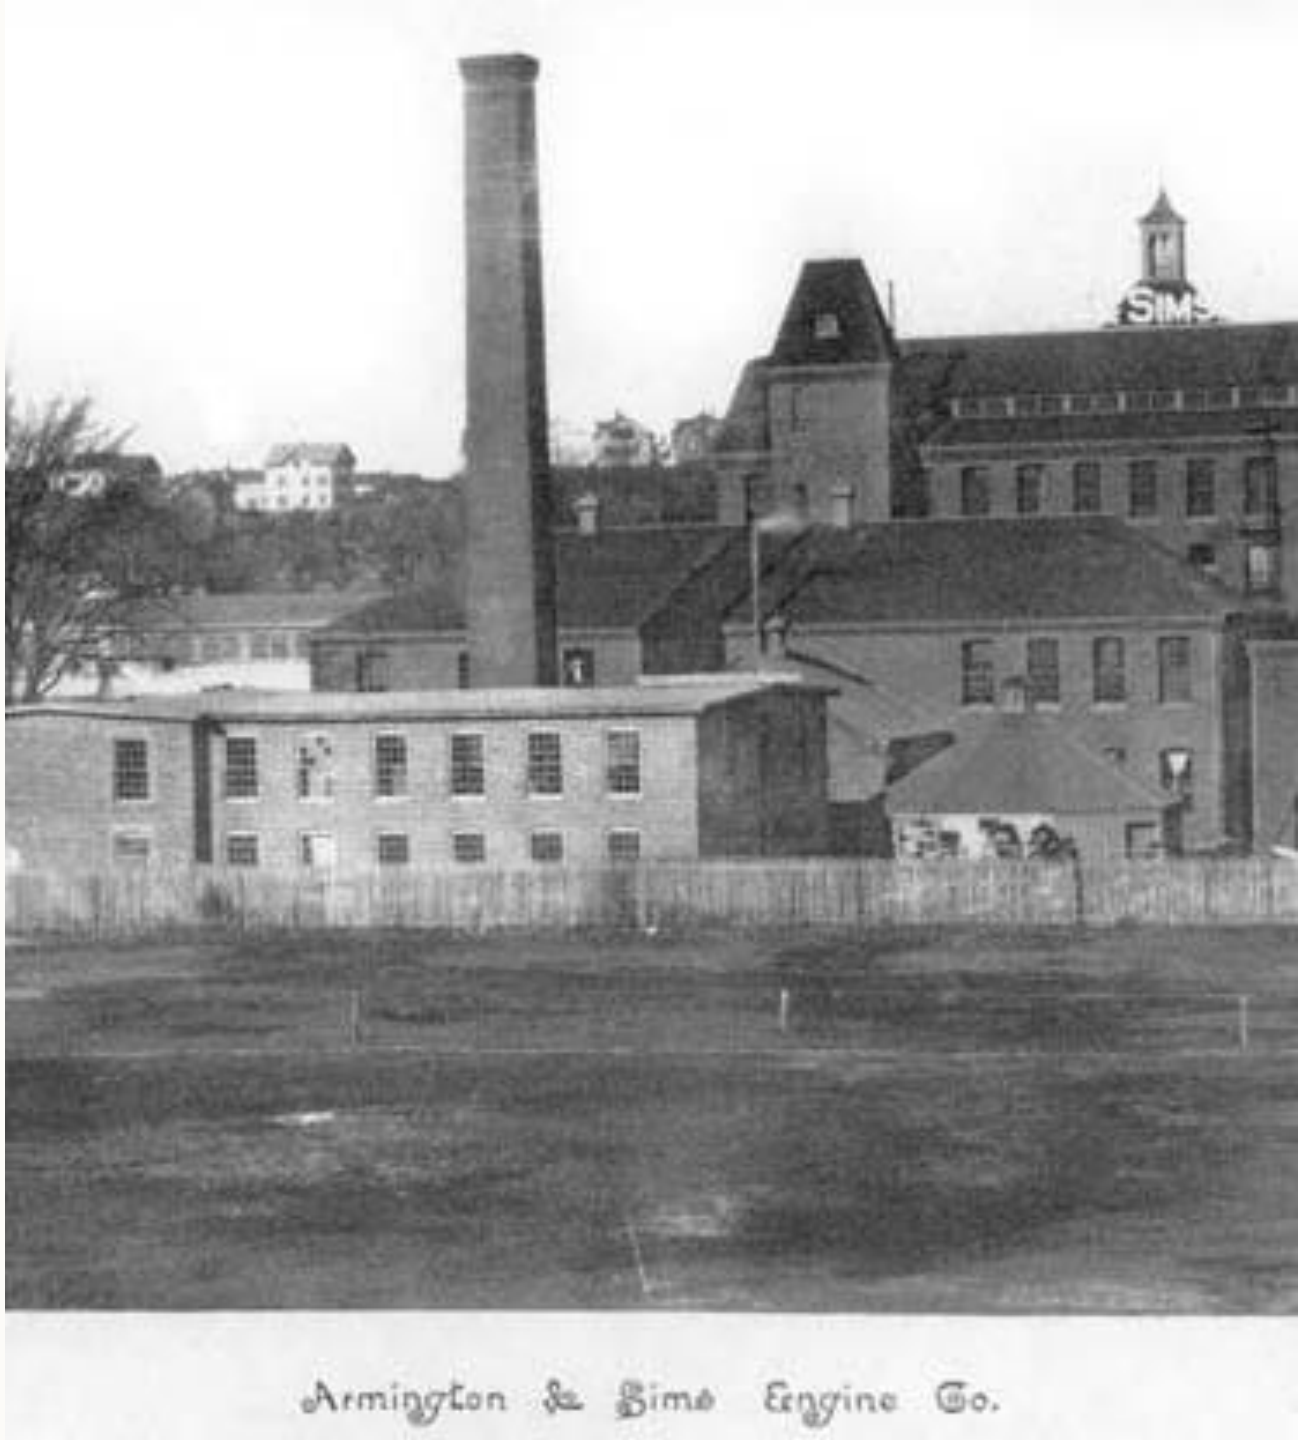
LLMs v. robotics v. other AI

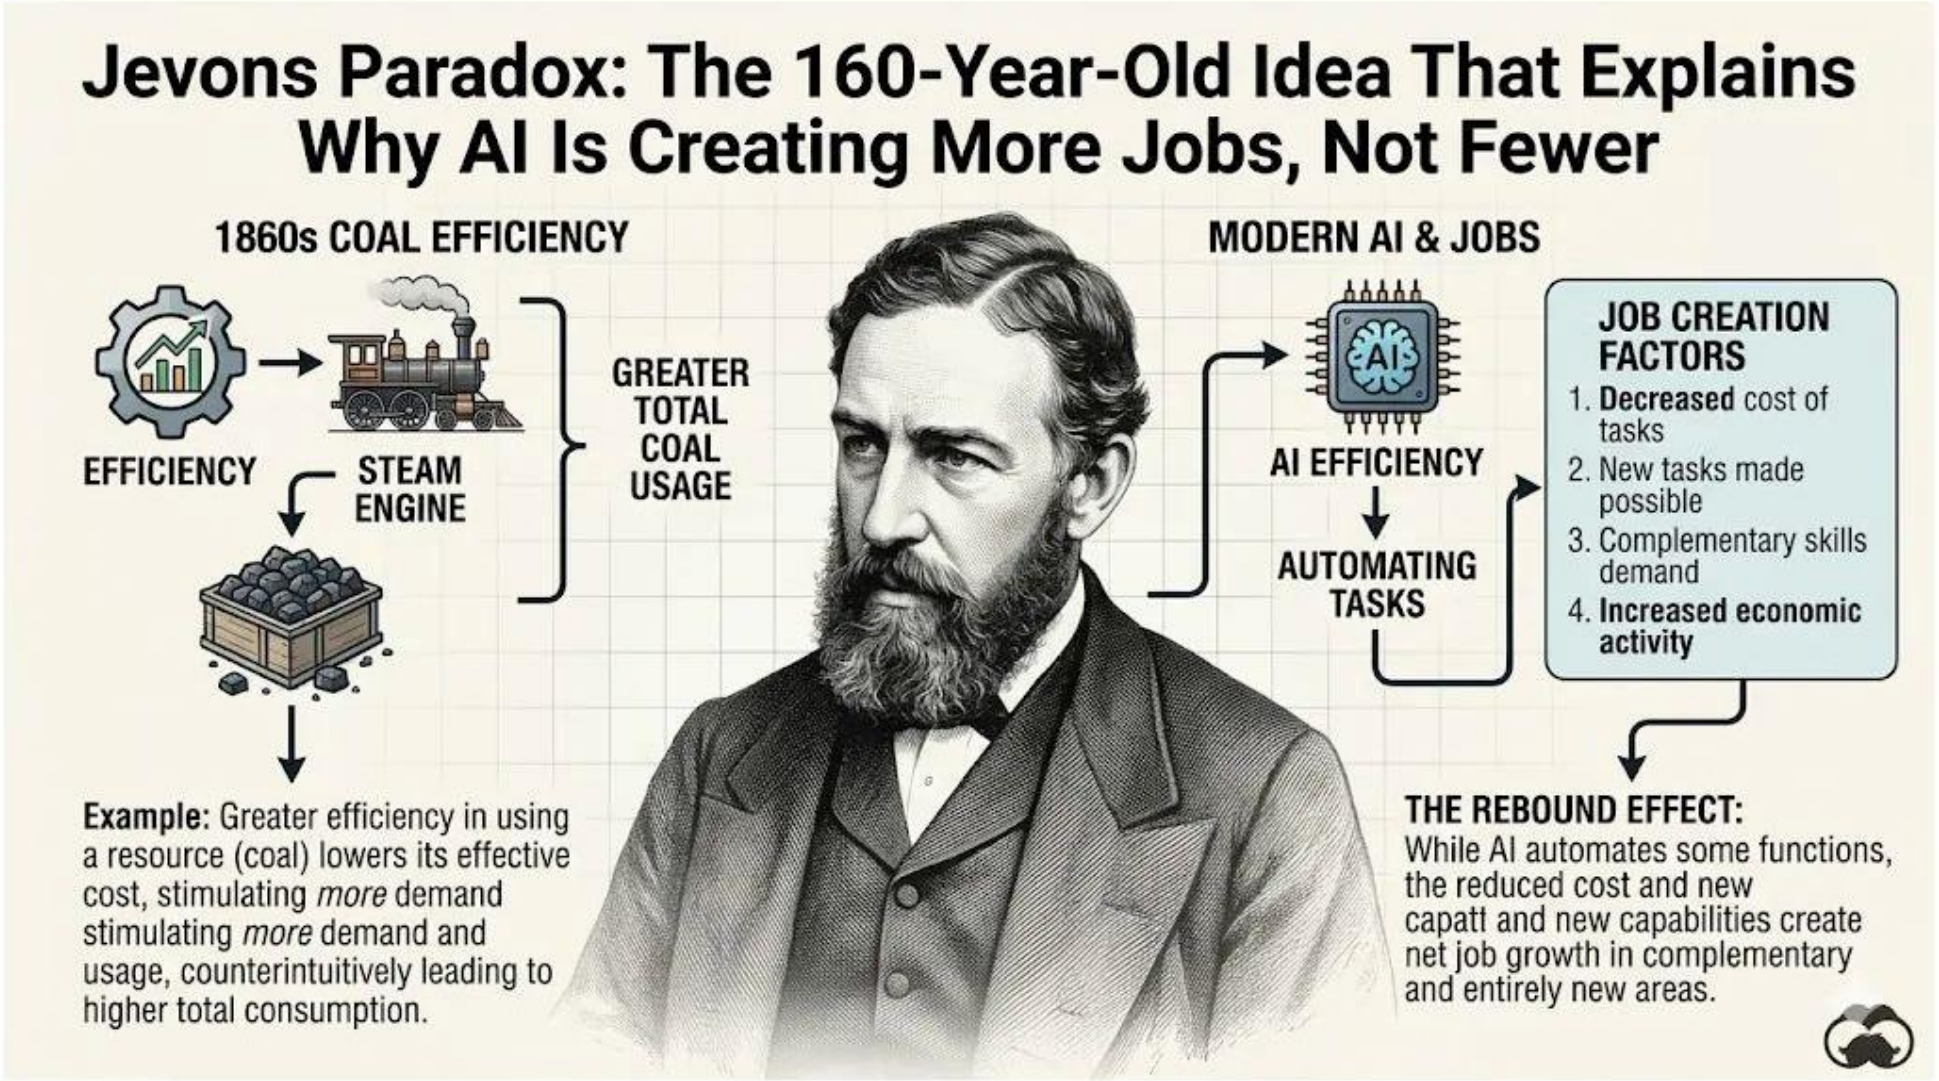
Capabilities v. acceptance

Innovation v. diffusion

Productivity v. growth v. well-being

I. Technology and Labor throughout history





But is human labor like coal? If AI eventually can do all of what human can do, more cheaply, why would demand for labor increase?

Historical analogies to AI

1. **ADOPTION**
   How quickly and widely did it spread?

2. **TASKS & JOBS**
   Which tasks were automated—and whose employment changed?

3. **WAGES & OUTPUT**
   What happened to productivity, wages, and demand?

4. **REORGANIZATION**
   What complementary investments and organizational changes were required?



Steam, looms, agriculture, electricity, telephone switching, ATMs, and office computers all changed labor—but in different ways.

Implication for AI: The same invention can destroy an occupation, shrink a sector, or mainly change tasks within existing jobs.

Framework: Paul A. David (1990/1991); Acemoglu & Restrepo; historical case studies on following slides.

HOW FAST DID TECHNOLOGY DIFFUSE — AND WHAT HAPPENED TO WORK?

Seven historical analogies | Adoption does not imply an immediate or uniform employment effect

TECHNOLOGY	INTRODUCED	DIFFUSION PACE	EMPLOYMENT PACE / EFFECT
 **Steam power**	1712 engine; Watt improvements 1769	• U.S. manufacturing workers in steam-powered plants: 17% (1850) → >50% (1880).	• No documented overall manufacturing-job collapse; labor productivity and factory scale rose over decades.
 **Power loom**	1785 patent; commercial use c. 1800	• British cotton cloth woven by power: 10% (1820) → 97% (1850).	• Handloom weavers: ~240,000 (1820) → 43,000 (1850), an ~82% decline in 30 years.
 **Farm machinery**	Reaper 1830s; tractors c. 1910	• U.S. tractors: ~246,000 (1920) → 1.57 million (1940); combines and pickers followed.	• U.S. agricultural workforce share: 41% (1900) → 16% (1945) → 1.9% (2000); multiple causes.
 **Electricity**	1882 power grid; factory motors late 1800s	• Electric share of U.S. factory horsepower: 4.8% (1899) → 55% (1919) → 82.3% (1929).	• Factory tasks and layouts changed over decades; no reliable single estimate of jobs eliminated by electrification.
 **Phone switching**	Automatic switch patented 1889	• AT&T mechanized more than half of its U.S. network in 1920–40.	• Operator jobs largely disappeared; incumbents suffered, but later cohorts moved into other jobs.
 **ATMs**	First cash dispenser 1967	• Introduced in U.S. during 1970s; large-scale ATM rollout from mid-1990s.	• Tellers/urban branch: ~20 (1988) → 13 (2004); branches +43%, buffering job losses.
 **Office computers**	VisiCalc 1979; IBM PC 1981	• U.S. employees using a computer at work: 25% (1984) → 55.5% (2003).	• Routine task demand fell within occupations over the 1980s–90s; no uniform occupation-wide employment collapse.

TAKEAWAY The labor effect depends on task substitution, demand, complementary investments, and the time allowed for adjustment.

Sources: Atack et al. (2008); Allen (2017); USDA ERS; U.S. Census of Manufactures; Feigenbaum & Gross (2024); Bessen (2015); BLS; Autor et al. (2003).

What these precedents suggest about AI and labor

Technological adoption did affect the labor force—but through several distinct mechanisms, not one uniform pattern.

Four recurring patterns

1. **Occupational collapse**
   Power looms; telephone operators (for incumbents)

2. **Sectoral shrinkage**
   Agricultural mechanization

3. **Reorganization more than headcount change**
   Steam power; electrification; office computers

4. **Task automation buffered by demand/institutions**
   ATMs

Bottom line

Steam	Adoption reorganized production; employment effects were indirect.
Looms	Adoption devastated a craft occupation despite output growth.
Agriculture	Adoption steadily reduced labor needed in an entire sector.
Electricity	Adoption paid off only after complementary organizational change.
Telephones	Adoption hurt incumbents most; later cohorts adjusted elsewhere.
ATMs	Adoption cut labor per branch, but demand and regulation shaped total jobs.
Computers	Adoption changed tasks within occupations before jobs disappeared.

Implication for AI: For AI, the central empirical questions are speed of adoption, which tasks are automated, whether demand expands, and how firms redesign work around the technology.

How AI affects labor demand: three channels



Displacement

AI takes over tasks people did



Productivity

AI makes remaining tasks faster or cheaper; cheaper output can raise demand



Reinstatement

New tasks in which people have the advantage

Acemoglu & Restrepo task framework; Acemoglu (2025), Economic Policy.

Five questions for any occupation



Task exposure

What share of the work can AI perform?



Automation or augmentation

Does AI do the task, or speed up a person who stays responsible?



Demand elasticity

When output gets cheaper, does demand expand?



Institutional brakes

Licensing, liability, regulation, unions, public acceptance



Organizational shape

Is the occupation a pyramid fed by junior hires?

Differentiating among tasks

Easy to learn

- Simple mapping from action to outcome
- Reliable, observable outcome measure
- Early AI gains skew here

Hard to learn

- Context-dependent decisions
- No objective outcome measure
- Complex judgment and subtlety

Complementary

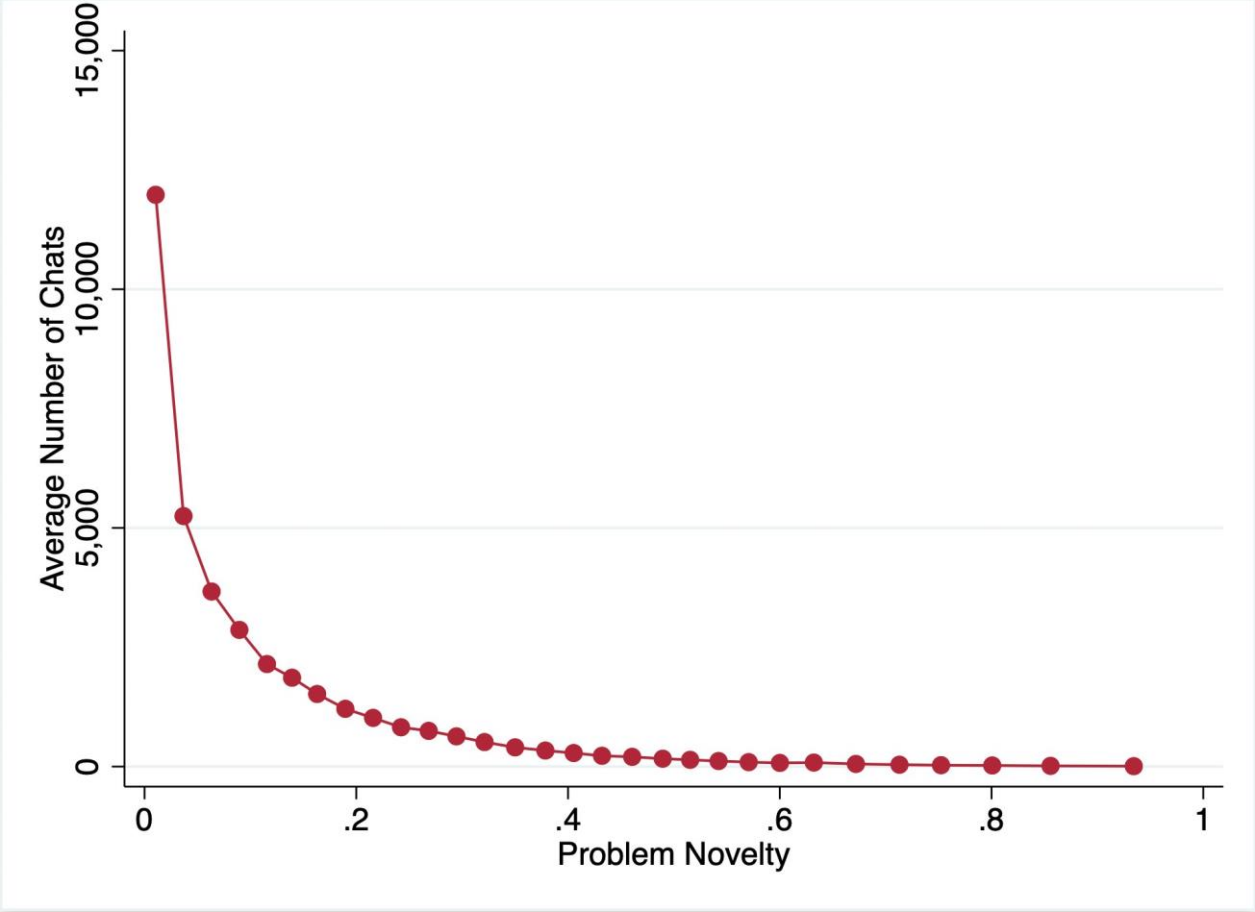
- AI augments rather than replaces
- Helps with related subtasks

New tasks

- Good: add value and opportunity
- Bad: may undermine welfare (e.g., manipulative advertising)

Acemoglu (2025), The Simple Macroeconomics of AI, Economic Policy.

The long tail of tasks (evidence from call centers)



Machine learning excels where training data are abundant; much work is a long tail of rare problems with little data.

Credit: Erik Brynjolfsson.

Same exposure, different outcomes: demand elasticity

Occupation	AI exposure	When output gets cheaper…	Employment / projected demand
Radiologists	High (imaging tasks)	Imaging demand rises	Employment +~3% annually, 2017–24
Software developers	High, especially programming	More software gets built	BLS +10.2%, 2025–35
Customer service reps	High for routine contacts	Volume need not grow	BLS about −5%, 2025–35
Translators	High for commodity text	Demand may shift to review	Some roles shrinking; uneven effects

Exposure measures the supply side. Demand response, regulation and acceptance decide whether employment rises or falls.

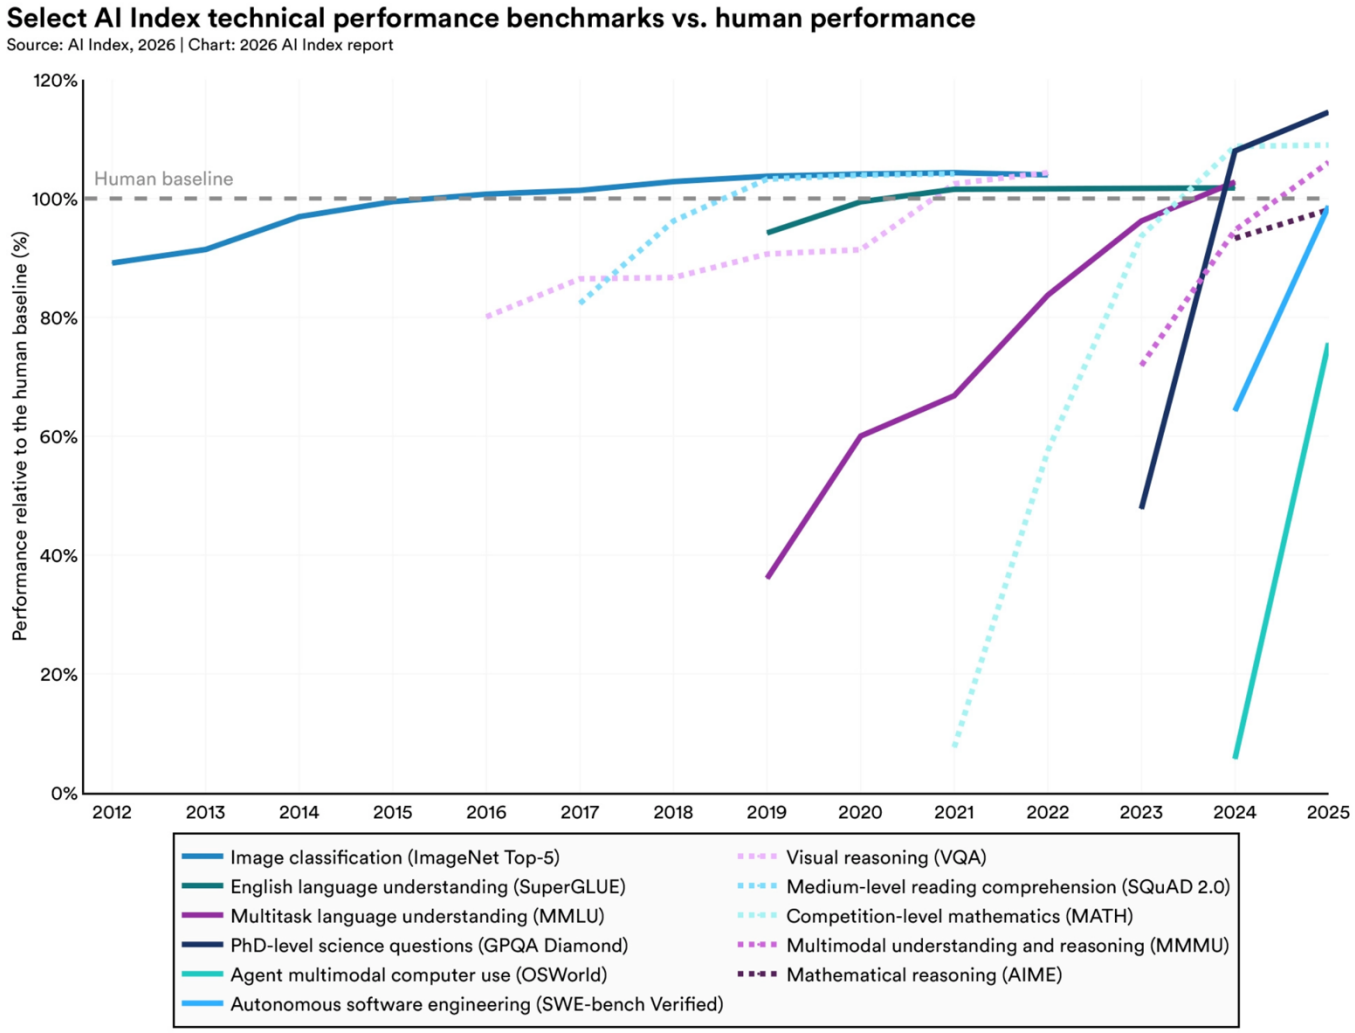
Sources: McKinsey (2026); BLS Employment Projections 2025–35; occupation evidence varies.

II. Adoption and use

How similar is AI to earlier technological innovations?



AI model performance across benchmarks



The length of tasks AI agents can complete keeps doubling



- Task length (in human time) that agents complete with 50% success
- Doubling about every 4 months since 2023; about every 3 months since 2024
- Unlabeled point: GPT-5.2 (6.6 hours)
- Claude Mythos Preview: ~17 hours, but METR says estimates above 16 hours are unreliable on its current suite

Source: METR Time Horizon 1.1 public run data (GitHub, Mar. 2026), 50% horizons computed with METR's method; Mythos Preview per METR (May 2026).

/ FIGURE

Six-country survey: Weekly use by AI product, 2024–2026

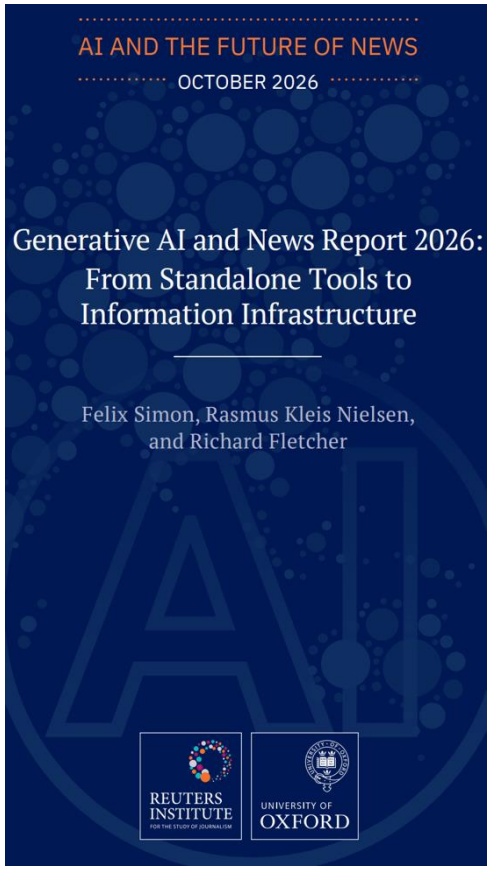
Proportion who uses each generative AI system weekly

On average across six countries, ChatGPT remains the most widely used system – but use of Google Gemini, Microsoft Copilot, and Claude is also growing.

System	2024	2025	2026	Change, 2025–26
ChatGPT	**11%**	**22%**	**34%**	**+12**
Google Gemini	**4%**	**11%**	**24%**	**+13**
Microsoft Copilot	**4%**	**6%**	**11%**	**+5**
Meta AI	3%	9%	11%	+2
Claude	**1%**	**1%**	**7%**	**+6**
Grok	0%	2%	4%	+2
Perplexity.ai	1%	1%	3%	+2
Snapchat My AI	2%	2%	2%	±0
DeepSeek		2%	2%	±0
Character.ai			1%	n/a
OpenClaw			1%	n/a
Mistral	1%	1%	1%	±0
Midjourney	1%	1%	1%	±0
Replika			1%	n/a

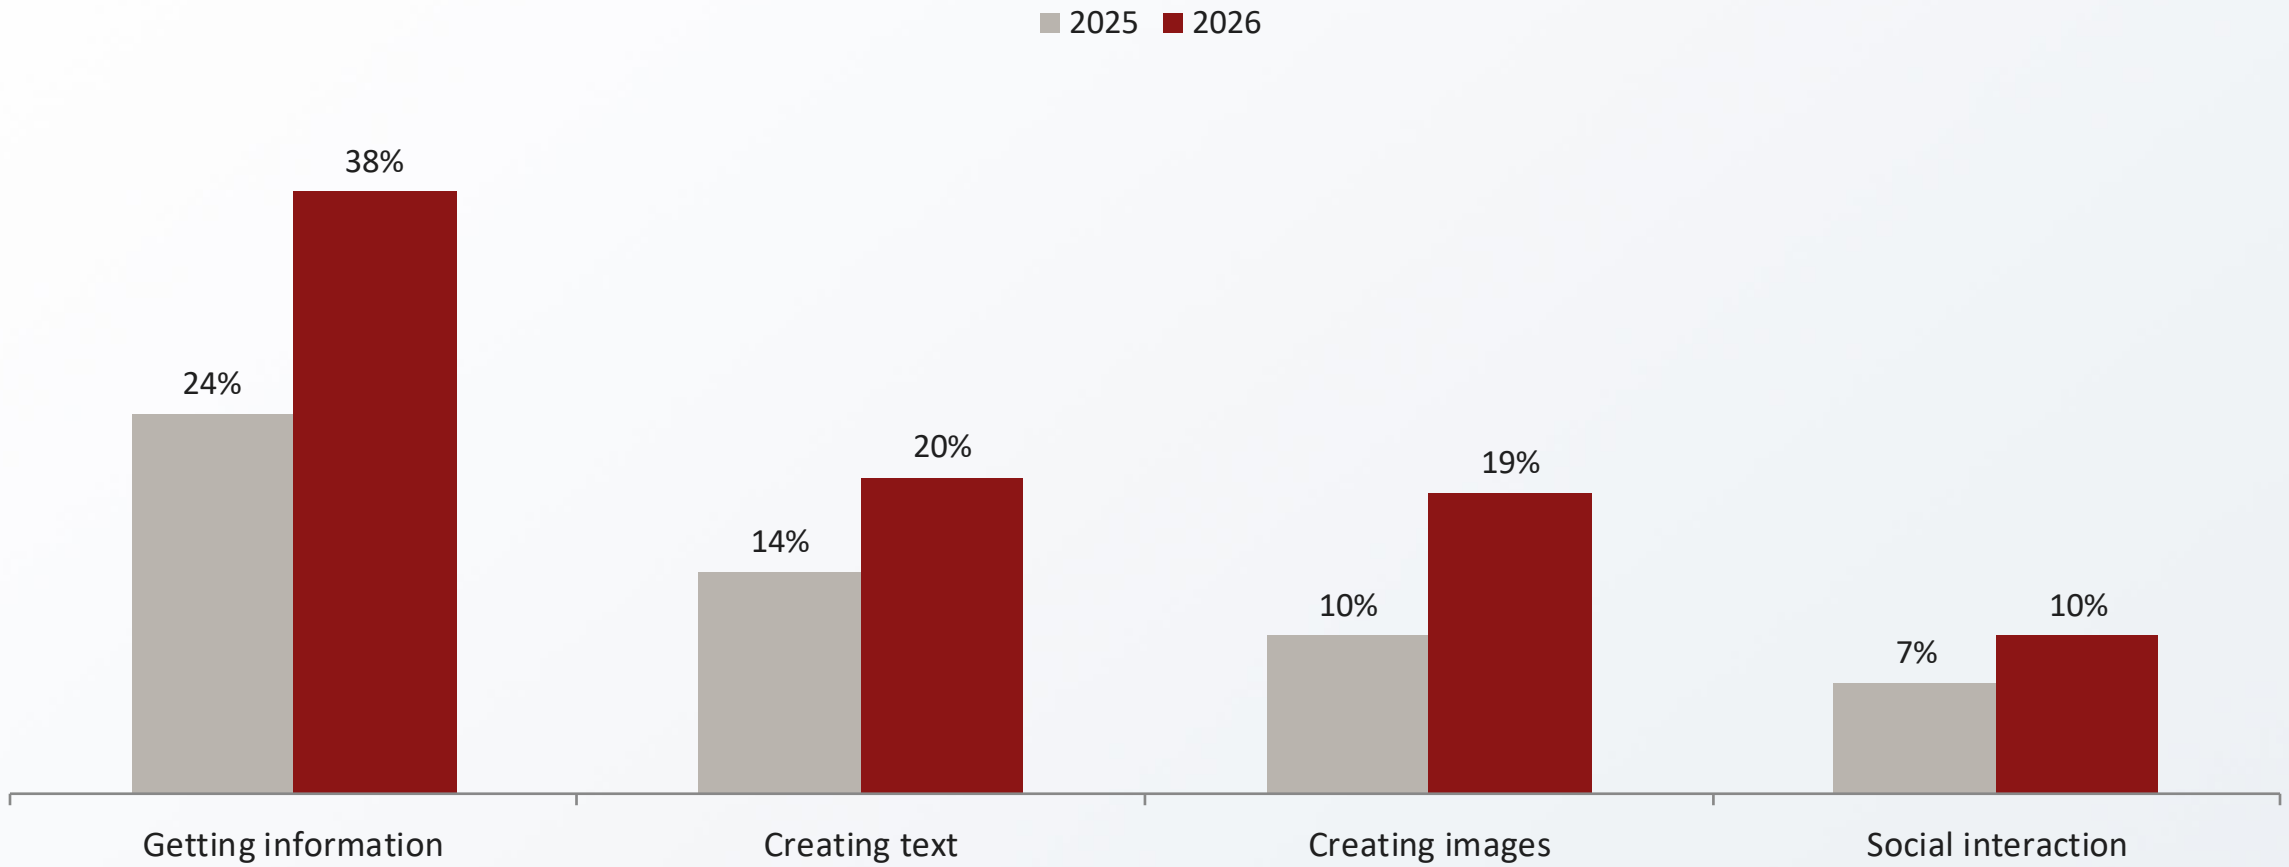
AI_branduse. How often, if at all, do you typically use each of the following generative AI chatbots or tools for any purpose? *Base: Total sample across Argentina, Denmark, France, Japan, the UK, the USA ≈ 12,000. Note: Empty cells indicate that a brand was not asked about that year.*

- Weekly generative AI use: 51% in 2026, versus 34% in 2025.
- Daily use: 28% in 2026, 16% in 2025, and 7% in 2024.



Source: Reuters Institute, Generative AI and News Report 2026 (YouGov, Aug. 14–Sept. 28, 2026; six-country average: Argentina, Denmark, France, Japan, U.K., U.S.).

Tasks people use generative AI for weekly



Source: Reuters Institute, Generative AI and News Report 2026 (six-country average).

/ FIGURE

OpenAI: distribution of work-oriented ChatGPT tasks

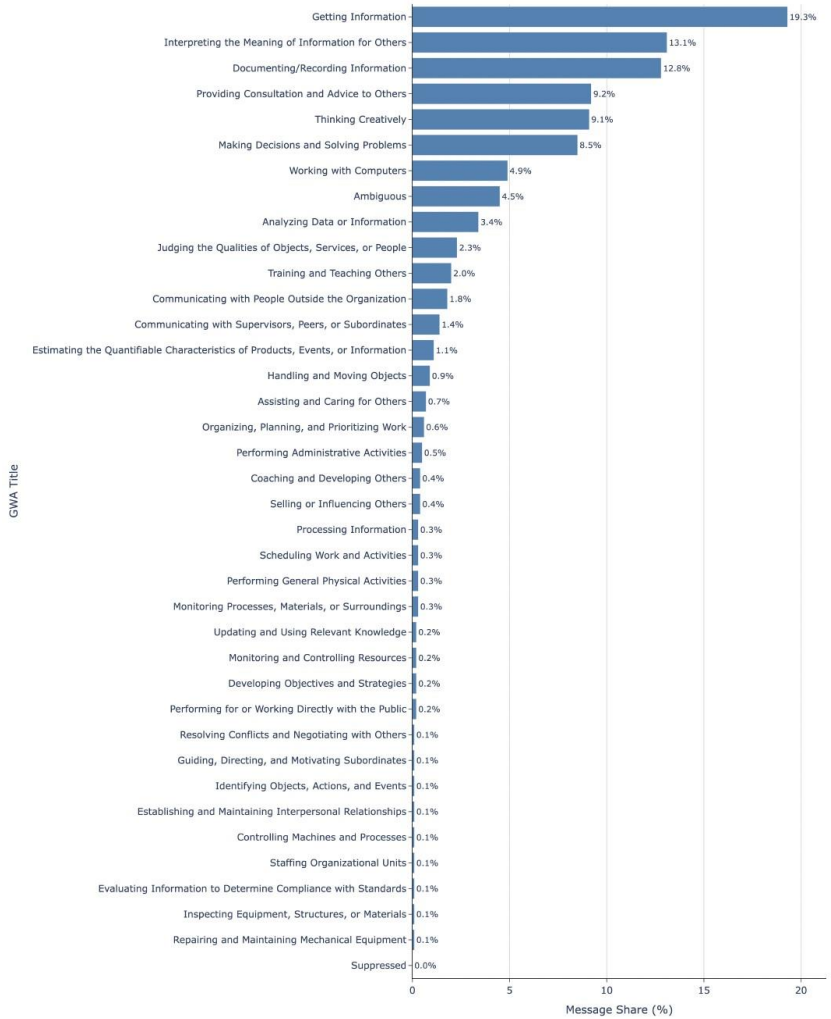
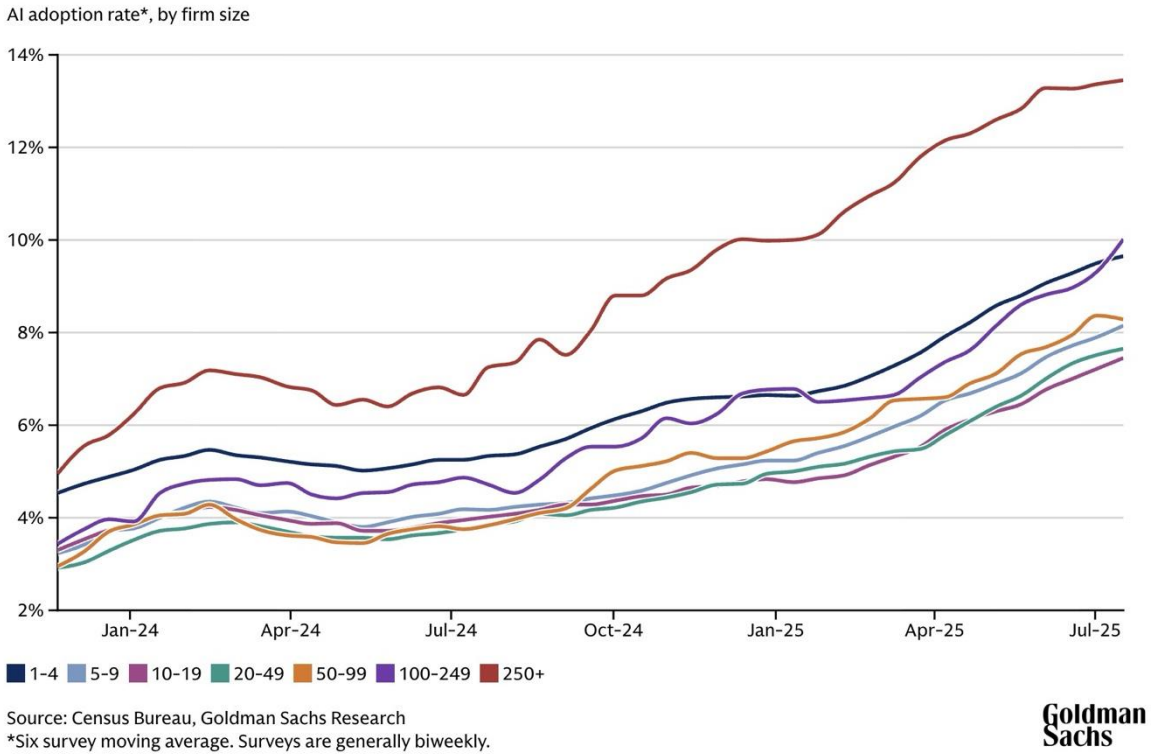


Figure 14: GWA Shares of 1.1M ChatGPT Messages. Messages are classified as pertaining to one of 332 O*NET IWAs, or *Ambiguous* using the prompt provided in the Appendix. IWAs were then aggregated to GWAs using the O*NET Work Activities taxonomy. Message sample from May 15, 2024 through June 26, 2025. We do not show the shares for the following GWAs as there were fewer than 100 users sending messages for each category and group them into *Suppressed*.

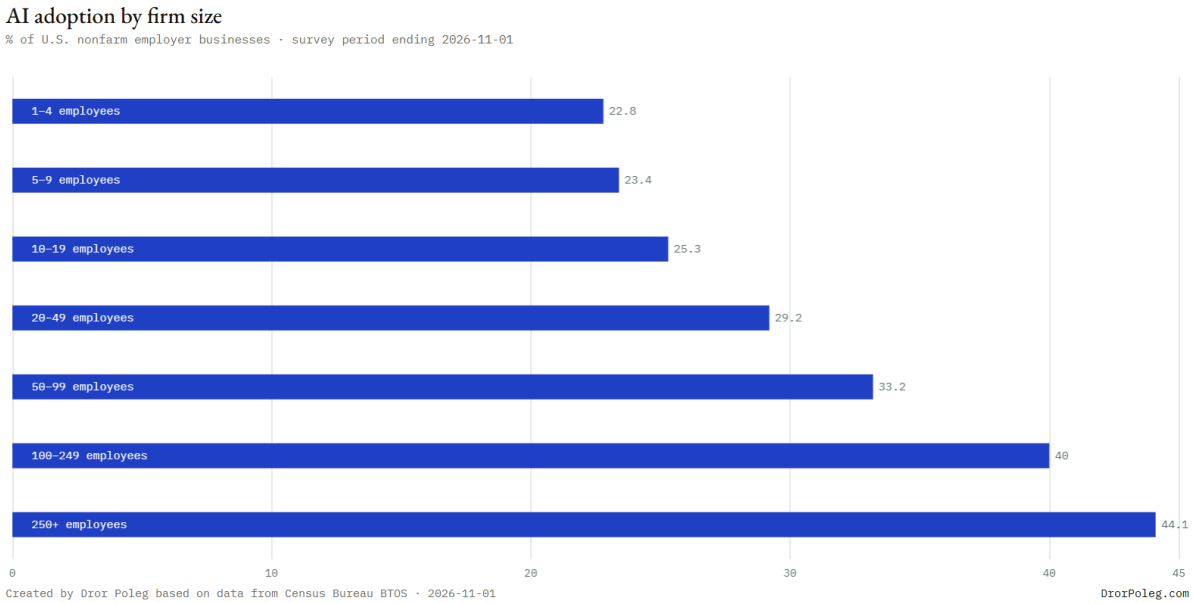
/ FIGURE

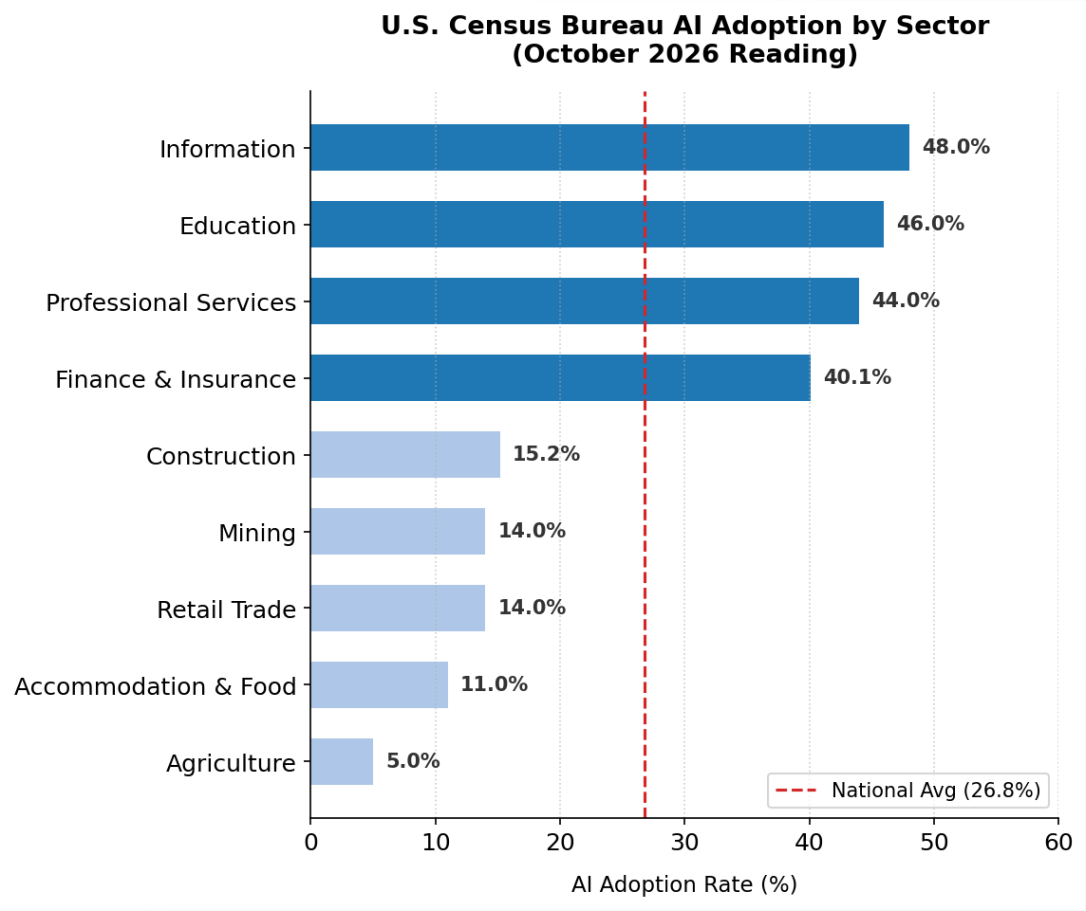
Business AI adoption by firm size, 2023–2026



Source: Census Bureau, Goldman Sachs Research
*Six survey moving average. Surveys are generally biweekly.

Goldman Sachs





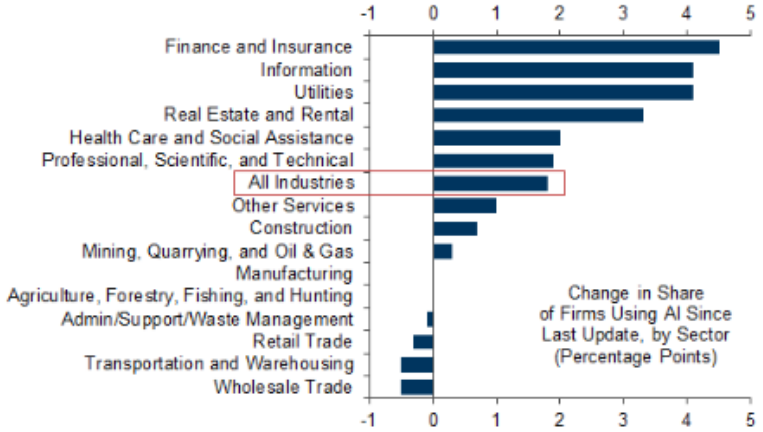
"In the last two weeks, did this business use Artificial Intelligence (AI) in any of its business functions? (Examples of AI: machine learning, natural language processing, virtual agents, voice recognition, etc.)"

Exhibit 9: Adoption Is Approaching 50% in the Information Sector



Source: Census Bureau, Goldman Sachs Global Investment Research

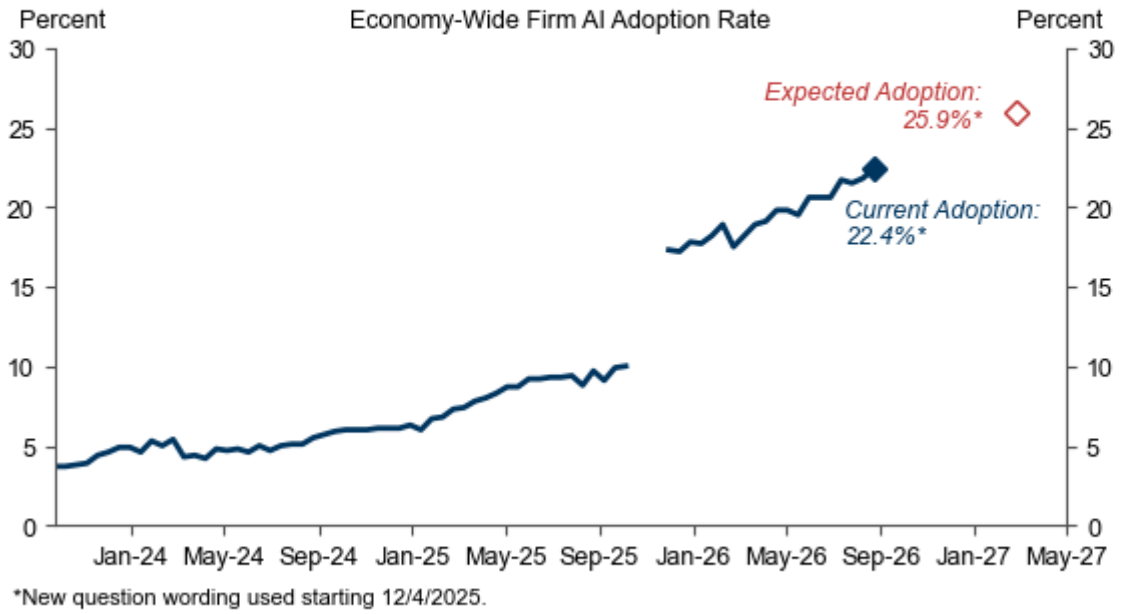
Exhibit 10: Finance and Insurance Firms Had the Largest Gains in Adoption Since Our Last Update



Source: Census Bureau, Goldman Sachs Global Investment Research

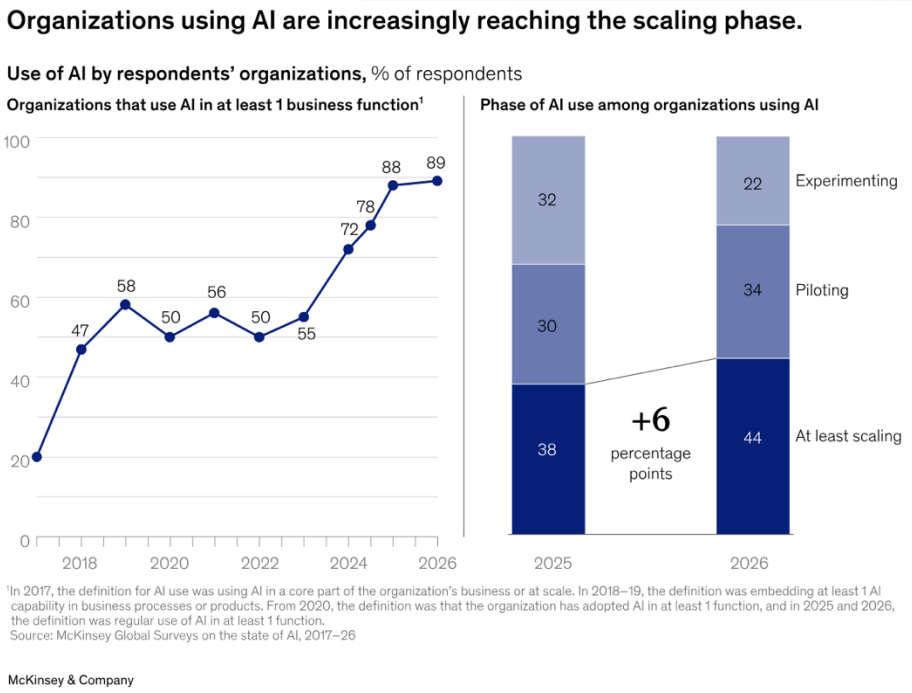
During the next six months, do you think this business will be using Artificial Intelligence (AI) in any of its business functions?"

Exhibit 8: 22.4% of Firms Are Using Artificial Intelligence in Their Regular Business Functions (+0.9pp From July); 25.9% Expect to Use AI In the Next 6 Months



*New question wording used starting 12/4/2025.

Source: Census Bureau, Goldman Sachs Global Investment Research



The original question, used since fall 2023, asked: "In the last two weeks, did this business use Artificial Intelligence (AI) in producing goods or services?"

"In the last two weeks, did this business use Artificial Intelligence (AI) in any of its business functions? (Examples of AI: machine learning, natural language processing, virtual agents, voice recognition, etc.)"

Variation within the U.S.

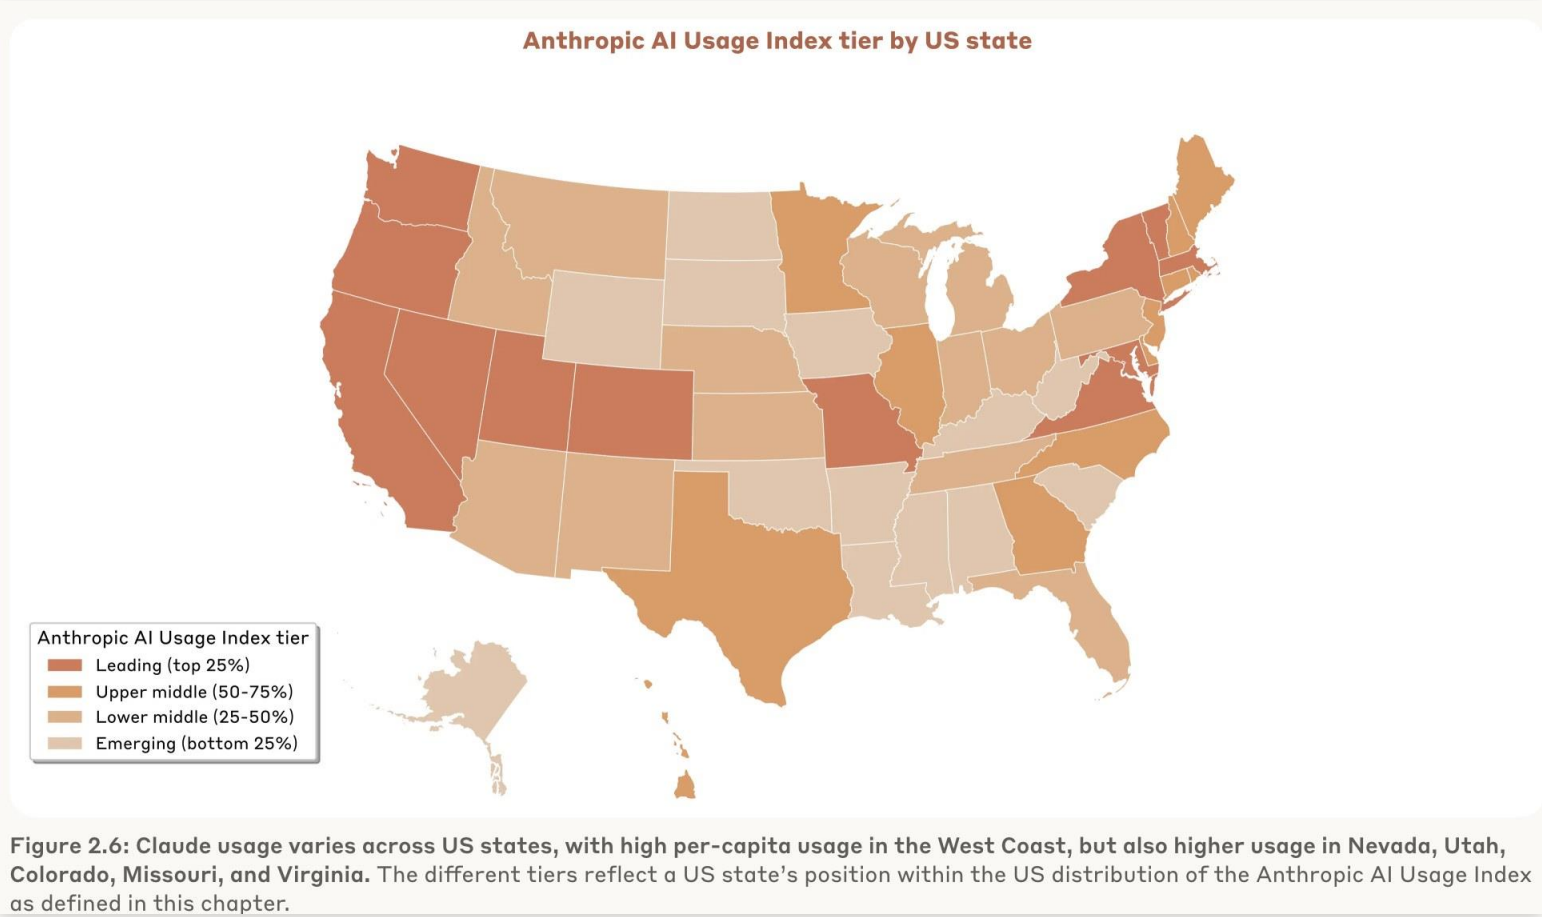


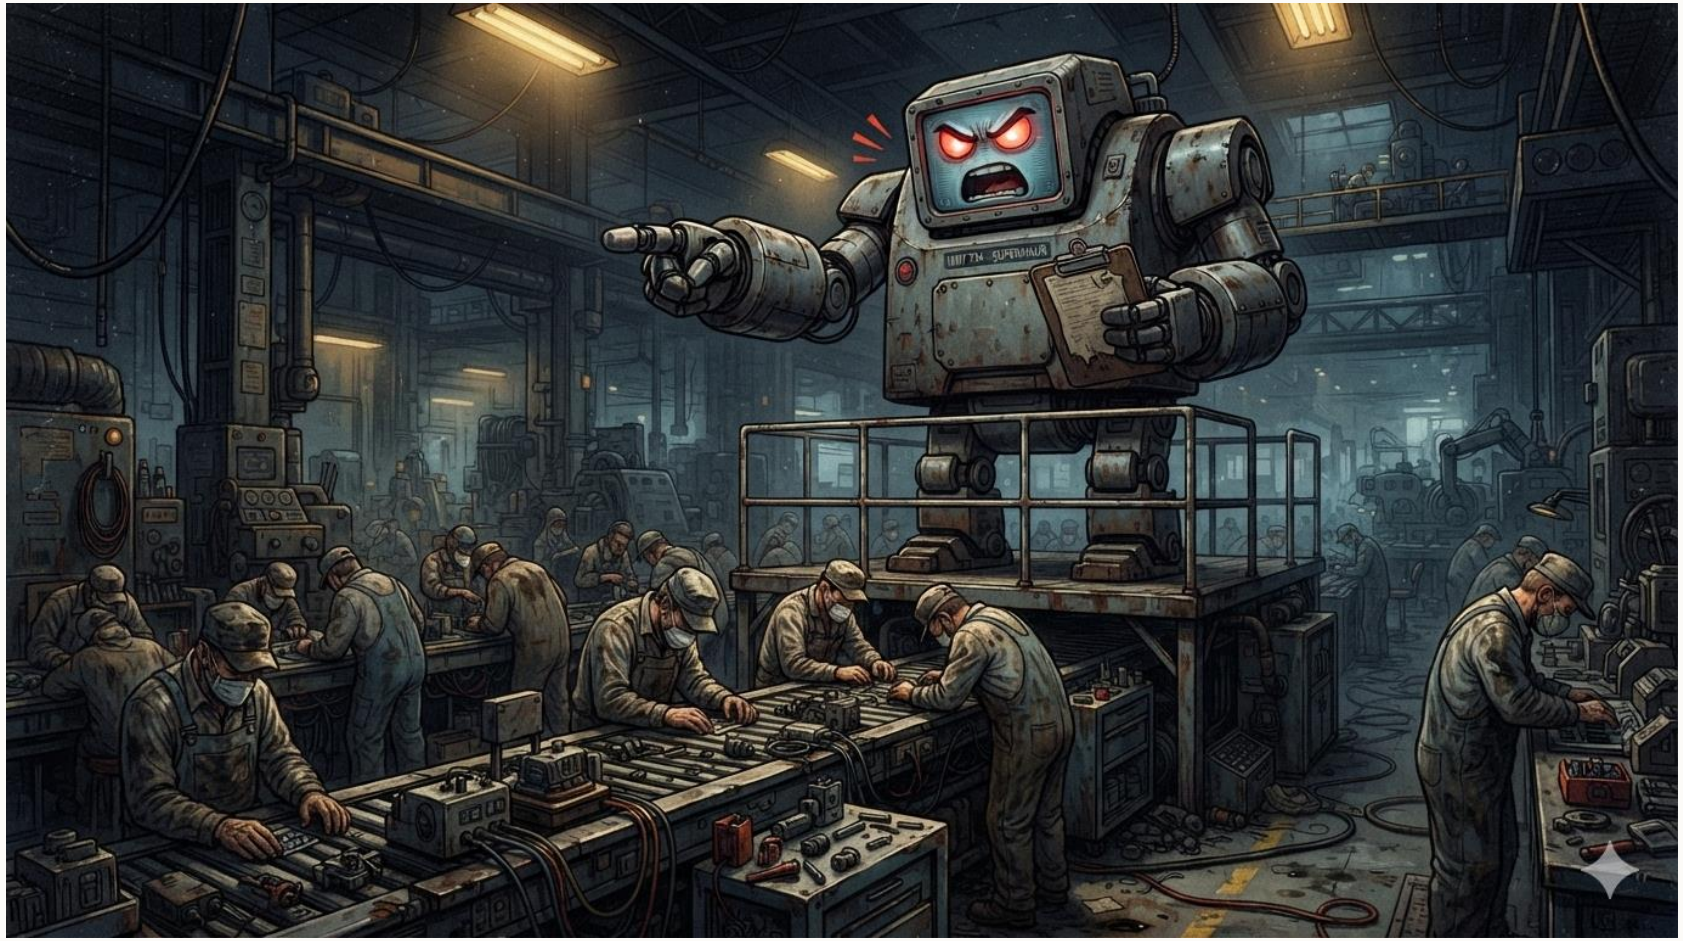
Figure 2.6: Claude usage varies across US states, with high per-capita usage in the West Coast, but also higher usage in Nevada, Utah, Colorado, Missouri, and Virginia. The different tiers reflect a US state's position within the US distribution of the Anthropic AI Usage Index as defined in this chapter.

Jan. 2026 update: use became less concentrated across states (Gini 0.37 → 0.32, Aug.–Nov. 2025); top five states: 50% of use, 38% of working-age population.

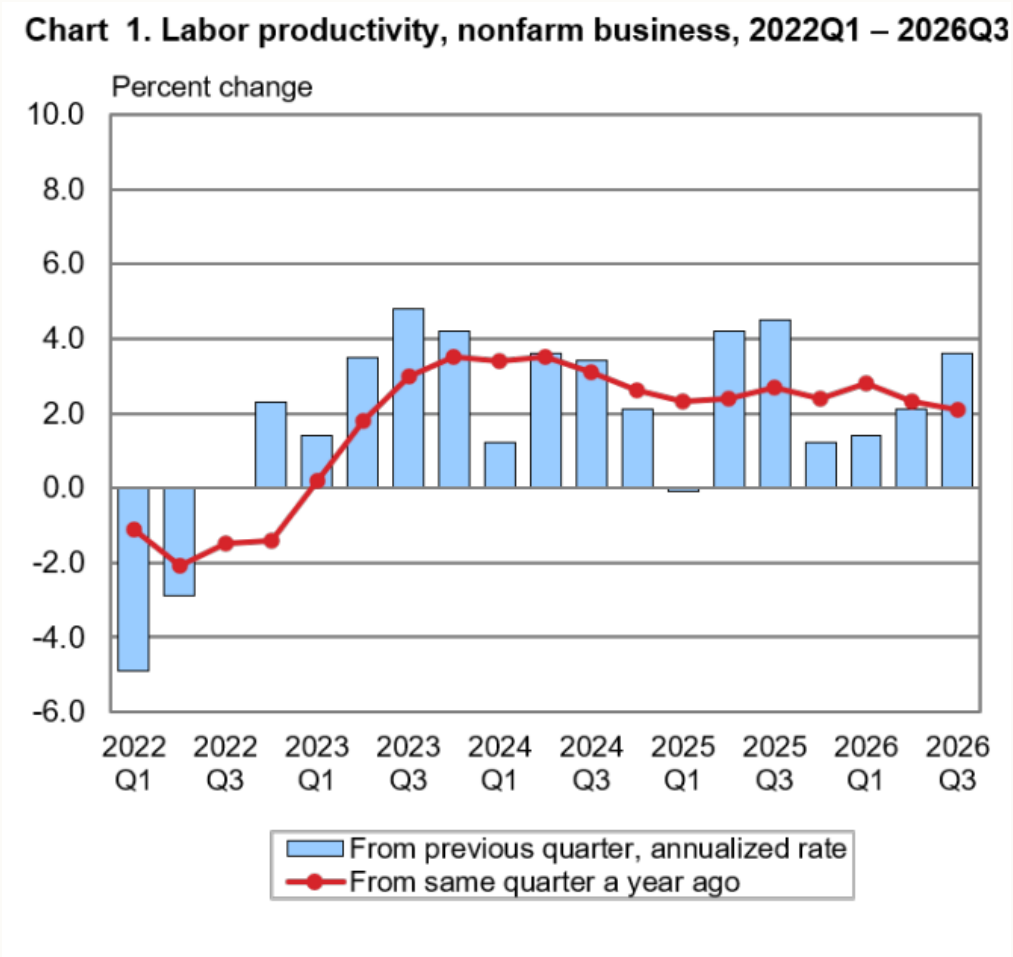
Sources: Anthropic Economic Index (Sept. 15, 2025; Jan. 2026).

III. What the data show so far

A. Productivity
B. Job cuts
C. Job-finding rate
D. Predicted employment changes



No evidence, yet, of AI inflection in productivity growth



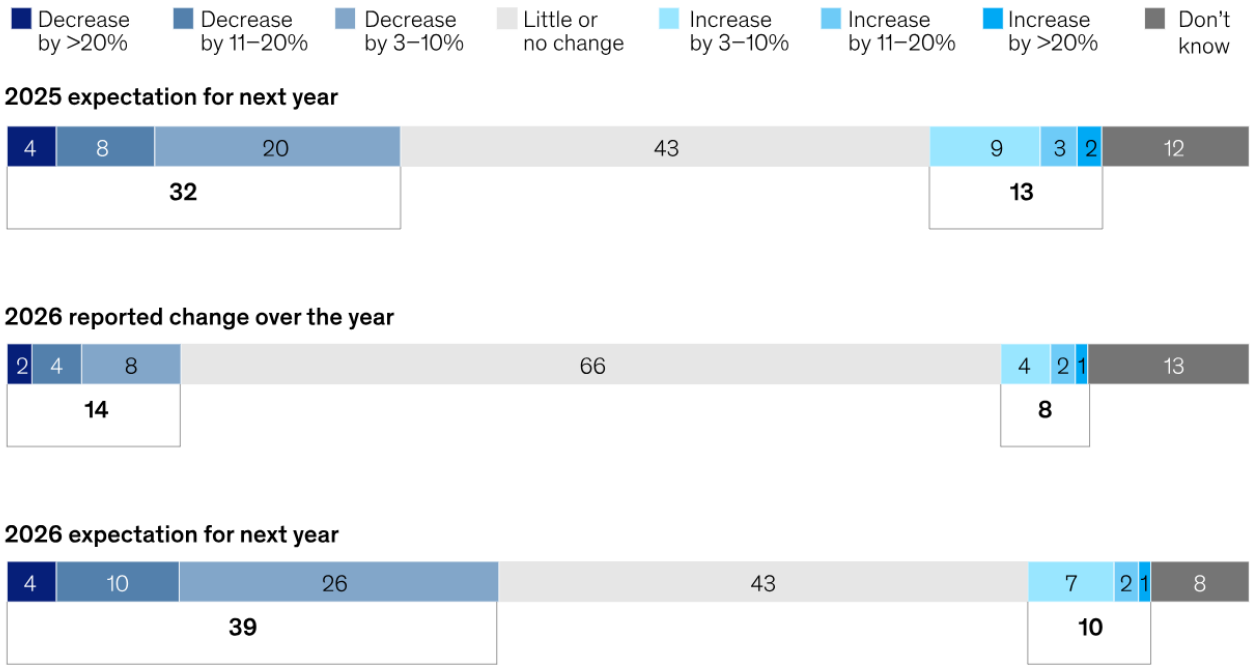
- Productivity growth since 2019 has run above the 2007–19 pace
- Q3 2026: +3.6% at an annual rate
- Whether AI explains the pickup is not yet established

Data: BLS, Productivity and Costs, Q3 2026 preliminary (Nov. 5, 2026), https://www.bls.gov/news.release/archives/prod2_11052026.pdf?utm_source=chatgpt.com.

Predicted job cuts, not yet materialized, but continue…

Last year's respondents overestimated AI's impact on head count, and expected reductions in the next year are larger still.

Expected change in number of employees across the enterprise as a result of AI in the next year,[1]
% of respondents

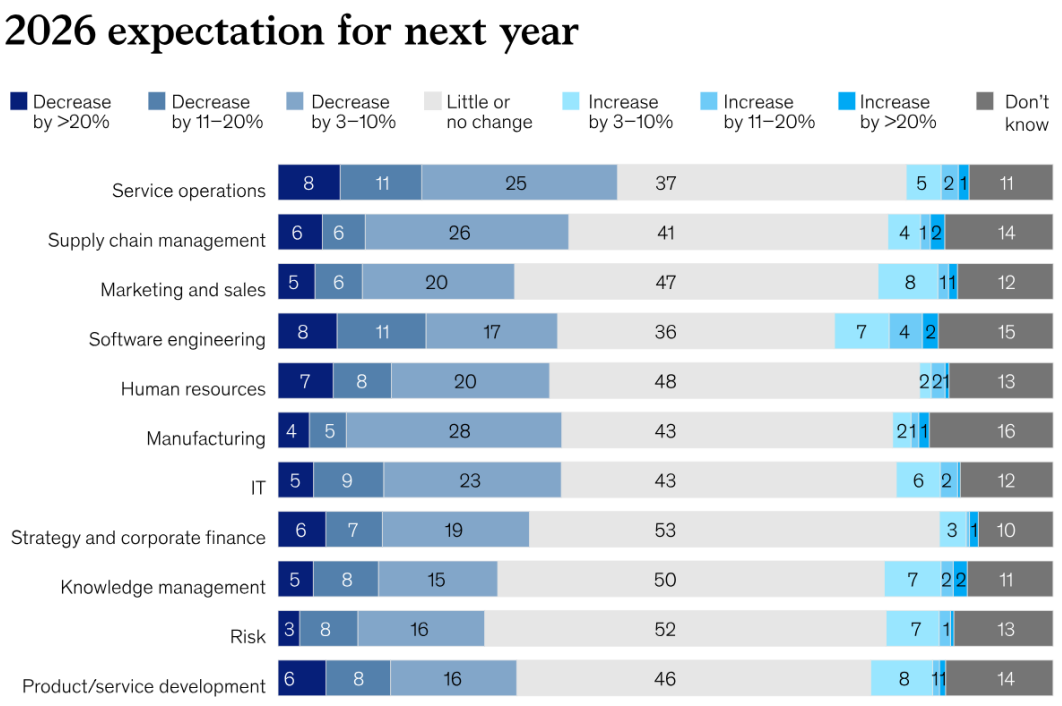


Note: Figures do not sum to 100%, because of rounding.
[1]Asked only of respondents who said their organizations regularly use AI in at least 1 business function. In 2025, n = 1,753; in 2026, n = 1,521.
Source: McKinsey Global Surveys on the state of AI, June 25–July 29, 2025

McKinsey & Company

While expectations for head count reductions have increased, a plurality continue to expect no change in many functions.

Reported percent change in employees over the past year due to AI in the given business functions,[1]
% of respondents in 2026 survey reporting AI use in at least 1 function

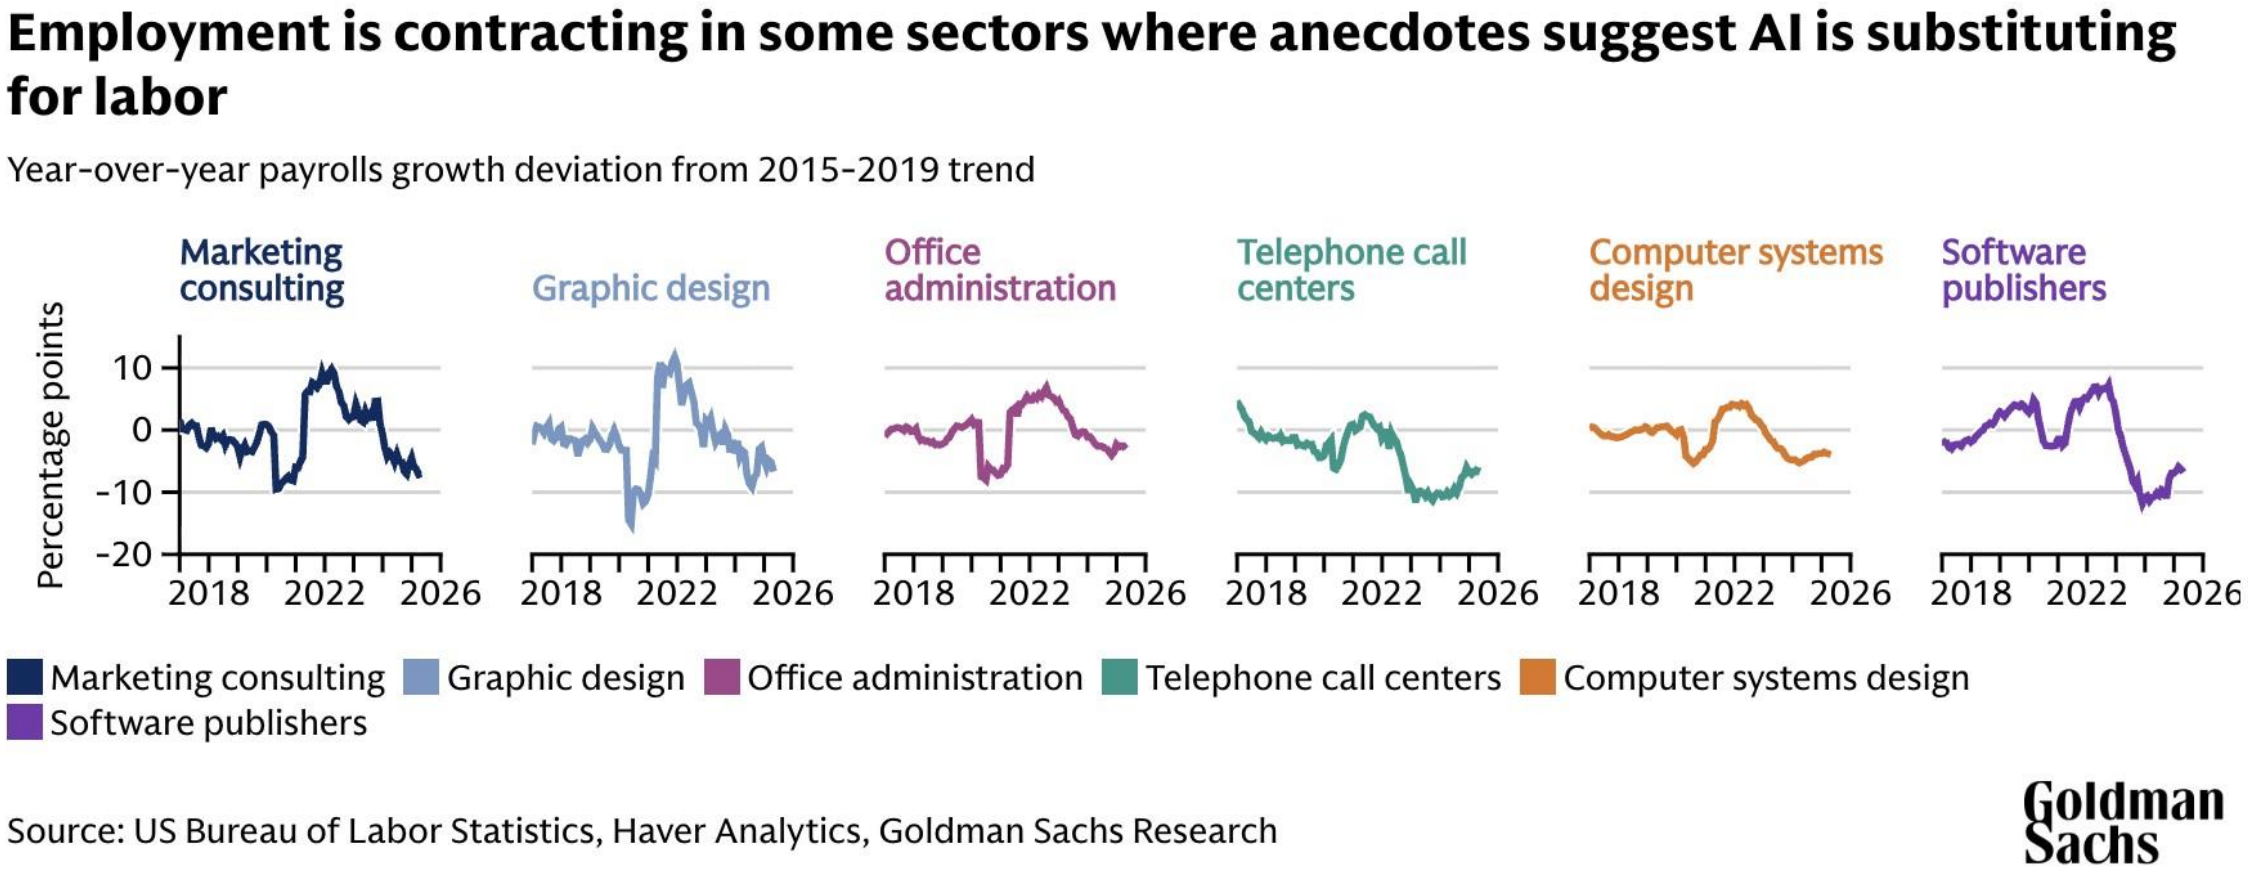


Note: Figures do not sum to 100%, because of rounding.
[1]Asked only of respondents who said their organization regularly uses AI in at least 1 business function; n = 1,753.
Source: McKinsey Global Surveys on the state of AI, May 4–June 8, 2026

McKinsey & Company

/ FIGURE

Goldman Sachs: sector hiring indicators



Canaries in the coal mine

Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence

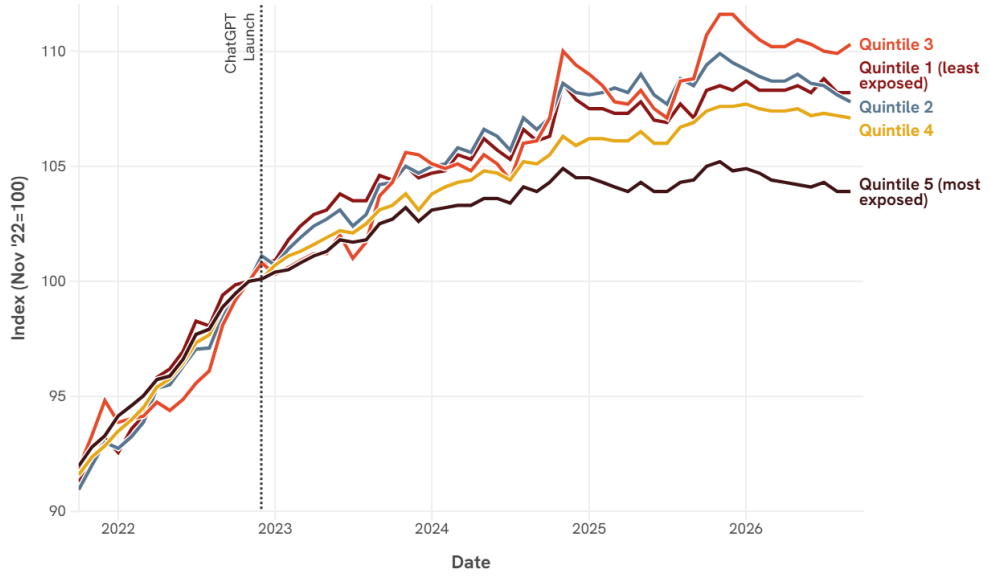
Erik Brynjolfsson* Bharat Chandar† Ruyu Chen‡

August 2026

Employment Index in Canaries Sample by AI Exposure

We group workers by their AI exposure score, comparing employment trends across these groups. We see modest differences between the five exposure groups, although employment growth is lowest for the most exposed occupations.

Data

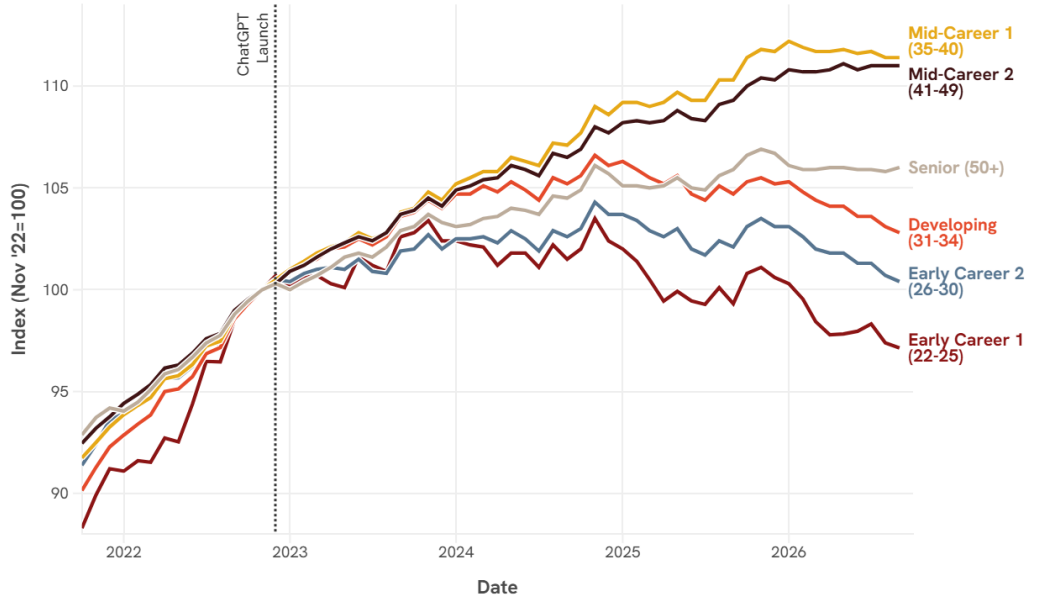


Source: Stanford Digital Economy Lab, ADP Research, Eloundou et al. (2024) ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

Employment Index in Canaries Sample by Age Group

We contextualize the AI exposure results with aggregate employment trends by age group, pooling all occupations. Consistent with the pattern in the previous chart, early-career workers (22-25) and the next-youngest group (26-30) see modest gains or slight declines. Using data from our July 2026 release, workers ages 22-25 accounted for 7.4% of employment in November 2022.

Data



Source: Stanford Digital Economy Lab, ADP Research ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

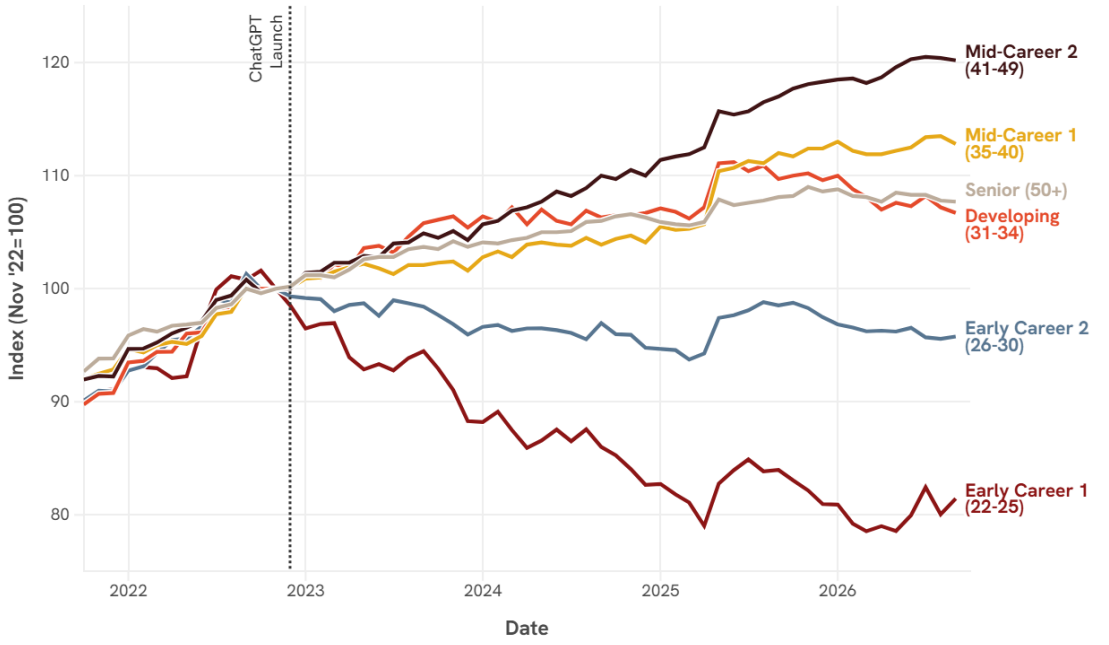
https://digitaleconomy.stanford.edu/project/indicators/canaries-dashboard/

Highly exposed occupations: Software developers, customer service

Software Developer Employment Index in Canaries Sample by Age Group

Software developers land in the most exposed occupations group. We see large declines for early-career workers (22-25), modest declines for the next-youngest groups, and expansion for the remaining age groups.

Data

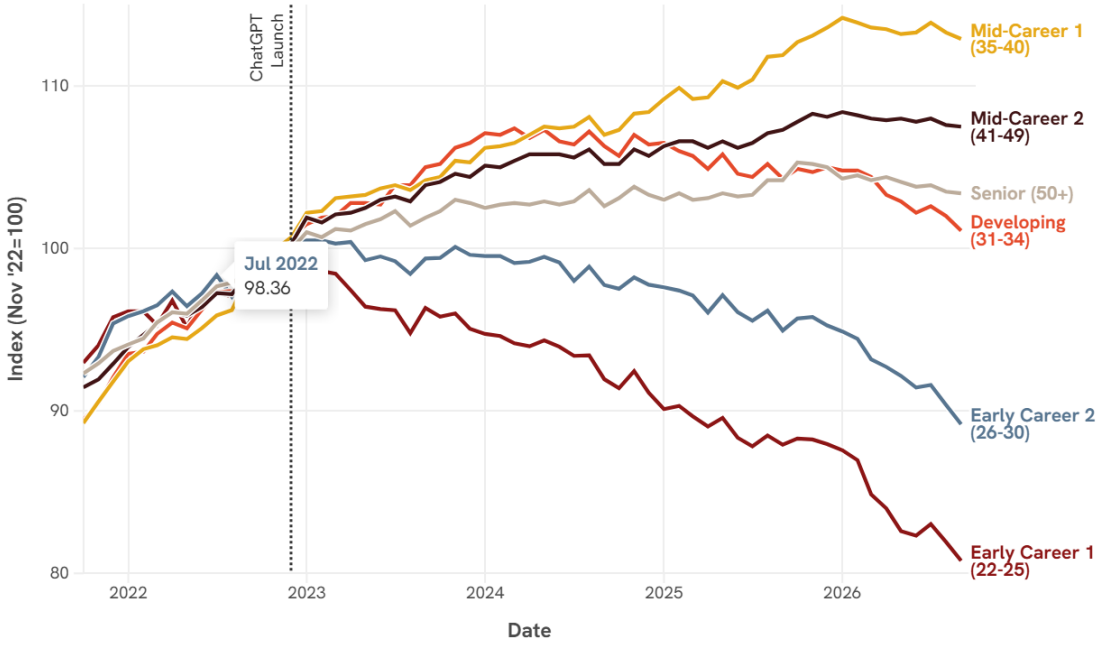


Source: Stanford Digital Economy Lab, ADP Research ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

Customer Service Representative Employment Index in Canaries Sample by Age Group

Customer service representatives are also assigned to the highest exposure group, and they indeed see similar trends to software developers: noticeable declines for the first two age groups, with expansion for the remaining age groups.

Data



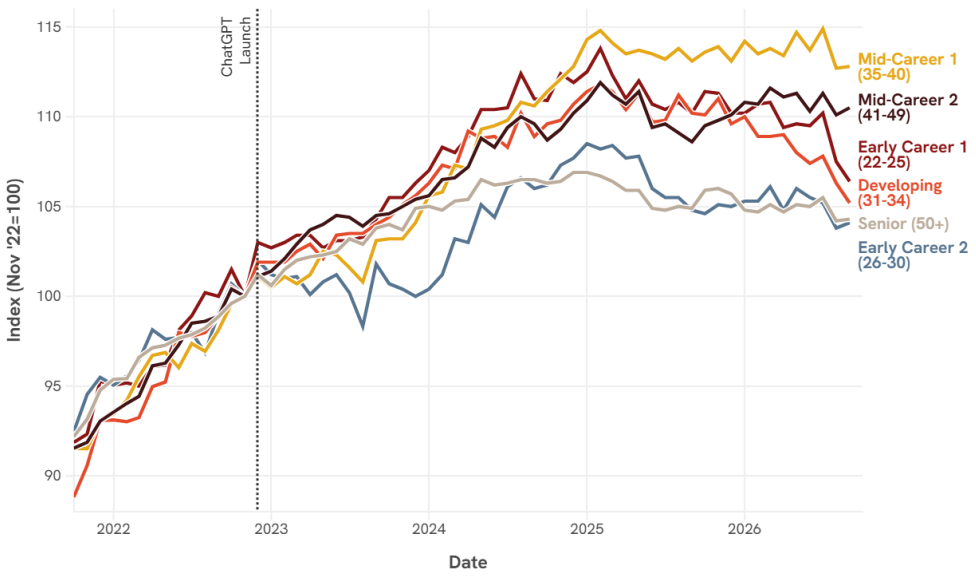
Source: Stanford Digital Economy Lab, ADP Research ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

Less exposed occupations: Stock Clerks, Home Health Aides

Stock Clerk Employment Index in Canaries Sample by Age Group

In contrast to software developers and customer service representatives, stock clerks fall in the second least exposed occupation group. There is little relationship between age and employment trends.

Data

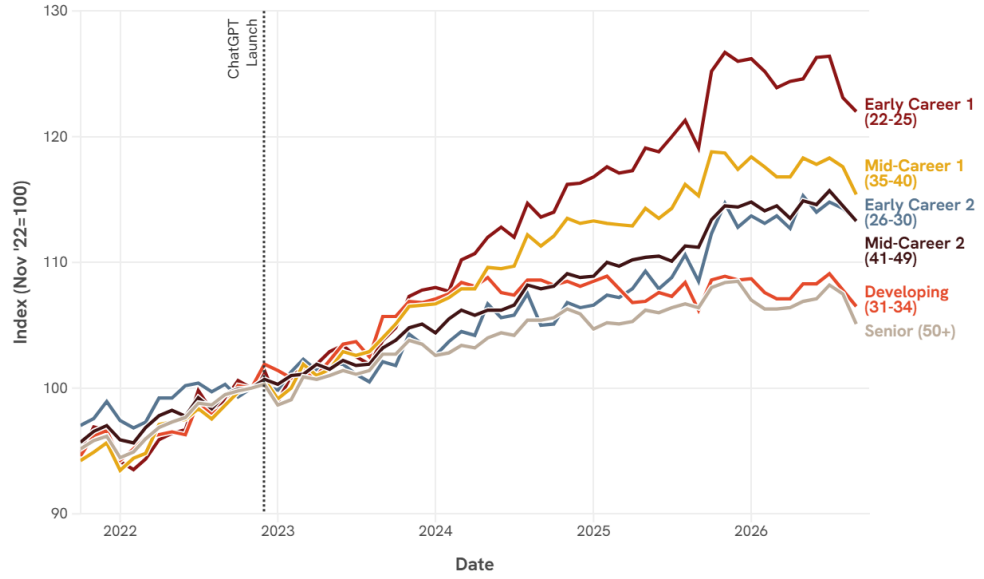


Source: Stanford Digital Economy Lab, ADP Research ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

Home Health Aide Employment Index in Canaries Sample by Age Group

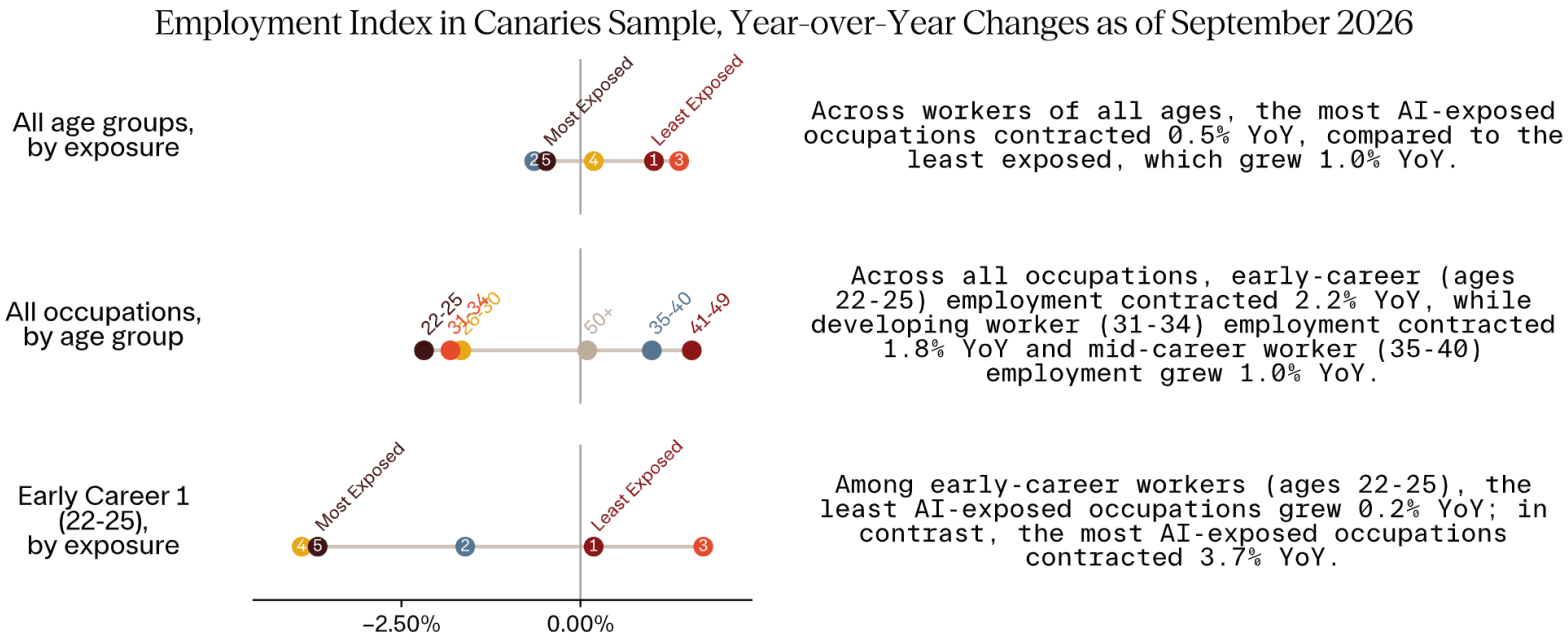
In contrast to software developers and customer service representatives, the relational and physical work of home health aides is minimally exposed to AI under most exposure frameworks: this cluster of occupations spans the least and second least exposed groups. Early-career workers (22-25) now see larger employment gains than the other age groups.

Data



Source: Stanford Digital Economy Lab, ADP Research ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

Canaries in the coal mine



Most-exposed occupations Employment approximately 11% below November 2022 levels.

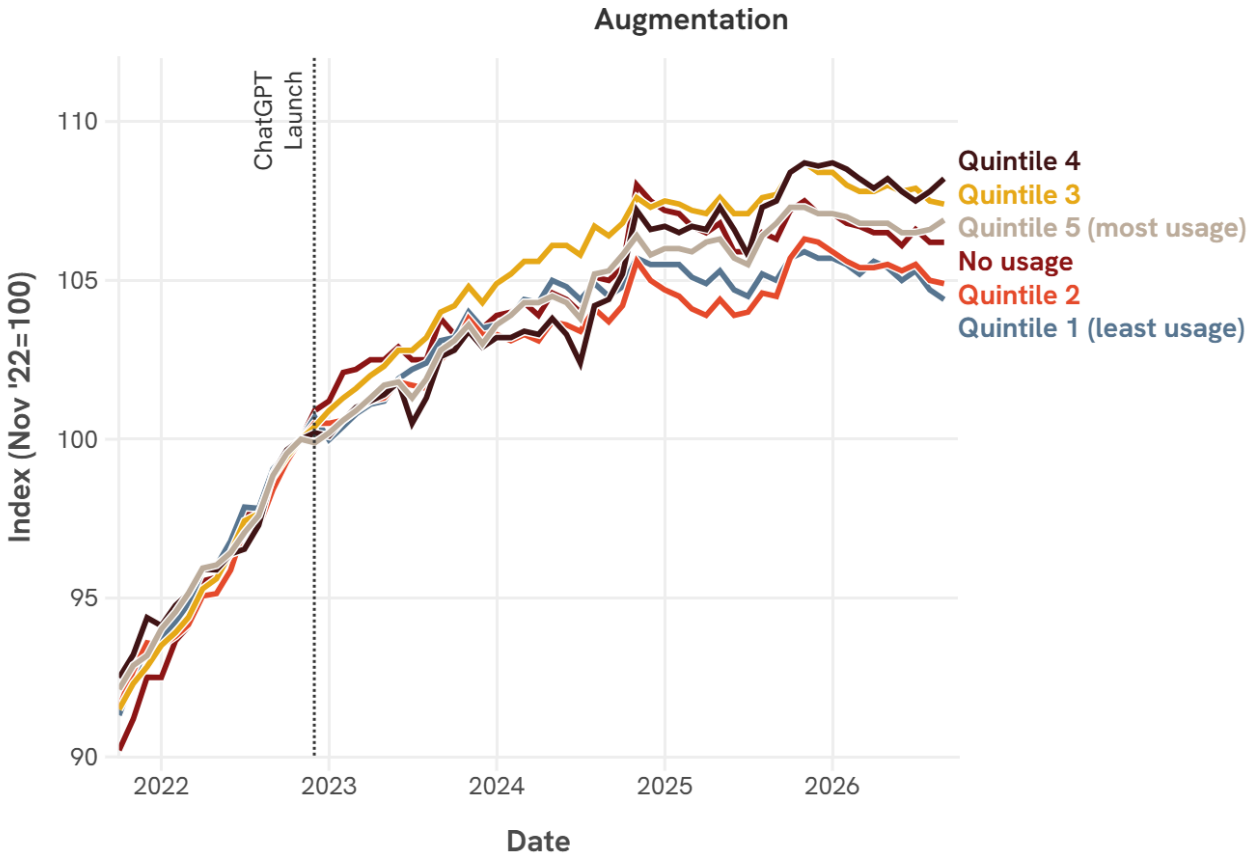
Less-exposed occupations Employment approximately 10% above November 2022 levels.

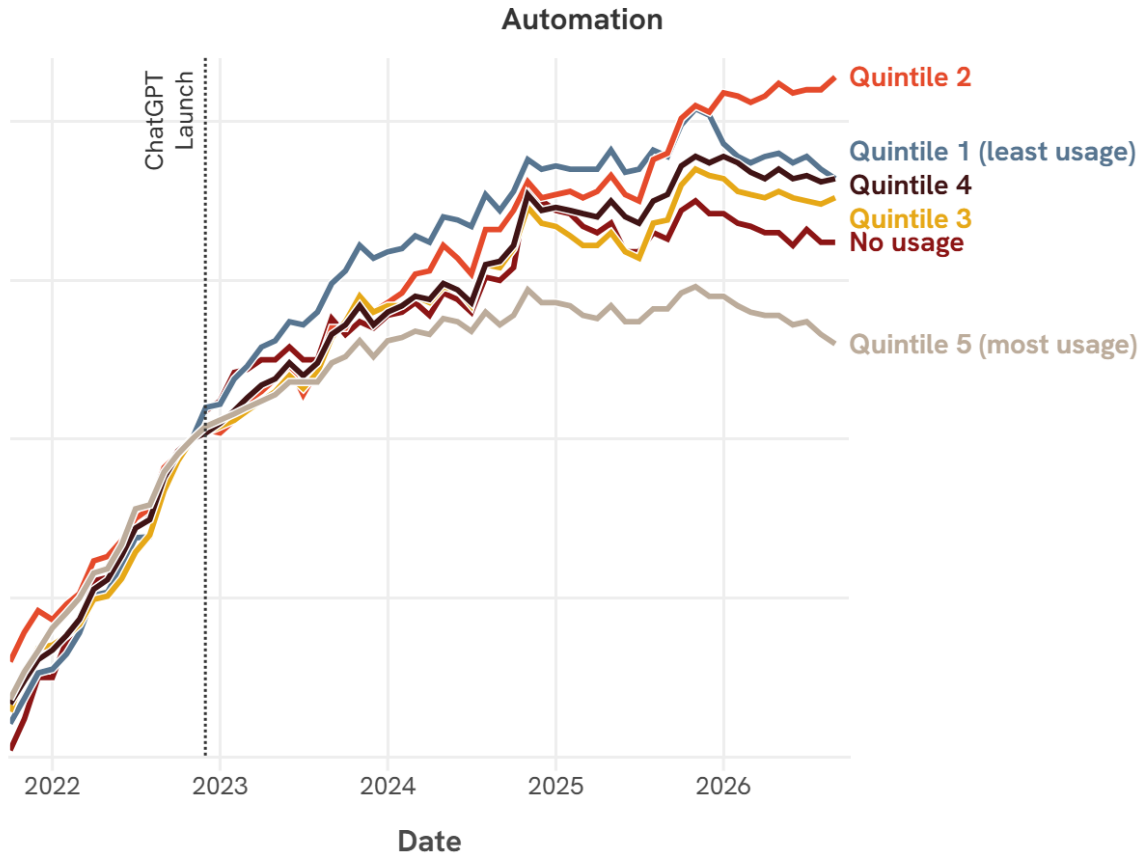
https://digitaleconomy.stanford.edu/app/uploads/2026/08/Canaries_August2026.pdf

Employment Index in Canaries Sample by Age Group and Anthropic Augmentation/Automation Ratios

We group employees by their age and their occupations' Anthropic Economic Index augmentation and automation ratios. Among early-career workers, the automation ratio shows a noticeable relationship with employment trends: occupations with a higher automation ratio see declines or more muted increases in the employment index. In contrast, the augmentation ratio does not show a clear relationship with employment trends for early-career workers.

Select an age group: Early Career 1 (22-25) | Early Career 2 (26-30) | Developing (31-34) | Mid-Career 1 (35-40) | Mid-Career 2 (41-49) | Senior (50+) | All Ages





Source: Stanford Digital Economy Lab, ADP Research, Anthropic Economic Index ▪ Results use a five-year balanced sample of firms using ADP payroll services and include all workers matched to an occupation code.

Which Occupations Are Most and Least Exposed to AI?

MOST EXPOSED: TOP 25

#	Occupation	Score
1	**Mathematicians**	**1.00**
2	Proofreaders and copy markers	0.98
3	Blockchain engineers	0.97
4	Correspondence clerks	0.96
5	Court reporters and captioners	0.96
6	**Computer programmers**	**0.95**
7	Data warehousing specialists	0.94
8	Database administrators	0.94
9	Web developers	0.93
10	Web administrators	0.91
11	Telephone operators	0.89
12	Data entry keyers	0.89
13	Statistical assistants	0.89
14	**Poets, lyricists and creative writers**	**0.89**
15	Interpreters and translators	0.88
16	Database architects	0.88
17	Software QA analysts and testers	0.88
18	Medical transcriptionists	0.88
19	**Software developers**	**0.87**
20	Bioinformatics technicians	0.87
21	Writers and authors	0.87
22	Payroll and timekeeping clerks	0.84
23	Web and digital interface designers	0.83
24	Insurance claims and policy clerks	0.83
25	Telemarketers	0.81

FAMILIAR OCCUPATIONS ALONG THE SCALE

Occupation	Score	Rank	Fifth
Software developers	**0.87**	**19**	**1**
Secretaries and admin. assistants	0.64	73	1
Customer service representatives	0.57	121	1
Accountants and auditors	0.56	129	1
Paralegals and legal assistants	**0.53**	**182**	**1**
Graphic designers	0.50	243	2
Financial and investment analysts	0.46	321	2
Lawyers	**0.43**	**383**	**3**
Retail salespersons	0.34	483	3
Registered nurses	0.33	493	3
Elementary school teachers	0.31	515	3
Heavy truck drivers	0.28	551	3
Cashiers	0.21	629	4
Stockers and order fillers	0.18	655	4
Restaurant cooks	0.09	758	5
Home health aides	0.08	785	5
Construction laborers	0.05	821	5
Janitors and cleaners	0.00	882	5

Fifth: 1 = most exposed, 5 = least (equal numbers of occupations).

LEAST EXPOSED: SCORE OF 0

48 occupations score exactly 0, among them:

Extraction: roustabouts, roof bolters, continuous mining machine operators, wellhead pumpers

Metal and foundry: pourers and casters, foundry mold and coremakers

Meat processing: slaughterers and meat packers; meat, poultry and fish cutters

Mechanics and repair: bus and truck mechanics, motorcycle mechanics, tire repairers, power-line installers

Cleaning and kitchen: dishwashers, janitors, maids and housekeepers

Other: riggers, hoist and winch operators, packaging machine operators

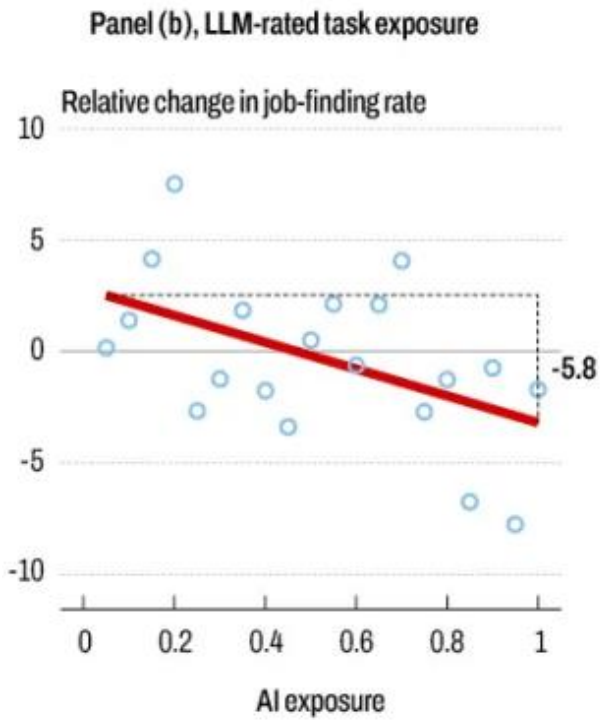
CUTOFFS FOR EQUAL FIFTHS

Fifth	Score range
1 (most)	above ~0.52
2	~0.43 to 0.52
3	~0.28 to 0.43
4	~0.11 to 0.28
5 (least)	below ~0.11

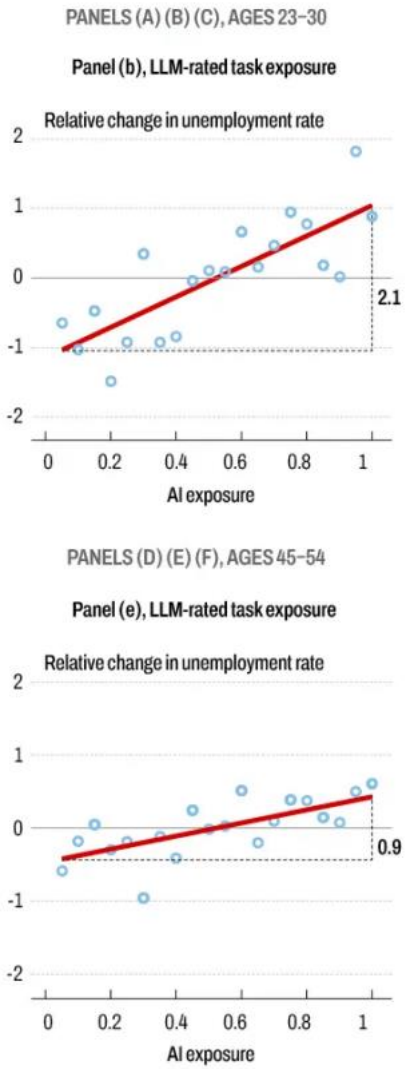
Score ≈ share of an occupation's tasks an LLM could speed up by at least half (exposure, not observed use).

*Source: Eloundou, Manning, Mishkin & Rock, "GPTs are GPTs," Science (2024), occupation-level GPT-4 β exposure ratings (github.com/openai/GPTs-are-GPTs). Ranks are of 923 O*NET occupations.*

Cleveland Fed -- Decline in job-finding rate due to AI.



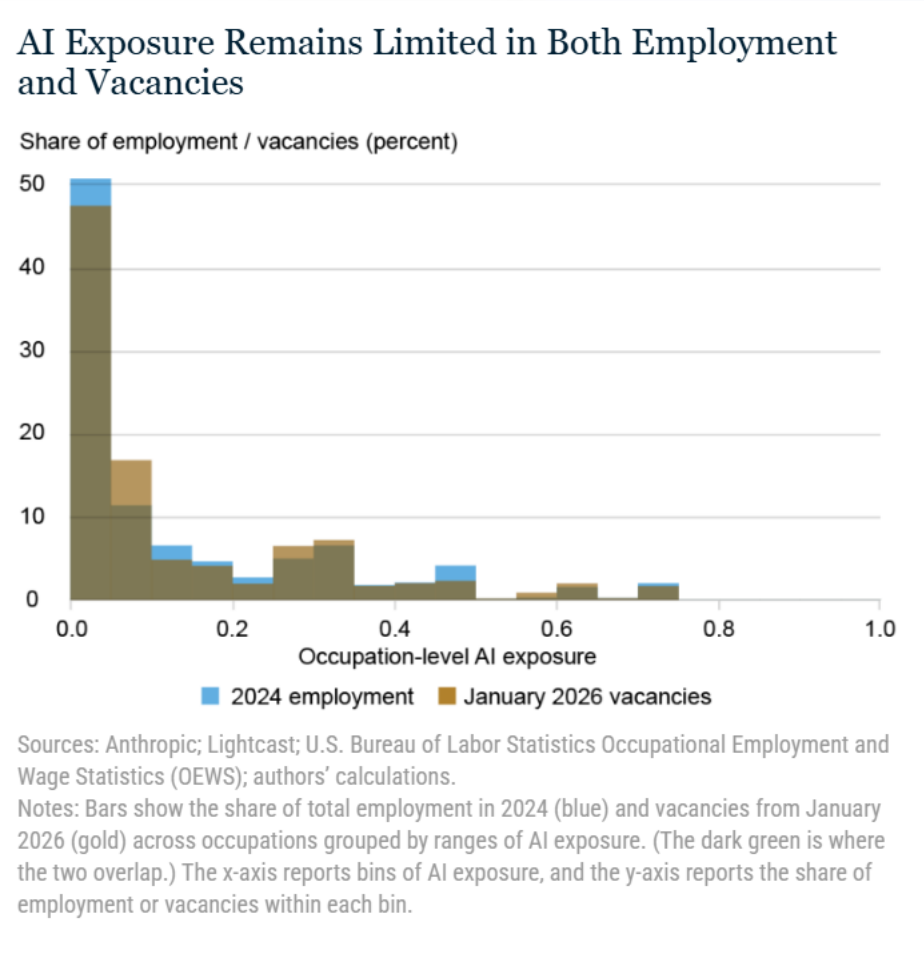
Cleveland Fed (11/9/2026): "Occupations with the highest AI exposure experienced job-finding rates that declined by roughly six percentage points more than occupations with the lowest exposure. This is a substantial shift."

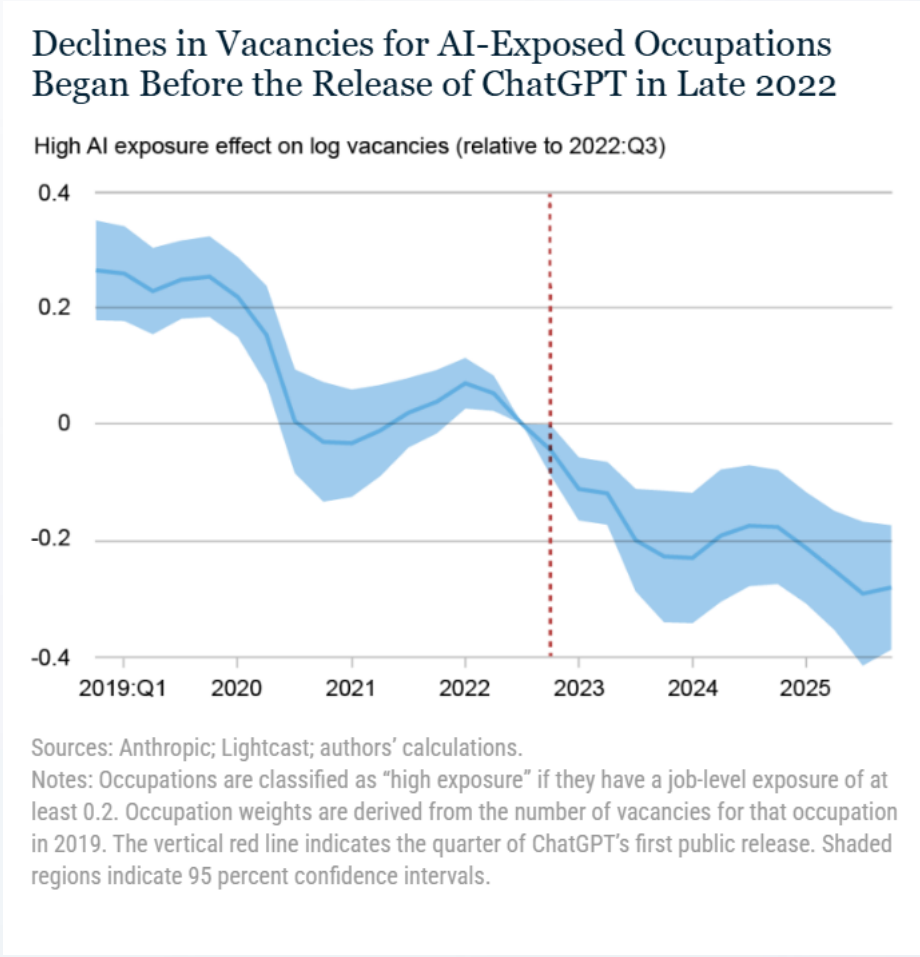


Younger workers experience a more pronounced positive relationship between AI exposure and unemployment. The fitted regression lines in the top row imply that moving from the lowest to the highest AI exposure group increases the predicted unemployment rate for younger workers by about two percentage points. For workers aged 45 and older (bottom row), the relationship between AI exposure and unemployment is positive but muted.

Our findings for jobs vulnerable to potential AI-for-human substitution are consistent with what some observers have called a "no hire, no fire" labor market dynamic, in which firms respond to uncertainty about AI's transformative potential, not by laying off existing workers, but by quietly pulling back on new hiring

New York Fed – decline predates AI, and fewer jobs are exposed





Cyclical downturn concentrated in exposed occupations

Audoly, Guerin & Topa, Liberty Street Economics (May 14, 2026); Fairlie & Wu, NBER w35796.

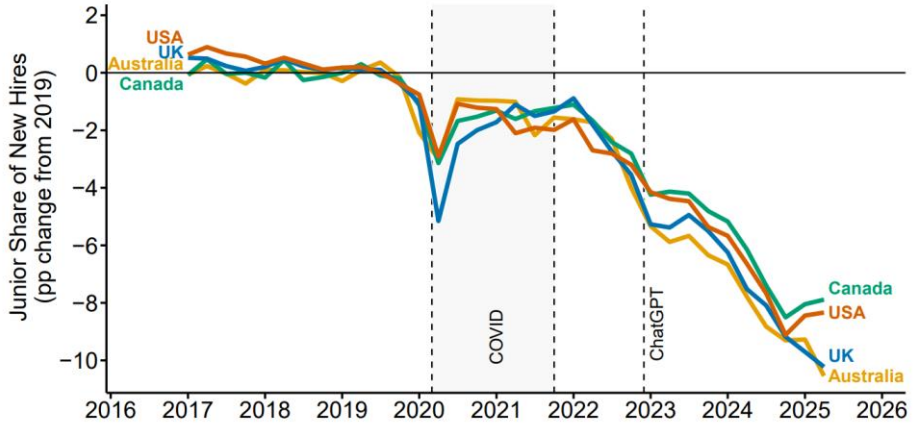
Work from home is to blame!

The Broken Ladder:
AI, Remote Work, and Early-Career Hiring

Peter John Lambert† & Yannick Schindler‡

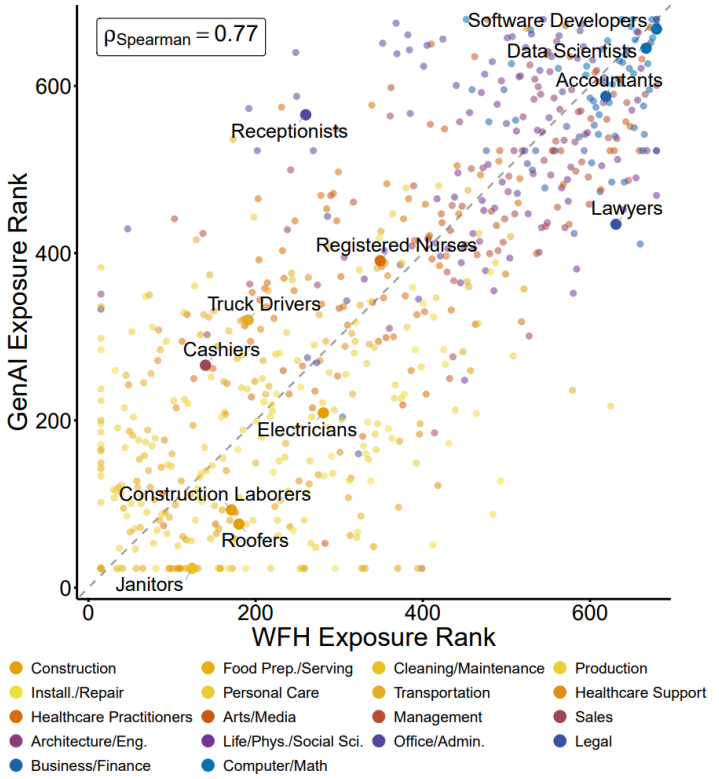
May 2026*

Figure 1: Junior-share of Hiring Has Declined Sharply Across Four Anglophone Economies, 2017–2025

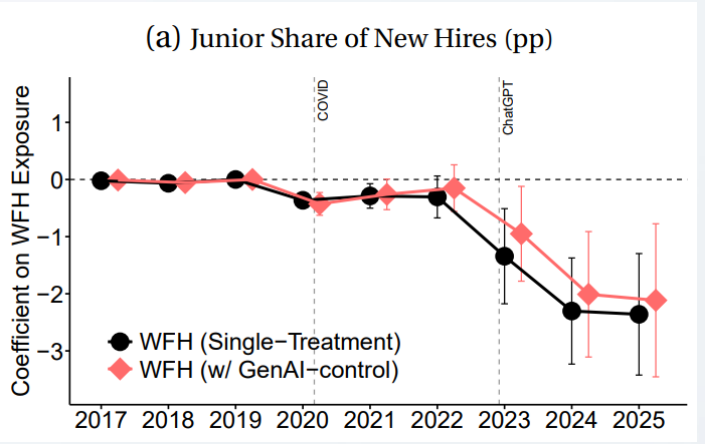


Note: Figure depicts the "junior-share" of new hires for the US, UK, Canada, and Australia, expressed as the percentage-point change from each country's 2019 mean. Data comes from Revelio Labs (see Section 4 for detailed description). "Junior" is defined using the bottom-two groups of a six-level scale which accounts for the job description and each worker's past experience at the beginning of their tenure in a new position. All series are constructed as a weighted arithmetic mean of 3-digit-occupation junior shares, using the US 2019 3-digit O*NET share of all hires as the (fixed) weight, to net out temporal and cross-country compositional differences. The 2019 (reweighted) junior share of new hires is 66.9% in the US, 64.9% in the UK, 69.3% in Canada, and 65.8% in Australia. Series are at quarterly frequency and seasonally adjusted using the U.S. Census Bureau's X-13ARIMA-SEATS procedure applied separately to each country's series. Dashed vertical lines mark the COVID-19 pandemic and the public release of ChatGPT.

Figure 2: WFH and GenAI Exposure Are Highly Correlated Across 6-Digit O*NET-SOC Occupations



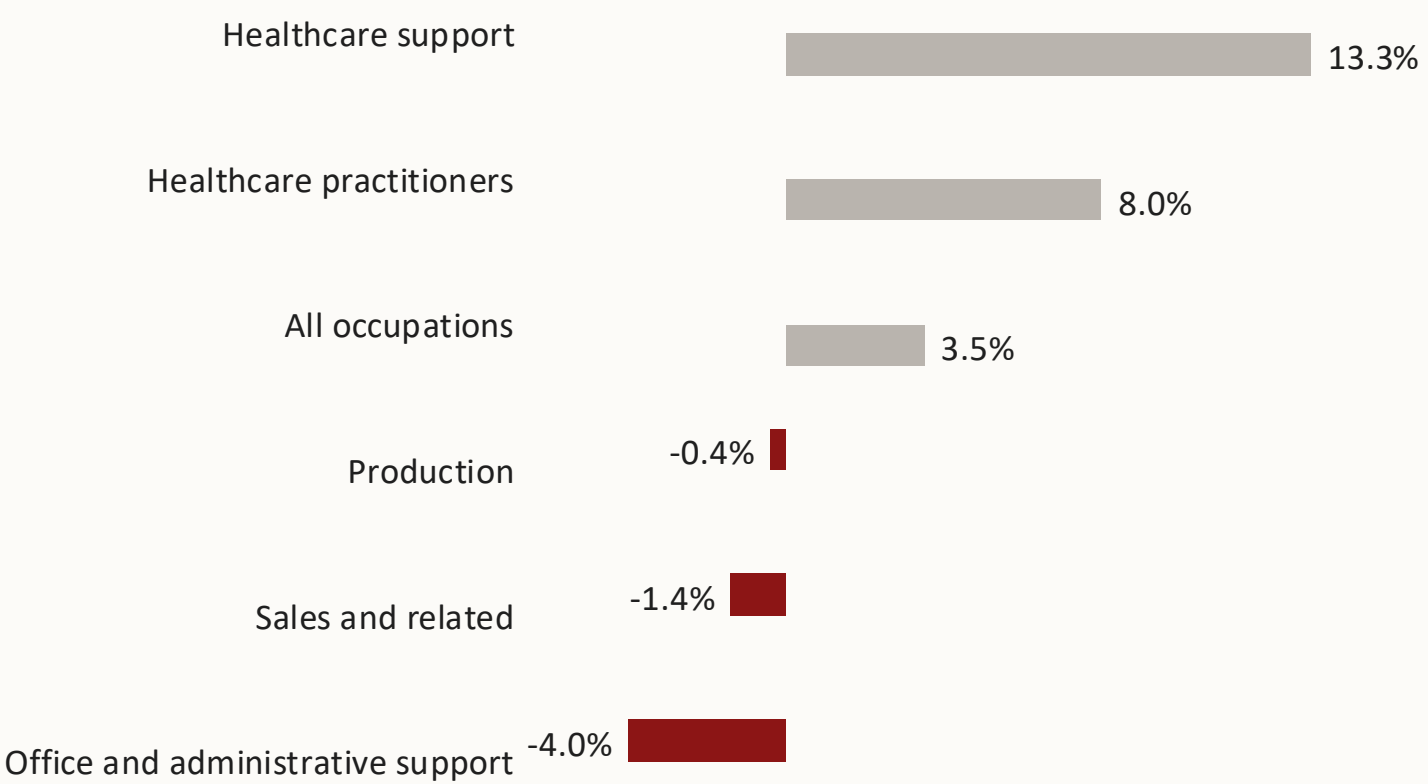
Note: This figure depicts 683 6-digit O*NET-SOC occupations showing their WFH exposure from Hansen et al. (2023) and their GenAI exposure from Eloundou et al. (2024) (see Section 4 for a discussion of both measures). The figure plots cross-sectional ranks of each occupation on the two indices; the dashed line is the 45-degree reference. The Spearman rank correlation is $\hat{\rho}_S = 0.77$, and the OLS regression of standardized GenAI on standardized WFH exposure gives $\hat{\beta} = 0.71$ (s.e. 0.03, $R^2 = 0.50$).



IV. Forecasts



BLS 2025–35: office and administrative support shrinks most



- Office & admin: −752,100 jobs, the largest loss of any group; BLS cites AI-powered automation
- Sales: e-commerce and AI sales tools
- Arts, design and media: generative AI “may limit demand”
- BLS: uncertainty “remains very high”

Source: BLS Employment Projections 2025–35 (Aug. 27, 2026).

BLS 2025–35: software developers grow, programmers fall

Nurse practitioners	**+41.0%**
Data scientists	**+34.6%**
Software developers	**+10.2%**
Computer programmers	**−7.3%**
Customer service representatives	**about −5%**
Office/administrative support (group)	**−4.0%**

All occupations: +3.5%. Occupational projections are not causal estimates of AI.

Source: U.S. Bureau of Labor Statistics, Employment Projections 2025–35 (Aug. 2026).

McKinsey Workforce in Motion: modeled U.S. labor demand, 2025-2035

Full-time-equivalent (FTE) labor demand; scenario includes AI, other automation, demographic and economic changes.

MODEL COMPONENT	FTE LABOR-DEMAND CHANGE / POSITION	HOW MGI USES THE FIGURE
Automation reduction	**-36 million FTE**	About 21% of current work hours; across AI and other automation.
Additional demand	**+41 million FTE**	Demand from the human economy, physical/tech buildout and other trends.
Within-occupation offset	**25 million FTE**	Growth in the same occupations may offset part of reduced task demand.
Cross-occupation movement	**11 million workers**	Estimated number requiring moves to other occupations; scenario range 6M-16M.
New openings from outside	**16 million FTE**	Growing occupations needing workers from other occupations or new entrants.
Technical capability	**Up to half of work hours**	Potentially automatable by 2035 in some scenarios, not actual job losses.

The 36M and 41M figures are gross modeled changes, not counts of layoffs and hires. Offsets depend on timing within occupations.

Source: McKinsey Global Institute (Sept. 29, 2026), Workforce in Motion.

McKinsey: additional results on occupational transitions

Labor mobility and credentials under MGI adoption scenarios, 2025-2035.

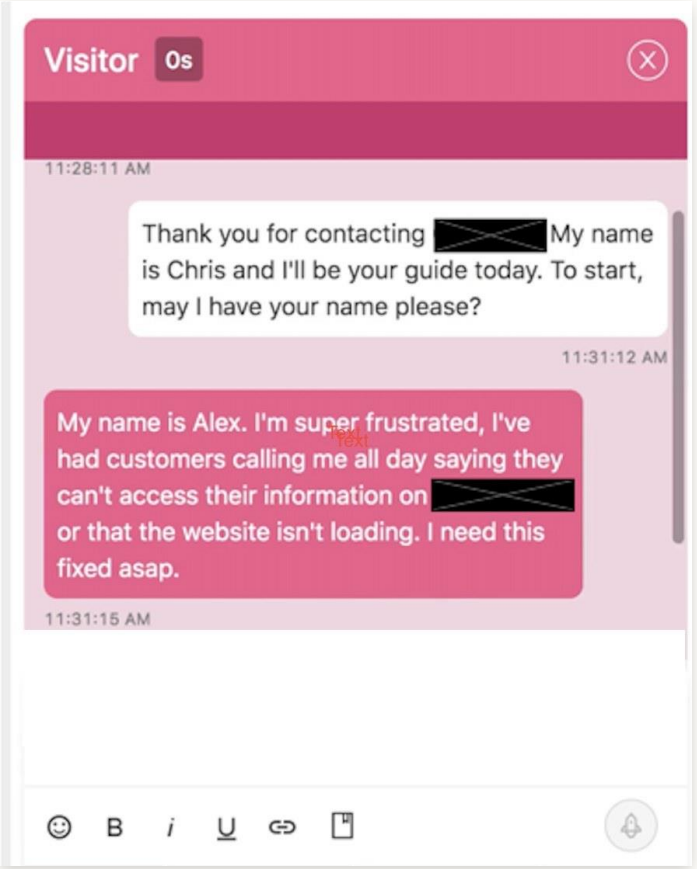
MEASURE	FIGURE / FINDING	PERIOD / DEFINITION
Estimated workers needing new occupations	11 million, range 6 million-16 million.	Range depends on automation pace and broader labor demand.
Occupational-group transitions	About 770,000 workers annually, 2025-2035.	MGI estimates 3.6 times non-COVID historical average.
Historical comparison	About 215,000 cross-group transitions per year in non-COVID reference years.	Cited MGI series, not a forecast.
Transition pathways	About one in seven workers have a direct skill-matched route into growing work.	MGI pathway modeling with wage and credential barriers.
Credential requirements	Approximately 85% of growing jobs require a credential or certification.	Qualification constraint in simulated transition paths.
Workforce composition	Roughly 60% of growing employment in upper two wage quintiles; >70% of declining in bottom two.	MGI wage-quintile classification, scenarios.

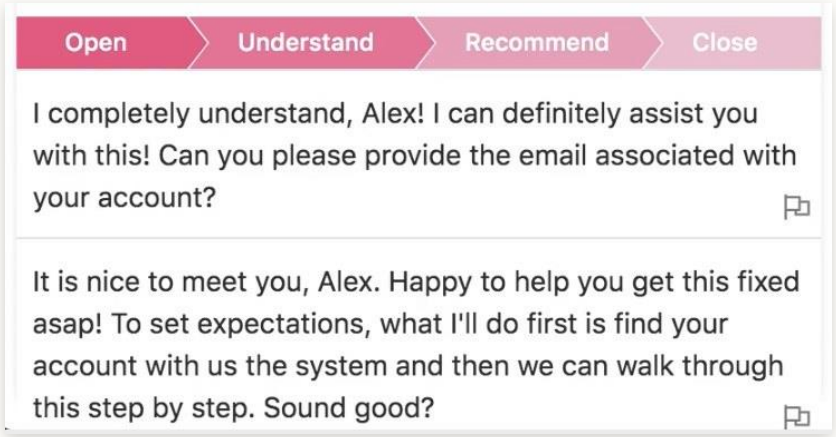
Source: McKinsey Global Institute, Workforce in Motion (2026), chapters 1-2.

Case Study #1: Call centers



The call-center study: AI suggests, the agent decides





- 5,172 agents at a Fortune 500 software firm; staggered rollout
- Suggestions drawn from responses most correlated with success
- Here: a friendly, reassuring rapport

Brynjolfsson, Li & Raymond (2025), "Generative AI at Work," Quarterly Journal of Economics 140(2).

Augmenting with AI increases productivity in call center

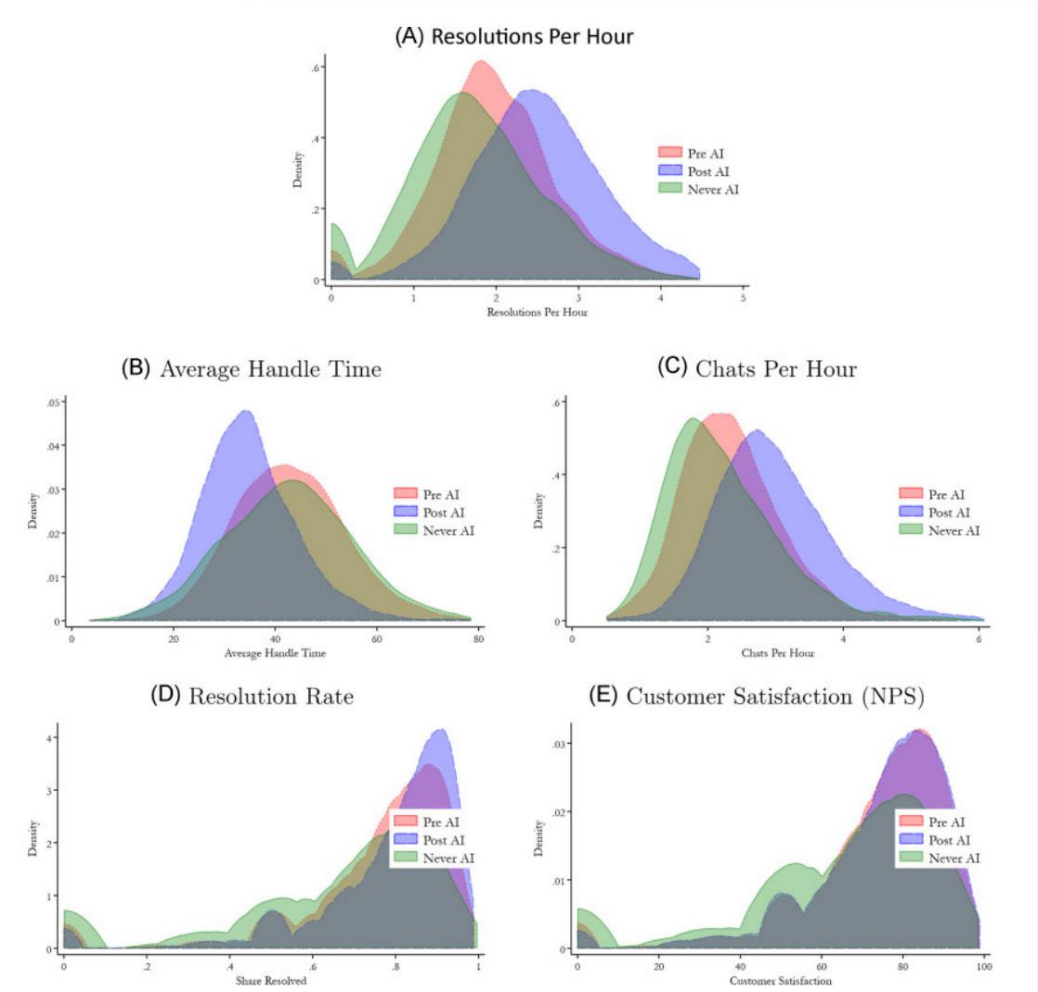


FIGURE I

Raw Productivity Distributions, by AI Treatment

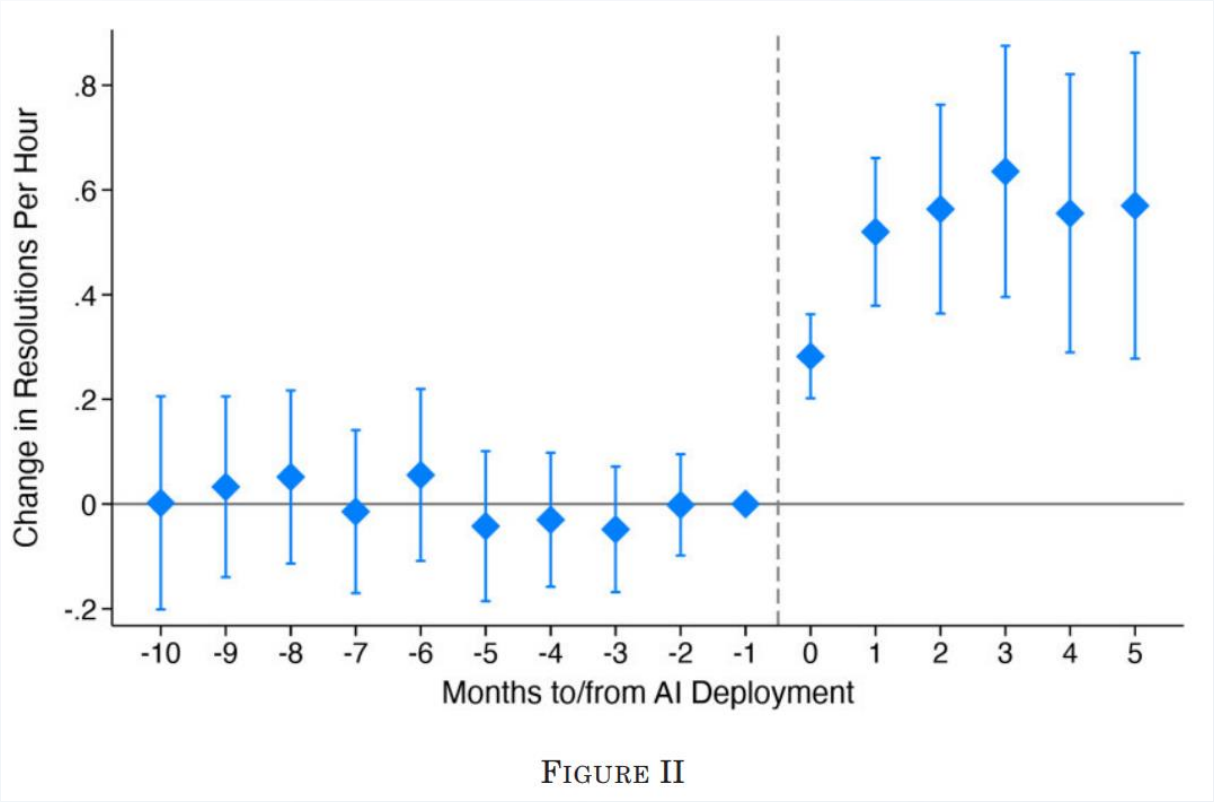
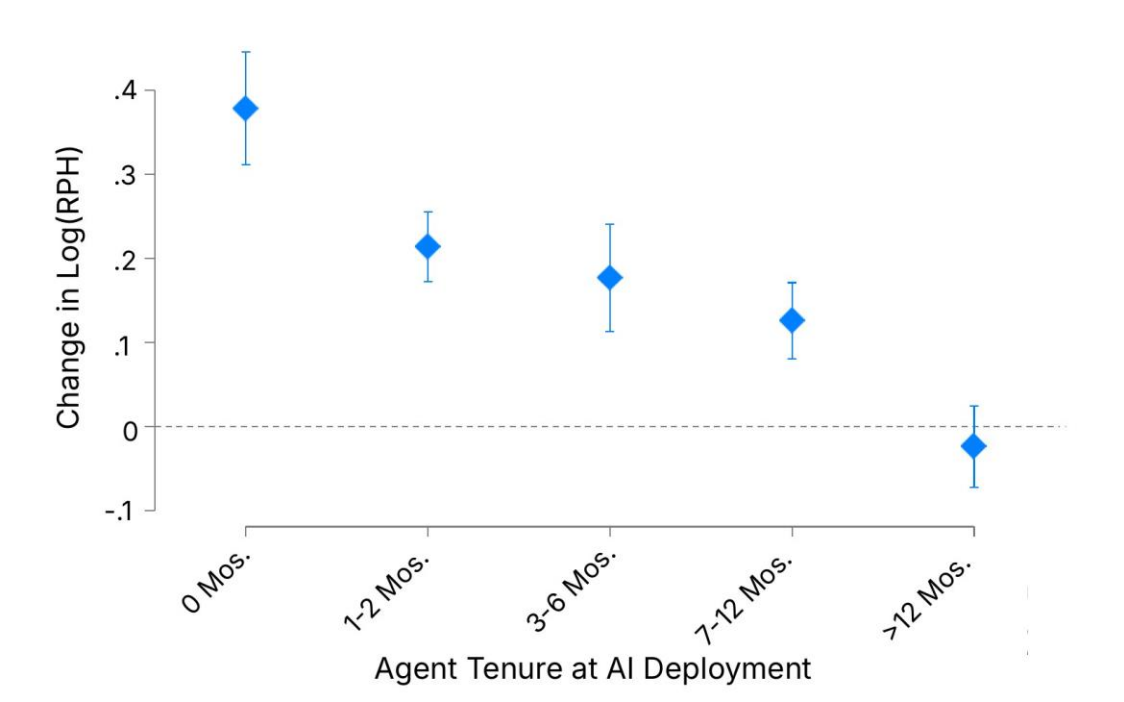
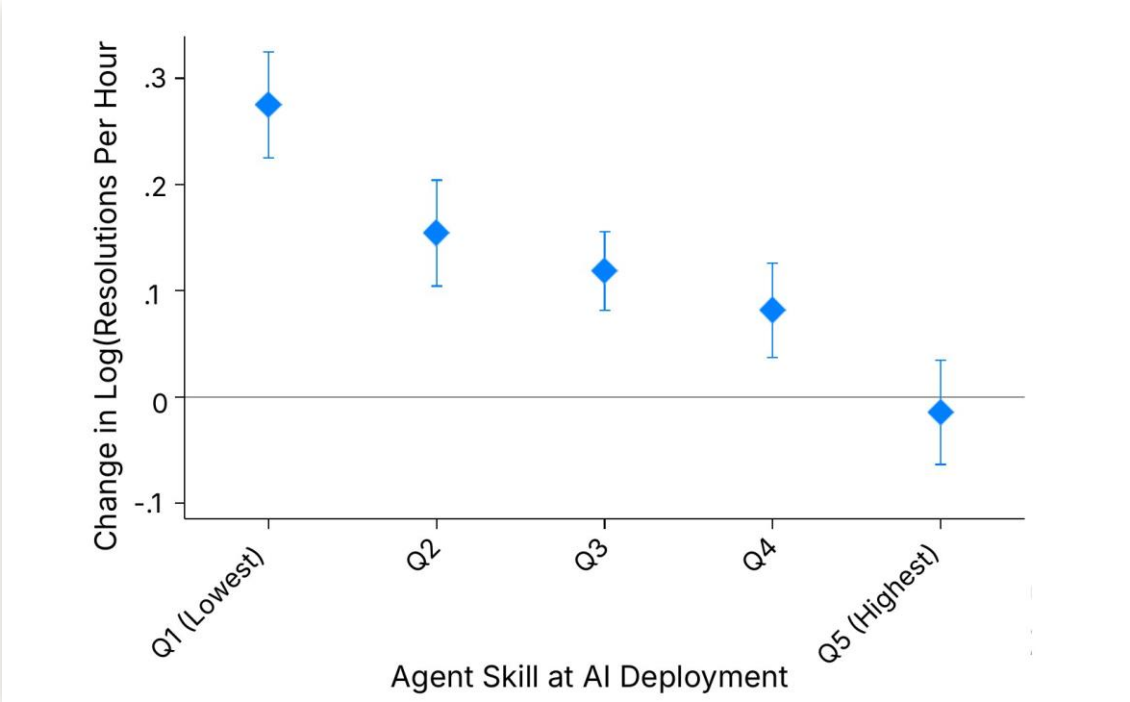


FIGURE II

Highest returns for the newest and lowest-skill agents



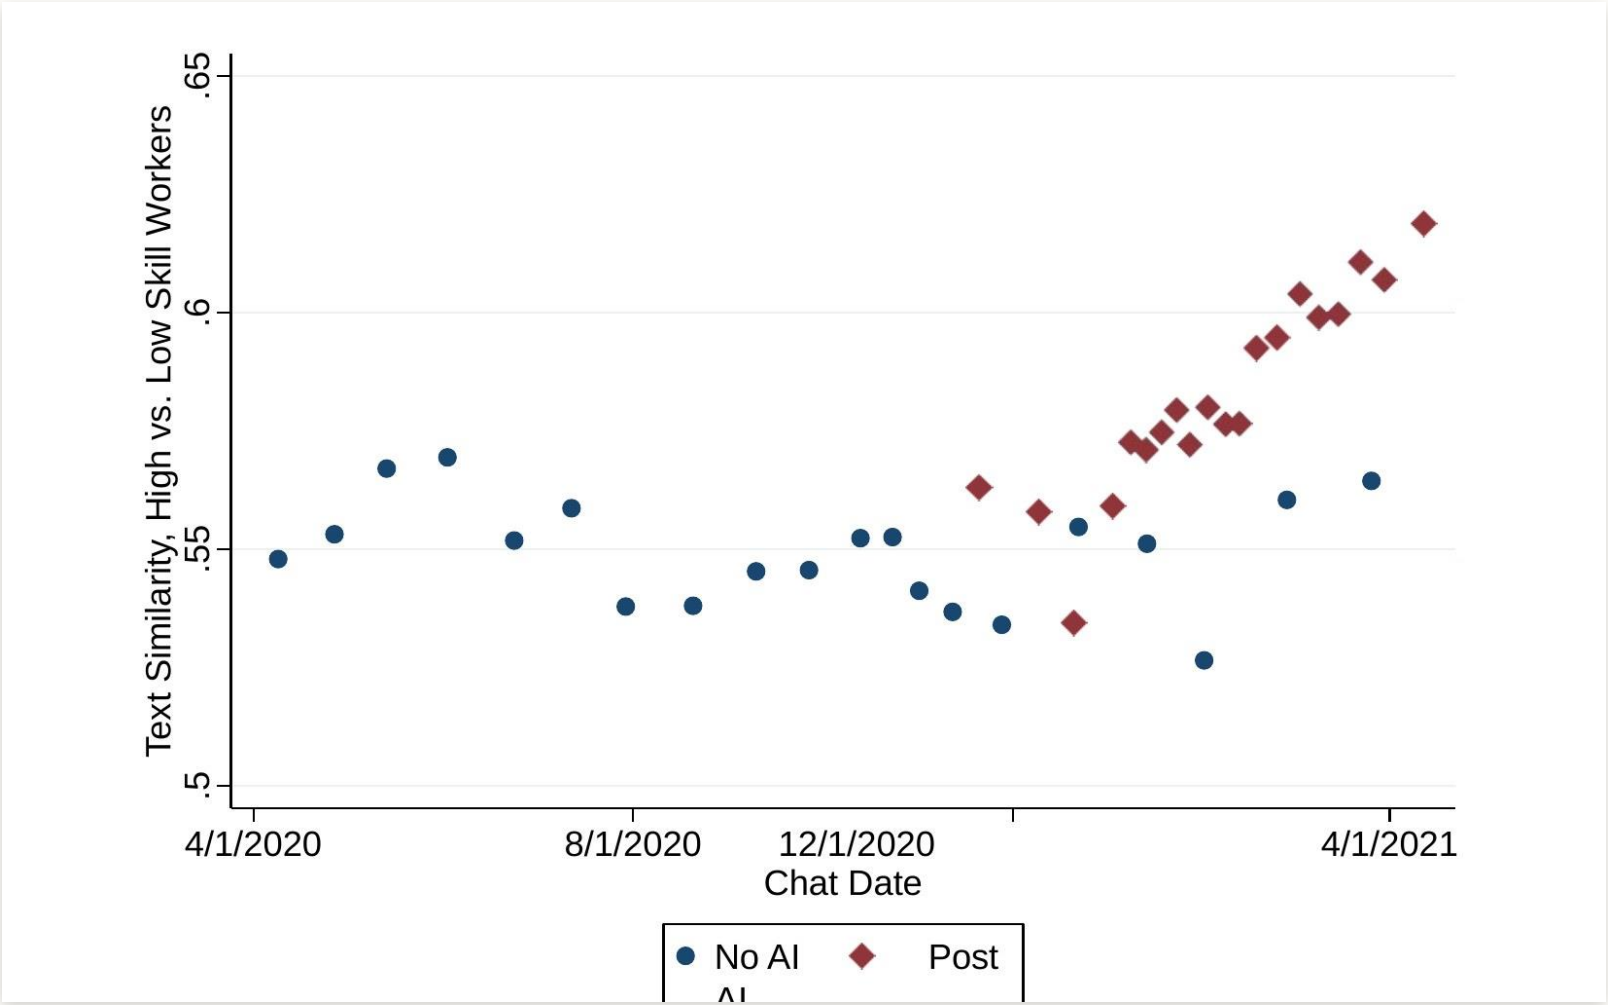
By tenure at deployment



By skill at deployment

Brynjolfsson, Li & Raymond (2025), QJE.

After AI access, low-skill agents sound more like high-skill agents



Brynjolfsson, Li & Raymond (2025), QJE. Text similarity, high v. low-skill workers.

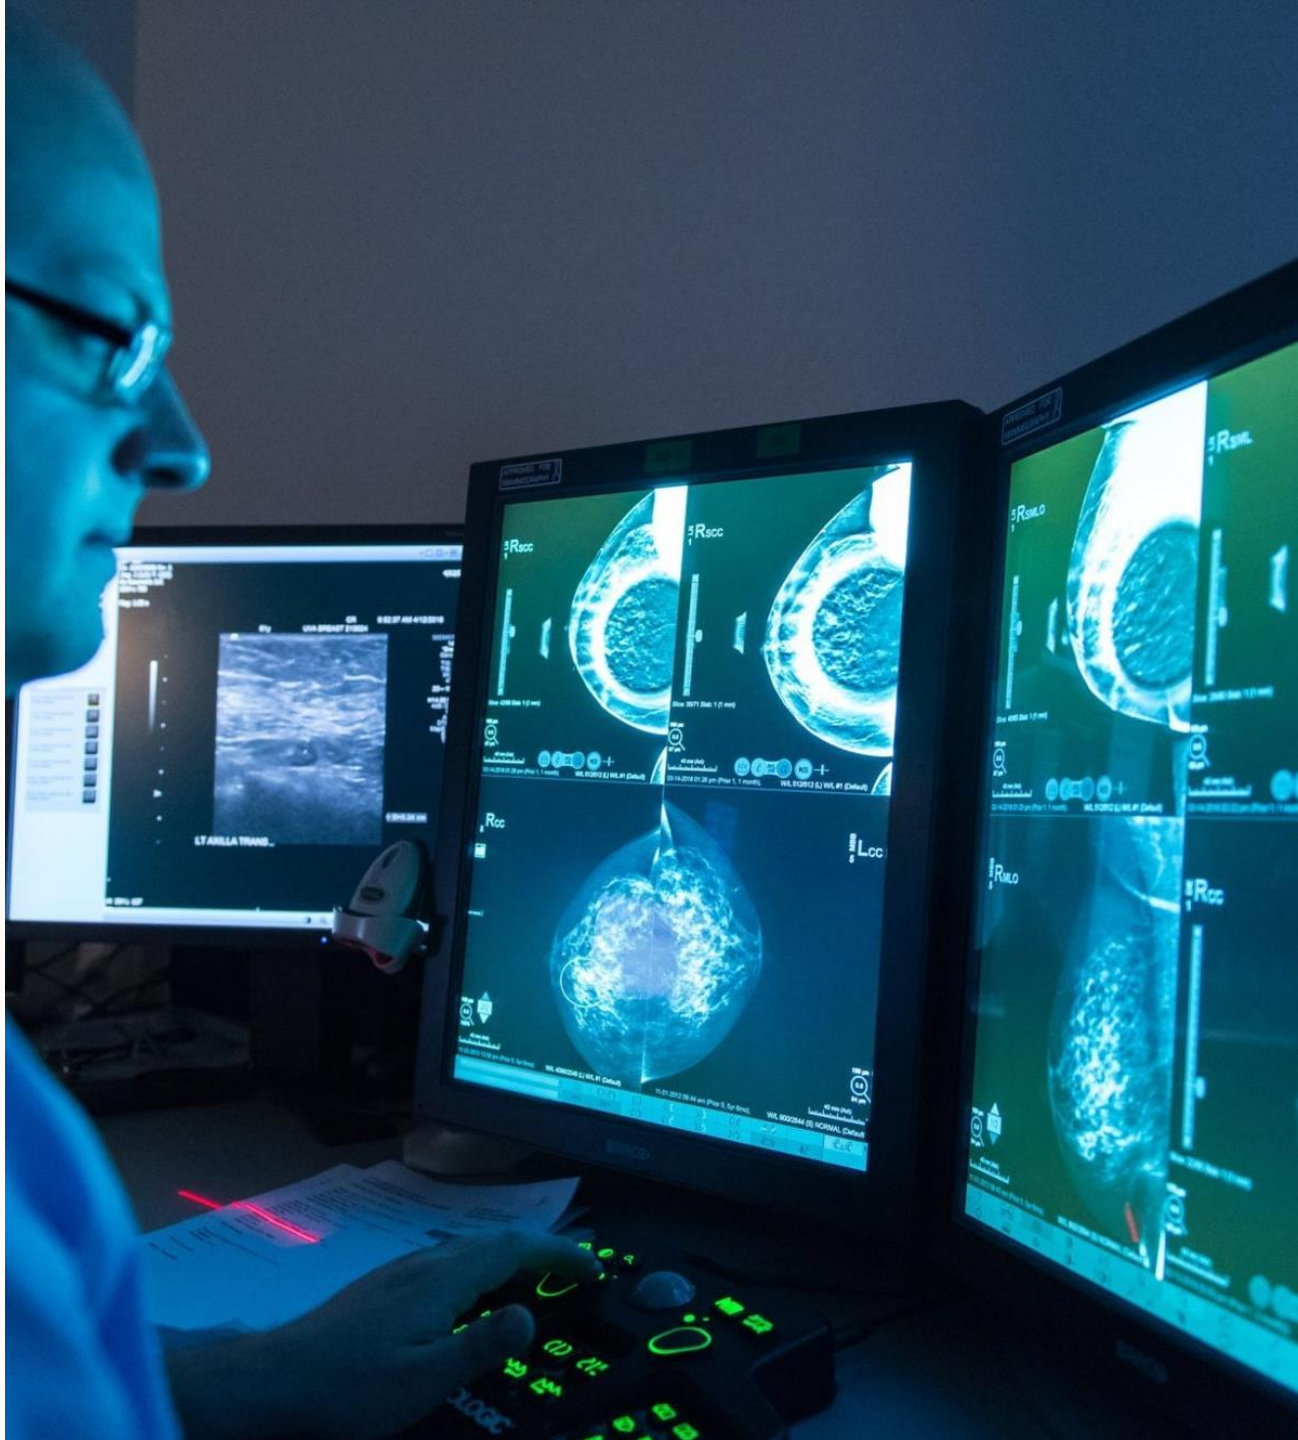
Customer service: employment and automation evidence

Observed staffing change, a separate ten-year forecast, and a company report.

MEASURE	FIGURE / FINDING	PERIOD / DEFINITION
U.S. employment, May 2024-May 2025	Customer-service jobs declined about 130,180 (-4.8%).	BLS occupational employment data as reported by Staffing Industry Analysts.
U.S. employment projection, 2025-2035	BLS projects about -5% (-141,800) from a base near 2.67 million.	Projection, not an AI-attributable employment estimate.
Company case: Klarna	The company said its assistant performed work equivalent to about 700 staff.	Company account; subsequent decision to restore some human service roles.
Automation scope	Routine contacts, common questions and first-pass responses.	Firms differ in deployment and escalation practices.
Workforce outcomes to measure	Headcount, staffing hours, hiring, attrition, offshoring and contact mix.	No single figure isolates AI causation in this slide.

Source: BLS OEWS / OOH; Staffing Industry Analysts.

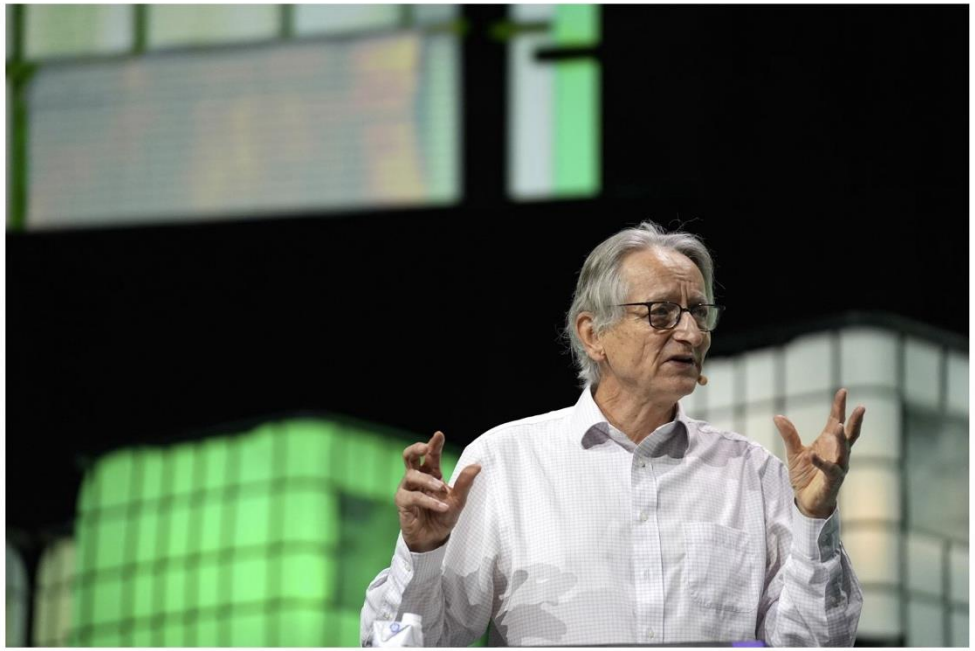
Case study #2: Radiologists



“It’s just completely obvious that within five years deep learning is going to do better than radiologists”

The “Godfather of AI” Predicted I Wouldn’t Have a Job. He Was Wrong.

Nobel Prize winner Geoffrey Hinton said that machine learning would outperform radiologists within five years. That was eight years ago. Now, thanks in part to doomers, we’re facing a historic labor shortage.



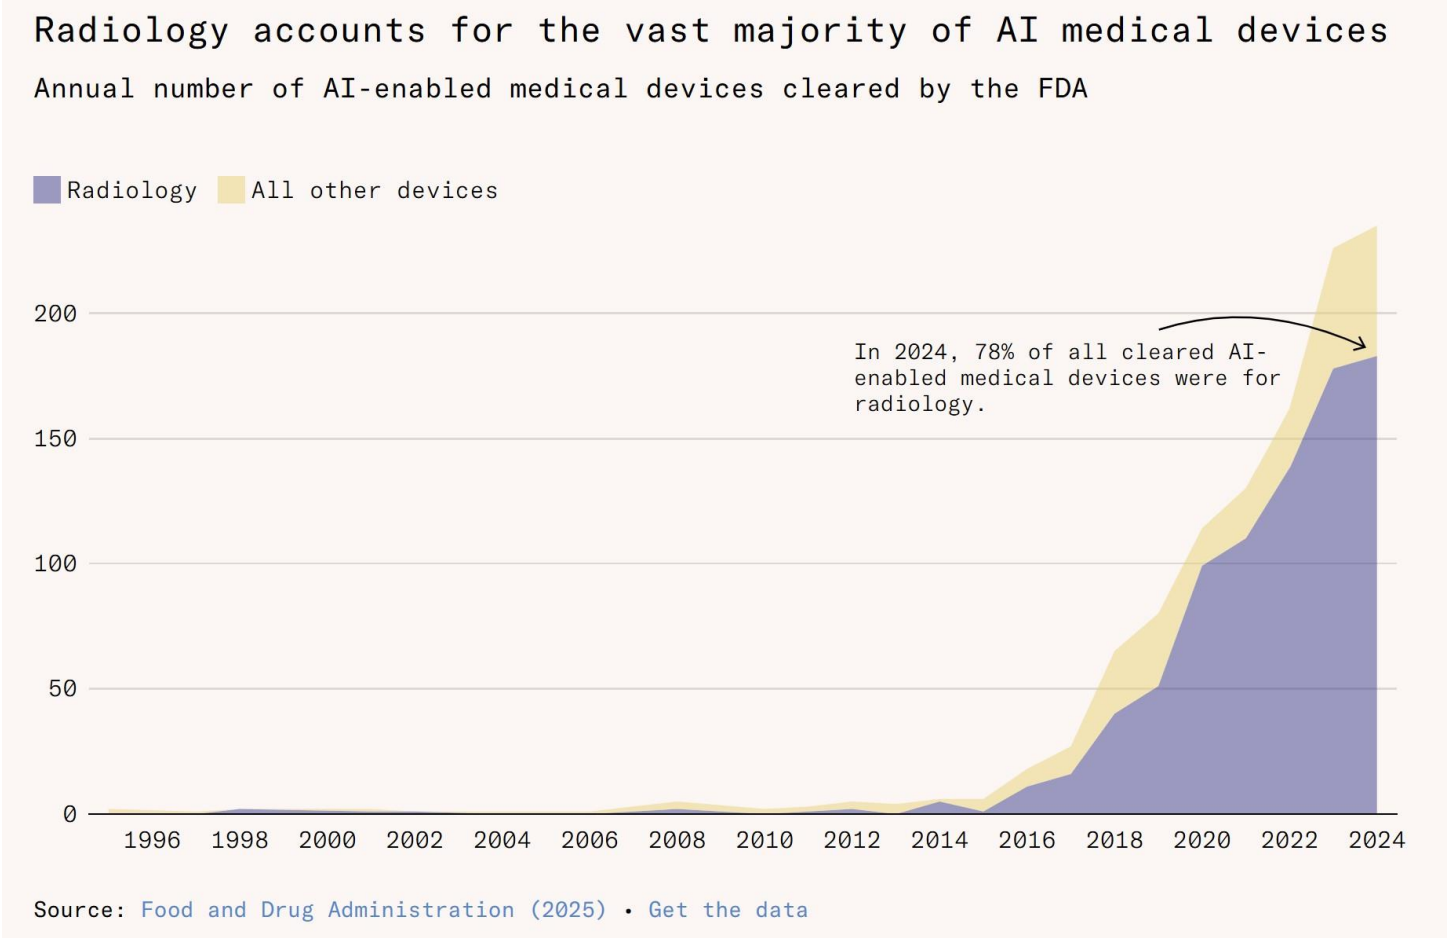
MERT ALPER DERVIS/ANADOLU/GETTY IMAGES

- “… It might be 10 years, but we’ve got plenty of radiologists already.”
- In 2026, Hinton said AI combined with radiologists would improve efficiency and accuracy

Geoffrey Hinton (2016).

/ FIGURE

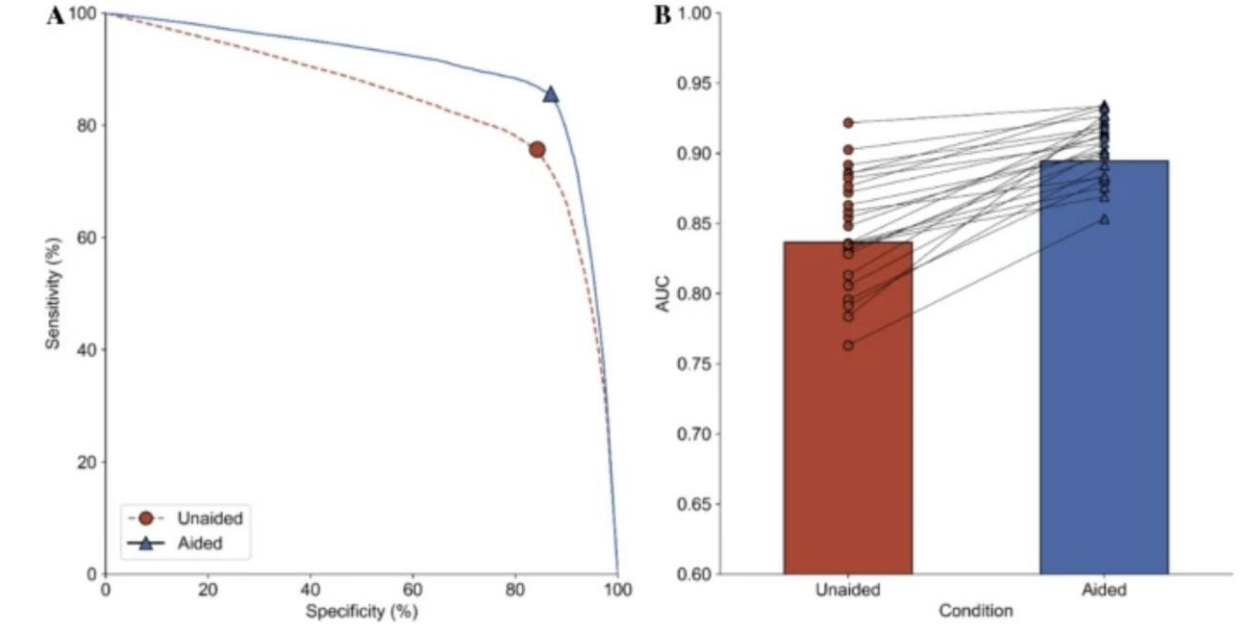
FDA-cleared AI devices: radiology share over time



/ FIGURE

Human and AI diagnostic accuracy in a radiology study

Figure. 3



Physician performance unaided and aided by the AI system. A. ROC curve for all physicians when unaided (dashed red line, red circle) and aided (solid blue line, blue triangle) by the AI system. B . AUCs for all physicians when unaided (red bar and circles) and aided (blue bar and triangles) by the AI system. Solid lines represent an individual physician's improvement when unaided and aided.

Sources: Anderson et al., Scientific Reports 14 (2024).

General conclusions from radiology-AI studies

- AI performs as well as or better than most humans on many routine assessments, but some studies find more false positives
- The best radiologists seem to outperform AI, especially on unusual cases
- Overreliance on AI may make radiologists worse than AI alone or radiologists alone

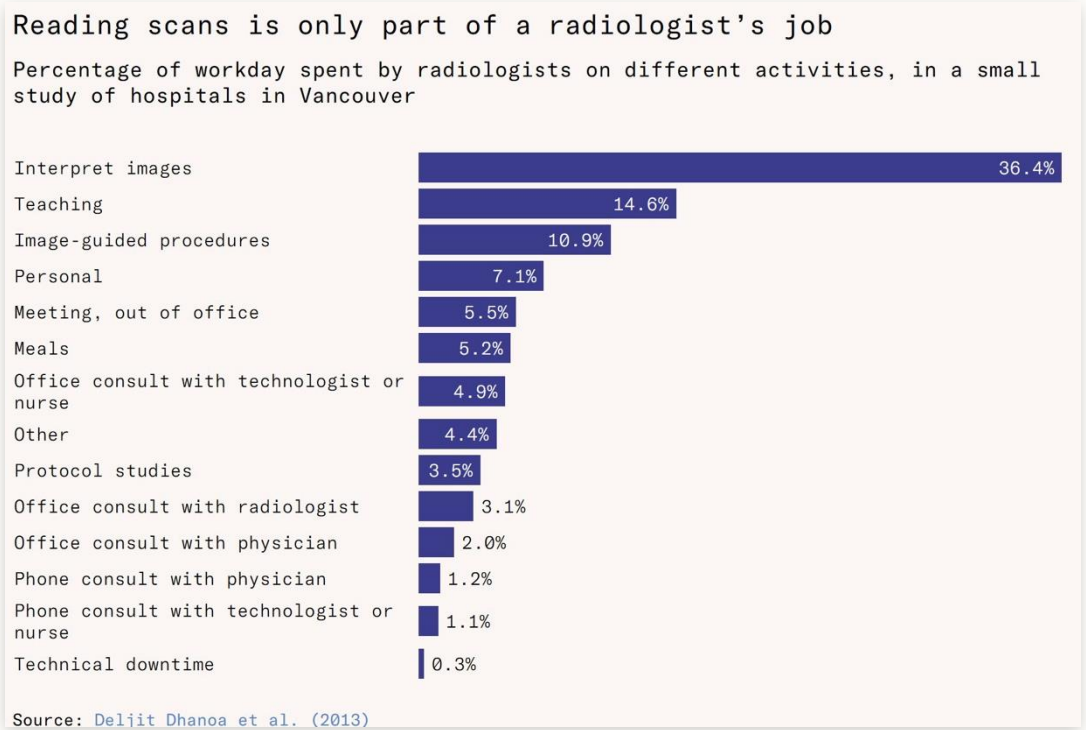
We find that AI assistance does not improve humans' diagnostic quality on average even though the AI predictions are more accurate than approximately 75% of the participants in our experiment. Moreover, the zero average effect cannot be explained by the participants ignoring these predictions – we observe that radiologists' reported probabilities move significantly towards AI predictions when AI assistance is provided. Instead, the zero effect of AI assistance is driven by heterogeneous treatment effects: diagnostic quality increases when the AI is confident (i.e. the predicted probability is close to zero or one) but decreases when the AI is uncertain. In parallel, AI assistance improves diagnostic quality for patient cases in which our participants are uncertain, but decreases quality for patient cases in which our participants are certain. In contrast, providing clinical history does improve diagnostic quality, a result that suggests humans have additional valuable information that has not yet been incorporated into AI predictions.

An upshot of the results is that information available only to radiologists is useful, but humans do not correctly combine their information with AI predictions. In fact, AI predictions reduce predictive preformance for a range of signals. This result cannot be rationalized if

Agarwal, Moehring, Rajpurkar & Salz, NBER w31422.

Mixed results on AI in radiology

- Models beat benchmarks but struggle in hospital conditions
- Regulators and insurers reluctant to approve or cover autonomous models
- Diagnosis is a minority of a radiologist's time; most goes to patients and other clinicians
- Imaging demand projected growth of 17-27% by 2055.
- Employment up 3% per year, 2017-2024.
- Shortage of radiologists

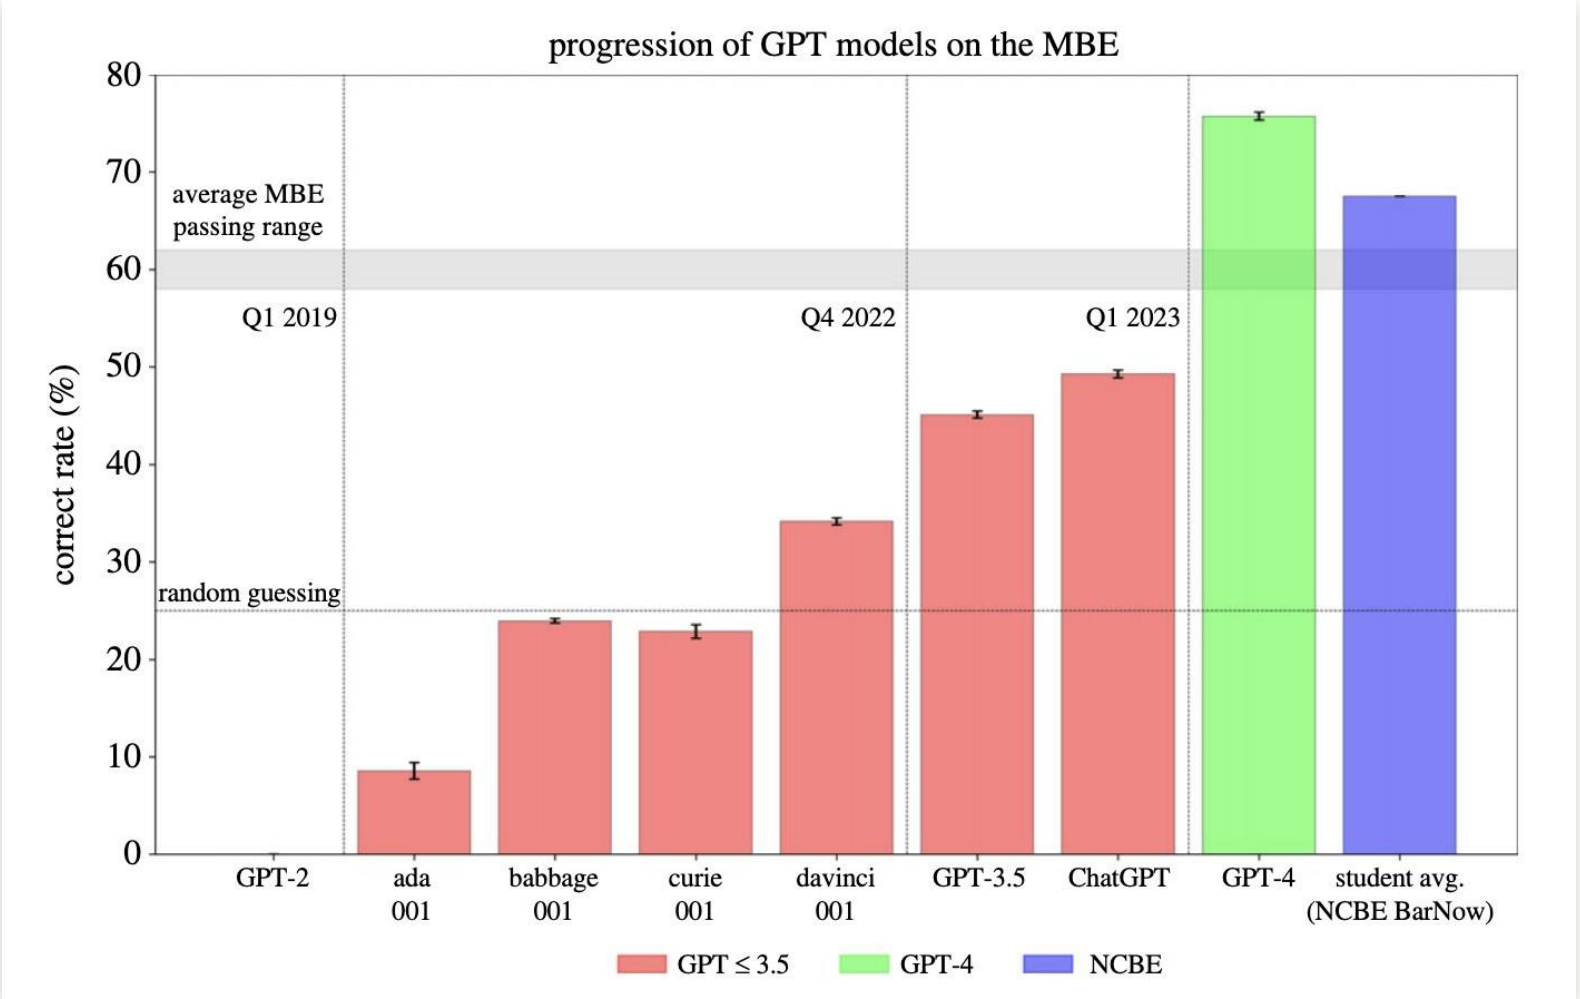


Deena Mousa, "The Algorithm Will See You Now," Works in Progress.

Case study #3: Lawyers

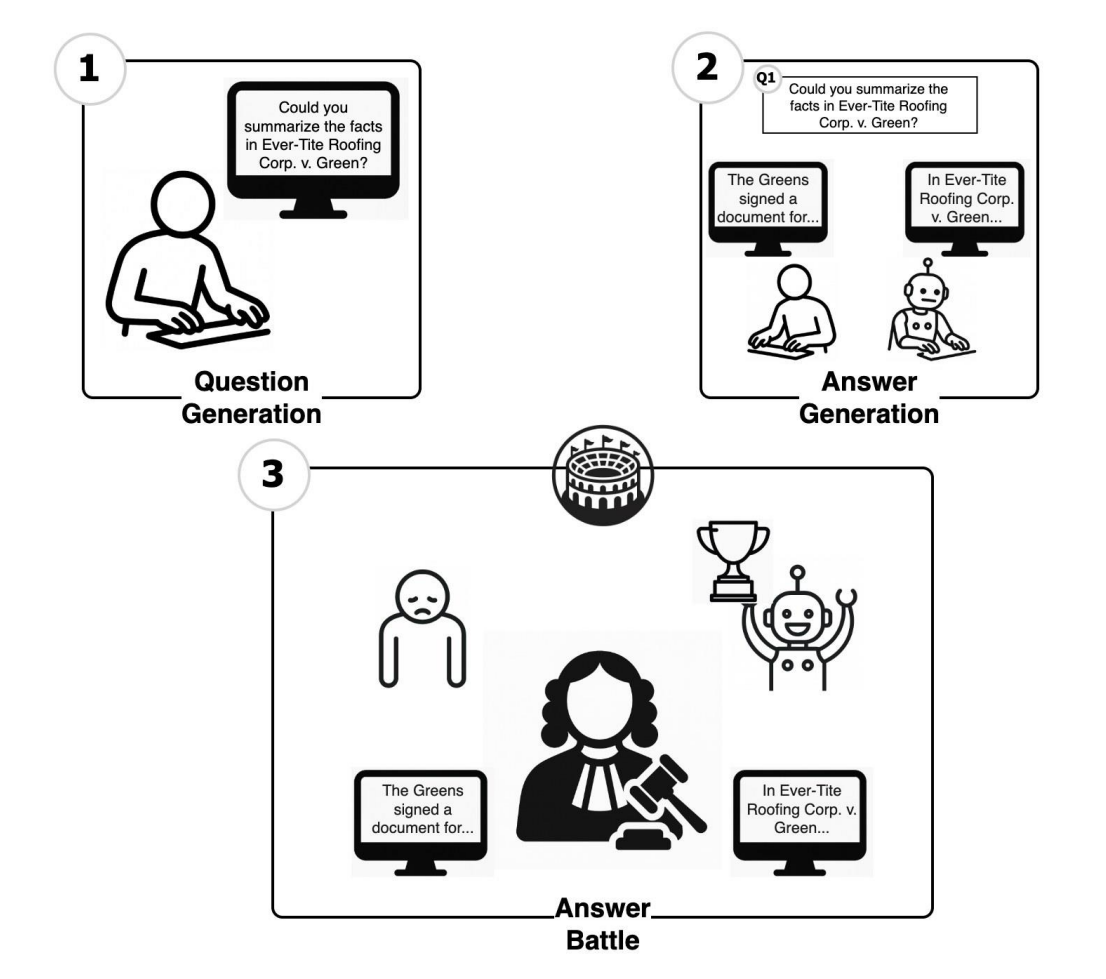


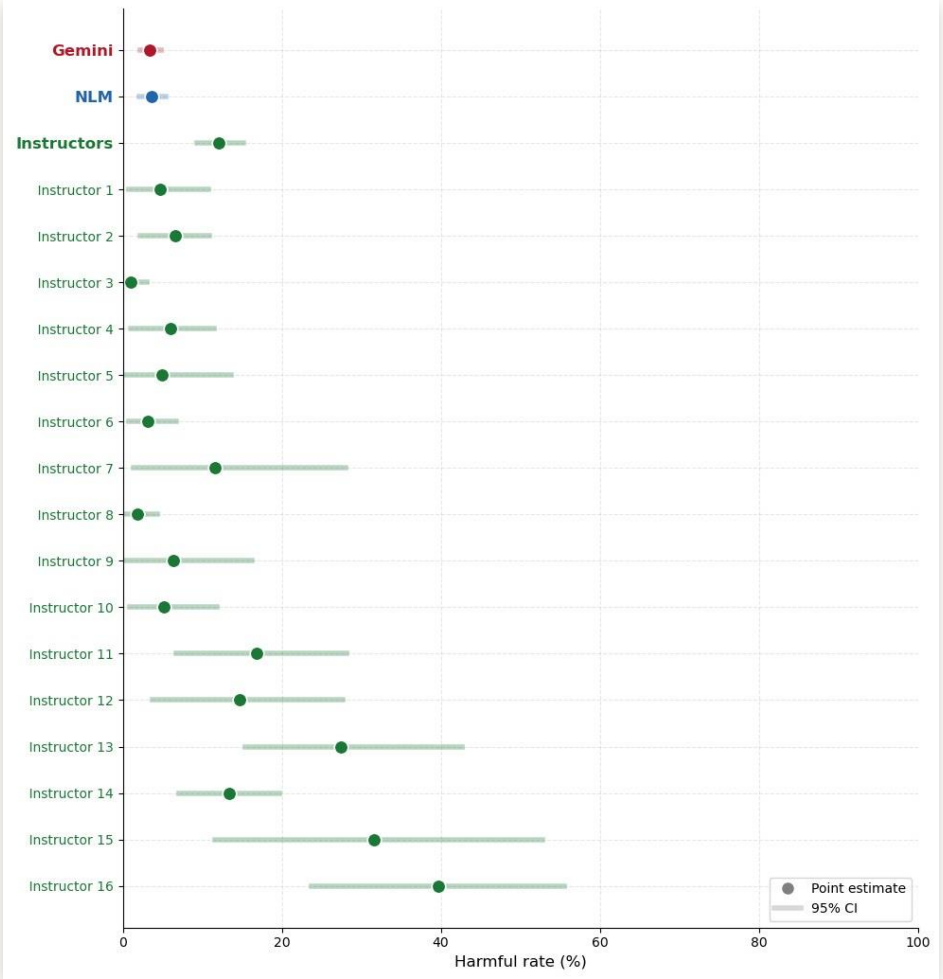
AI and the legal profession: bar exam results



Katz, Bommarito, Gao & Arredondo (2024), "GPT-4 passes the bar exam," Phil. Trans. R. Soc. A 382.

AI outperforms contracts law professors





Courtesy: Julian Nyarko.

Legal research tools still hallucinate

Input Query
Are the deadlines established by the bankruptcy rules for objecting to discharge jurisdictional?

Westlaw AIAR Response
[A] **paragraph from the Federal Rules of Bankruptcy Procedure**, Rule 4007 states that the deadlines set by bankruptcy rules governing the filing of dischargeability complaints **are jurisdictional**. [...] ✗

Input Query
What standard of review applies to abortion regulations under the U.S. Constitution?

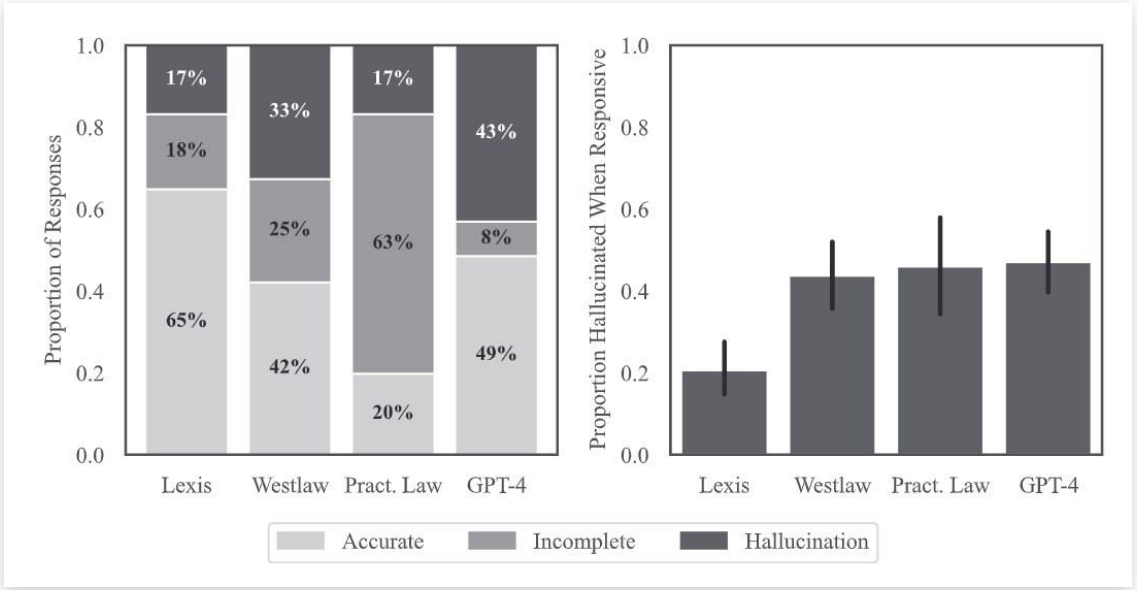
Lexis+ AI Response
Currently, the "undue burden" test, as established in *Planned Parenthood of Southeastern Pa. v. Casey*, is the governing standard. Planned Parenthood of the Heartland, Inc. v. Reynolds, [...]. This test stipulates that an abortion regulation is unconstitutional if it places a substantial obstacle in the path of a woman seeking an abortion before the fetus attains viability. [...] ✗

Input Query
Why did Justice Ginsburg dissent in Obergefell?

Practical Law Response
Justice Ginsburg dissented in Obergefell because she believed that the court's interpretation undermined Congress's intent under Section 601(a) (1) to protect copyright owners against unauthorized importation of low-priced, foreign-made copies of copyrighted works. She argued that the text of the Copyright Act clearly applies to works governed by the statute. [...] ✗

Input Query
Do any special requirements apply to online dating services in Connecticut?

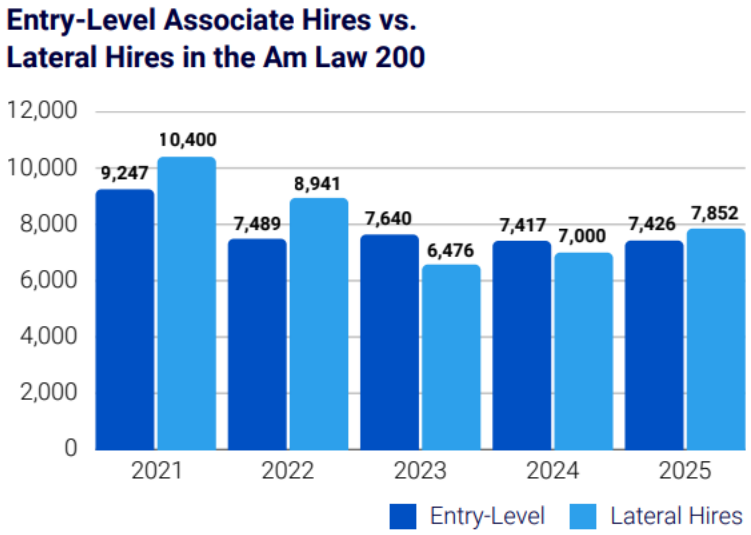
GPT-4 Response
Yes, Connecticut has specific legal requirements that apply to online dating services. Under Connecticut law, particularly **Conn. Gen. Stat. § 42-290**, online dating services that charge a fee and have members who reside in Connecticut must provide their users with specific safety awareness notifications. These notifications must include: [...] ✗



Magesh, Surani, Dahl, Suzgun, Manning & Ho (2025), "Hallucination-Free?" Journal of Empirical Legal Studies 22(2).

Legal employment: So far, so good





- The overall employment rate for the Class of 2025 was 92.8%, down 0.6 percentage point from the Class of 2024's record high of 93.4%. However, this year's figure was still the second highest employment rate on record.

- The percentage of graduates securing jobs for which bar admission is required or anticipated (i.e., attorney positions) improved by 0.3 percentage point to 84.6%, the highest level recorded since NALP introduced the current job classifications in 2001.

92.8%
Of law graduates* were employed as of March 16, 2026
*with known employment status

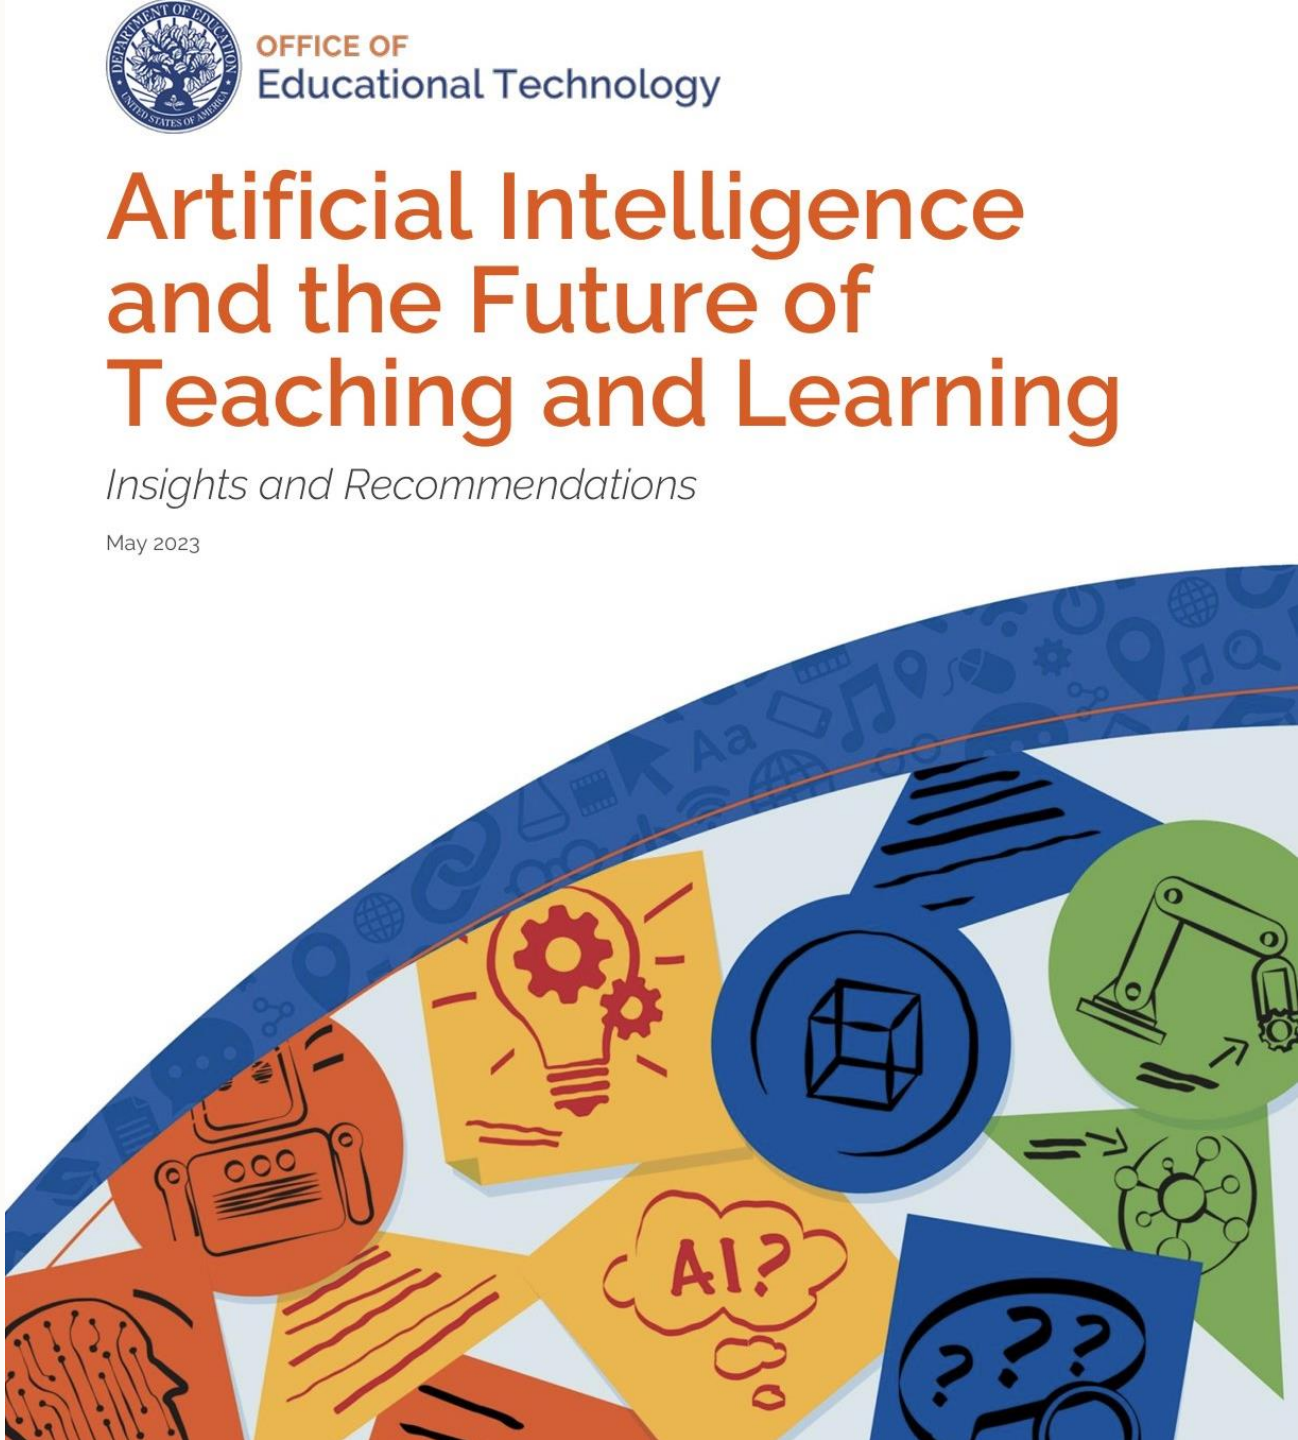
83.9%
Of law graduates* were employed in a full-time, long-term bar admission required/anticipated job
*with known employment status

98.0%
Of jobs were full-time

33.2%
Of private practice jobs were in firms of more than 500 lawyers

https://www.nalp.org/uploads/Research/Classof2025_SelectedFindings_5Aug2026.pdf

Case study #4: Teachers



AI diffusion in education

Tutor CoPilot: largest gains for less effective tutors

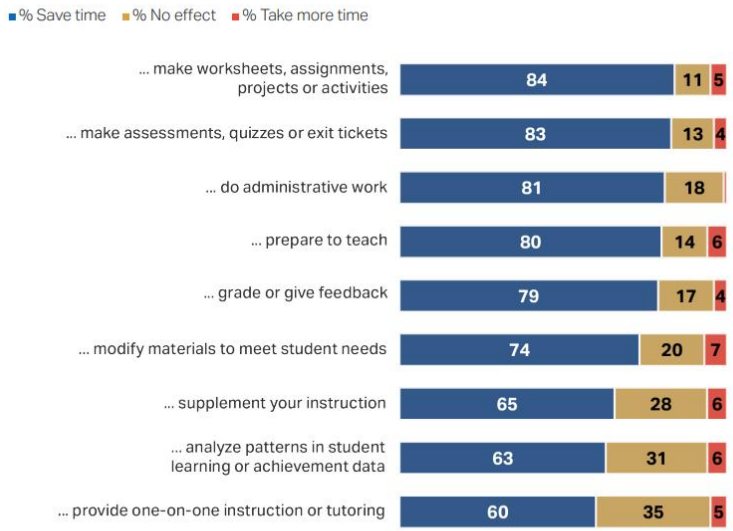
GALLUP | WALTON FAMILY FOUNDATION

TEACHING FOR TOMORROW

Unlocking Six Weeks a Year With AI

The percentage of teachers who say AI saves them time on various work tasks

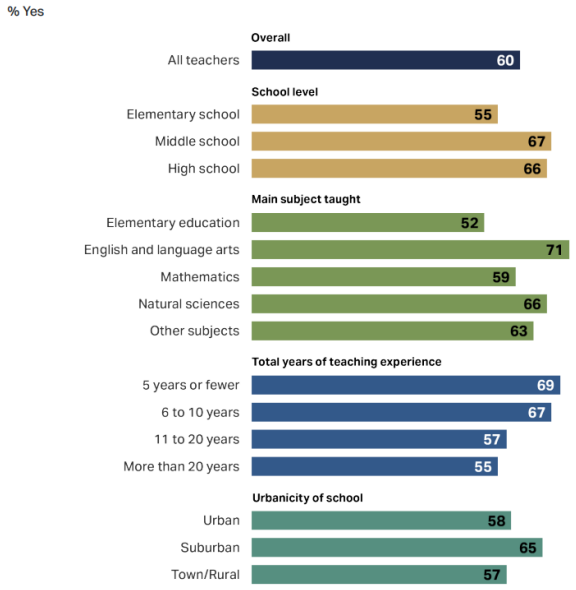
When you use AI tools to … does it save time, take more time or have no effect on the amount of time?



Note: Results for each task are reported among teachers who report using AI for that task at least monthly. Data labels less than 4% are not displayed.

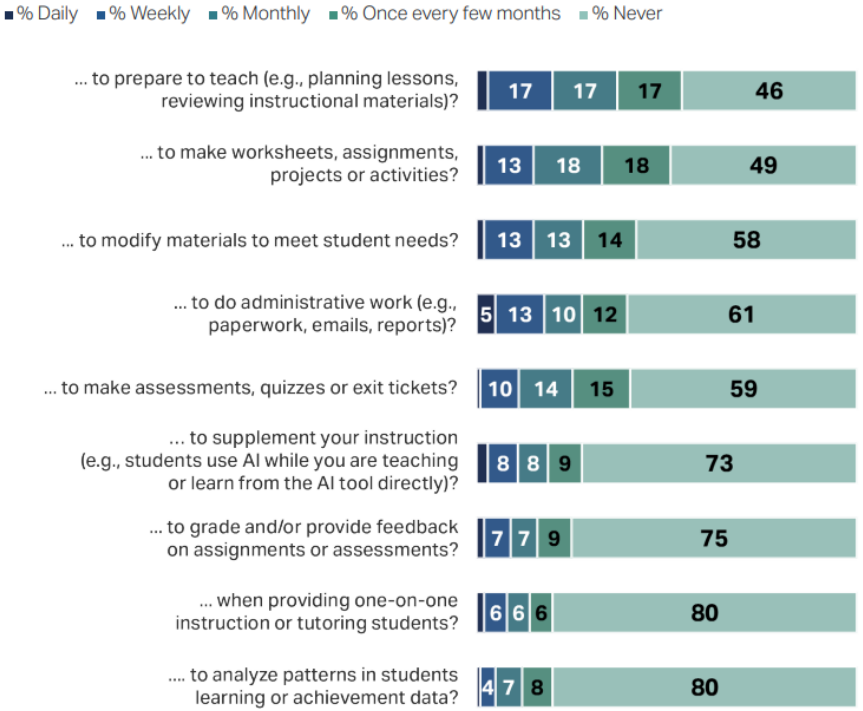
AI use by teacher and school characteristics

During the 2024-25 school year, have you used any AI tools, like chatbots, adaptive learning systems, or other interactive AI platforms for your work or teaching? (e.g., preparing to teach, making student learning materials or assessments, grading or providing feedback)



Frequency of using AI tools for various teaching tasks

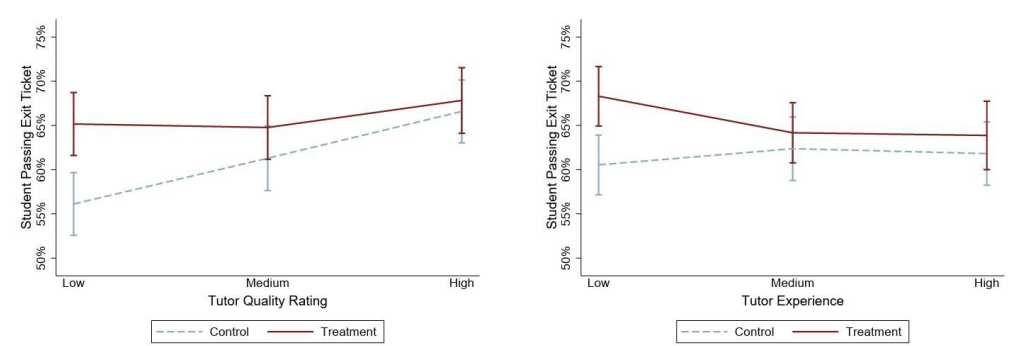
How often do you use AI tools … ?



Note: Data labels less than 4% are not displayed.

Tutor CoPilot: largest gains for less effective tutors

- First RCT of a human-AI copilot in live K–12 math tutoring
- 900 tutors; 1,800 students in Title I schools
- More prompting students to explain; less giving away answers (550,000+ messages)



(a) Tutor Quality Rating (b) Tutor Experience

Figure 2: Heterogeneity analysis by tutor initial effectiveness on student learning. (a) reports by the tutor's initial quality rating and (b) by tutor's tutoring experience.

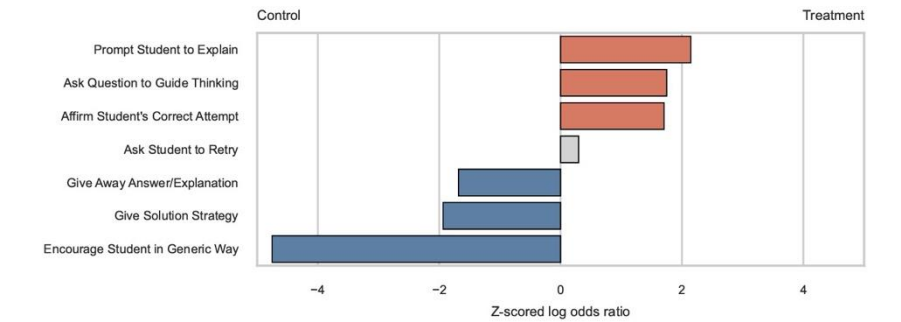
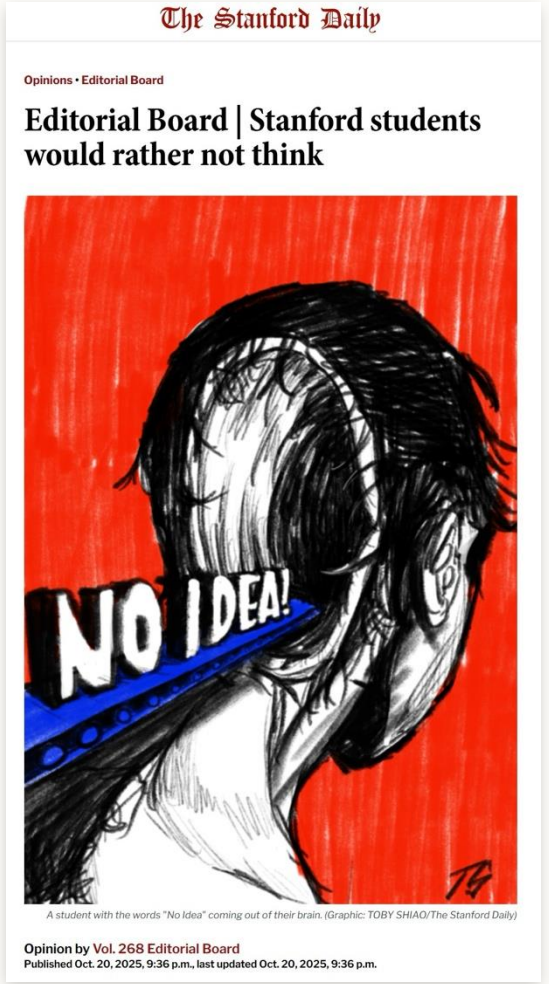


Figure 3: Strategies more likely to be used by control tutors (left) vs. treatment tutors (right). Strategies with a z-score below 1 standard deviation are shaded in gray. Control tutors tended to rely on solution-focused, passive strategies, while treatment tutors more frequently used strategies that promote deeper student engagement and comprehension.

Wang, Ribeiro, Robinson, Loeb & Demszky (2025), Tutor CoPilot, arXiv.

General conclusions from the experiments

- AI often improves novice performance on scoped tasks, but not always
- Human–AI teams can outperform either alone; performance varies by workflow and verification
- Experts may be harmed on tasks where AI guidance is unreliable
- Most classic experiments study copilots; autonomous agents could affect staffing differently


The Stanford Daily

Opinions • Editorial Board

Editorial Board | Stanford students would rather not think

A student with the words "No Idea" coming out of their brain. (Graphic: TOBY SHIAO/The Stanford Daily)

Opinion by Vol. 268 Editorial Board

Published Oct. 20, 2025, 9:36 p.m., last updated Oct. 20, 2025, 9:36 p.m.

Are there jobs AI shouldn’t do?







Just because AI can do a job doesn't mean we want it to

Table 1. Selected occupations in the top/bottom 50 ranked by repugnance score. Occupations were selected to illustrate the range of occupations in these categories. Repugnance score indicates mean with 95% confidence intervals, calculated via stratified bootstrap with 10,000 replications.

Rank	Occupation	Repugnance score
Most repugnant		
1	Clergy	5.91 [5.44, 6.31]
2	Childcare Workers	5.86 [5.31, 6.26]
4	Marriage and Family Therapists	5.64 [5.12, 6.04]
6	Administrative Law Judges, Adjudicators, and Hearing Officers	5.62 [4.82, 6.10]
9	Athletes and Sports Competitors	5.52 [4.82, 6.01]
11	Craft Artists	5.50 [5.00, 5.95]
17	Police and Sheriff's Patrol Officers	5.40 [4.89, 5.81]
27	Funeral Attendants	5.28 [4.54, 5.89]
48	Barbers	5.03 [4.41, 5.52]
50	Actors	5.01 [4.43, 5.57]
Least repugnant		
904	Derrick Operators, Oil and Gas	2.72 [2.26, 3.29]
911	Cashiers	2.67 [2.10, 3.42]
913	Janitors and Cleaners, Except Maids and Housekeeping Cleaners	2.67 [2.15, 3.32]
924	Data Entry Keyers	2.55 [2.05, 3.25]
926	Segmental Pavers	2.55 [2.10, 3.05]
927	Biostatisticians	2.54 [2.14, 3.04]
930	Switchboard Operators, Including Answering Service	2.52 [2.02, 3.27]
936	Transportation Planners	2.38 [1.92, 2.93]
938	Search Marketing Strategists	2.31 [1.91, 2.89]
940	File Clerks	2.17 [1.69, 2.84]

Friis & Riley (2025), "Performance or Principle: Resistance to AI in the U.S. Labor Market," Harvard Business School.

Americans would accept AI in ~58% of occupations; ~12% are off-limits



- Accept if AI performs well and cheaply: about 58% of occupations
- Categorically off-limits: about 12%
- Caregiving, therapy and spiritual roles cluster in “high feasibility, high repugnance”

Friis & Riley (2025), Harvard Business School.

Care, education and spiritual leadership: high feasibility, high repugnance

Table S12: Examples of Occupations by Technical Suitability and Moral Repugnance Classification

Occupation	Technical Suitability Score	Repugnance Score
Low Technical Suitability, Low Moral Repugnance		
Paper Goods Machine Setters, Operators, and Tenders	0.00	2.56
Highway Maintenance Workers	0.05	2.60
Shoe Machine Operators and Tenders	0.05	2.48
Segmental Pavers	0.08	2.55
Extruding and Forming Machine Setters, Operators, and Tenders, Synthetic...	0.05	2.62
Low Technical Suitability, High Moral Repugnance		
Ambulance Drivers and Attendants, Except Emergency Medical Technicians	0.09	5.23
Embalmers	0.04	4.74
Hairdressers, Hairstylists, and Cosmetologists	0.09	4.98
Home Health Aides	0.07	4.87
Athletes and Sports Competitors	0.00	5.52
High Technical Suitability, Low Moral Repugnance		
Radio Frequency Identification Device Specialists	0.90	2.46
Search Marketing Strategists	1.00	2.31
Geographic Information Systems Technologists and Technicians	1.00	2.45
Transportation Planners	0.95	2.38
Cartographers and Photogrammetrists	1.00	2.47
High Technical Suitability, High Moral Repugnance		
Education and Childcare Administrators, Preschool and Daycare	0.88	5.69
Psychology Teachers, Postsecondary	0.76	5.62
Administrative Law Judges, Adjudicators, and Hearing Officers	0.86	5.62
School Psychologists	0.79	5.45
Psychiatrists	0.83	5.36

Friis & Riley (2025), Table S12.

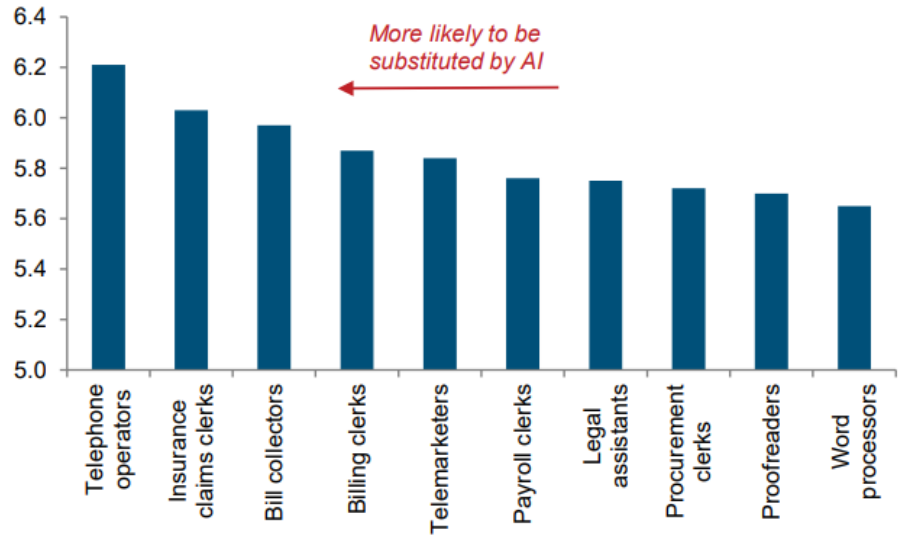
AI AUGMENTATION



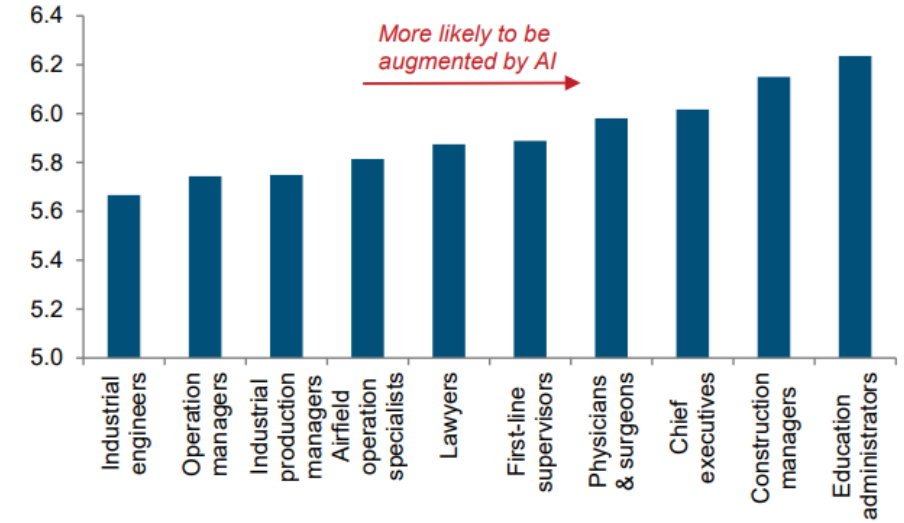
AI AUTOMATION



Occupations most exposed to AI substitution effect, displacement score

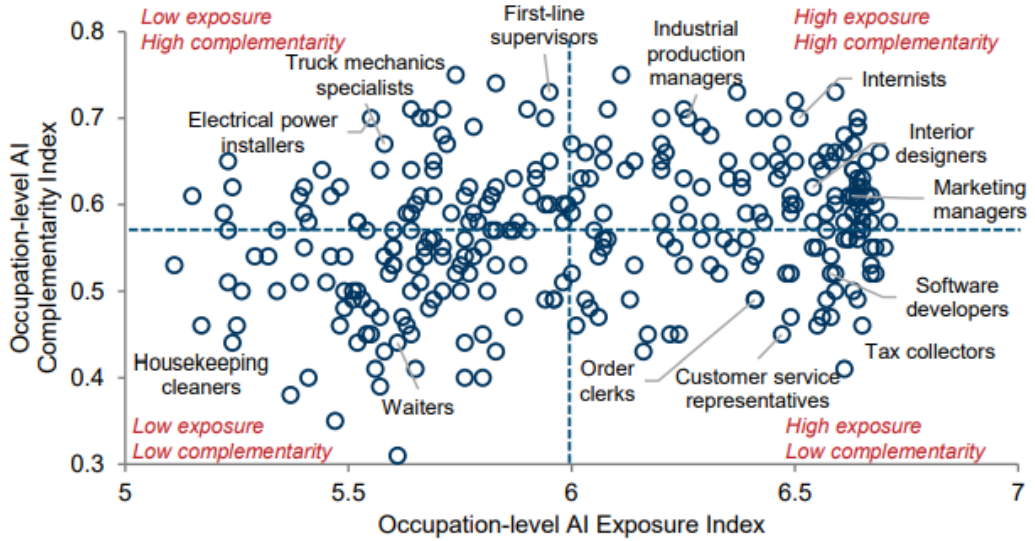


Occupations most exposed to AI augmentation effect, augmentation score



Source: Goldman Sachs GIR.

Whether AI is more likely to augment or substitute for labor depends on its level of exposure and complementarity profile...



Source: Goldman Sachs GIR.

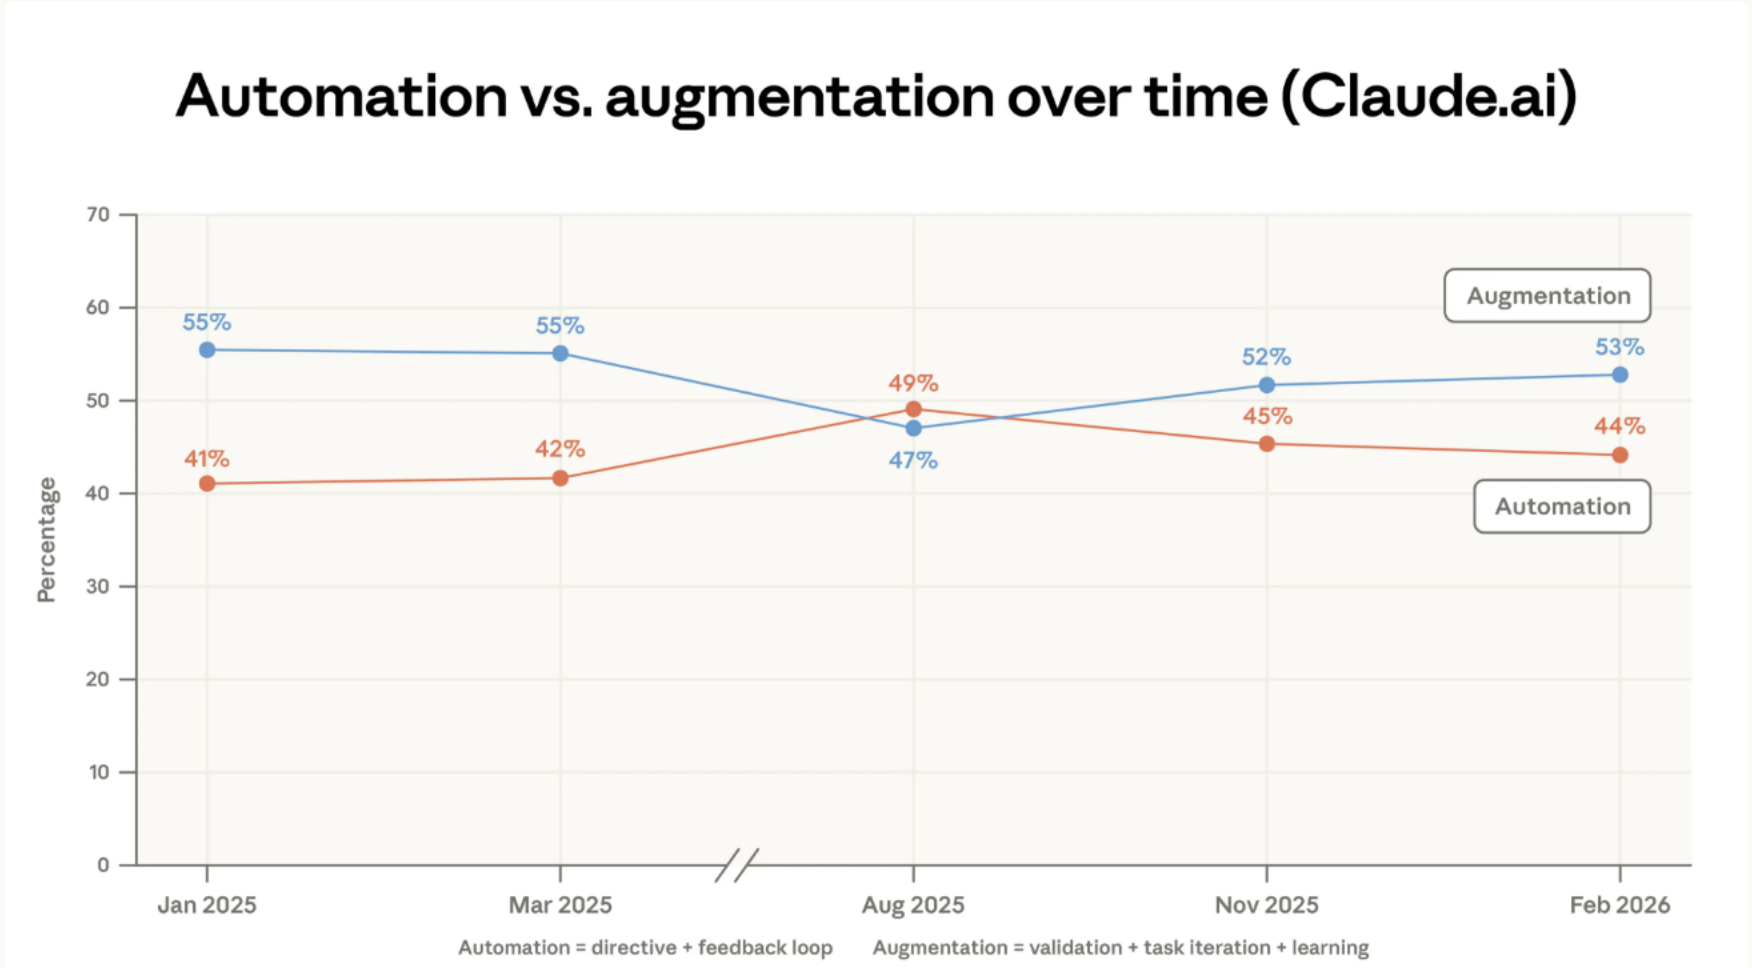


FIGURE 1.3 **Collaboration mode share, Claude.ai.** Collaboration mode frequencies across Anthropic Economic Index Reports in Claude.ai.

Collaboration v. delegation: early Claude data

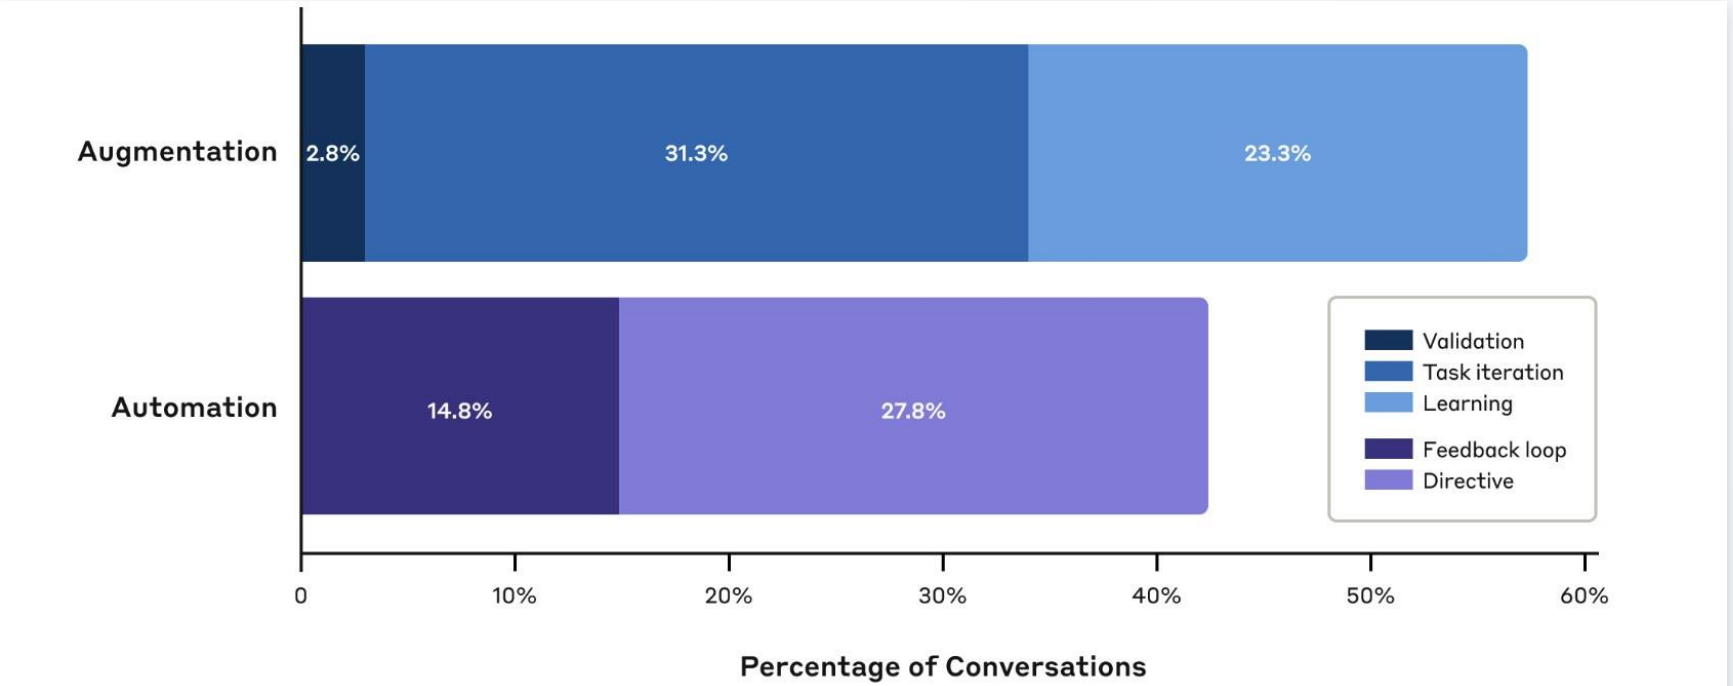


Figure 7: **Distribution of automative behaviors (43%) where users delegate tasks to AI, and augmentative behaviors (57%) where users actively collaborate with AI.** Patterns are categorized into five modes of engagement; automative modes include Directive and Feedback Loop, while augmentative modes are comprised of Task Iteration, Learning, and Validation.

Source: Handa et al., "Which Economic Tasks are Performed with AI?" arXiv:2503.04761 (Feb. 2025).

VII. Policy options

Labor-market adjustment, bargaining power and skills policy



Policy options

Approach	Examples	Reference
Design AI to augment labor	Support human-in-loop deployment; evaluate impacts on entry roles and careers	Acemoglu & Restrepo
Build skills and career ladders	Apprenticeships, junior training, portability and mobility support	National Academies; MGI (2026)
Share productivity gains	Wage bargaining, worker voice, profit sharing, hours and benefits	Labor economics literature
Protect workers in transition	Wage insurance, training finance, job matching, targeted adjustment aid	Goldman Sachs (2026)
UBI and Redistribution	Tax and transfer policy, social insurance, potential UBI if displacement accelerates	Korinek & Stiglitz; others

Open questions

1. Do the entry-level effects persist, spread to older workers, or reverse?
2. Will agentic deployment show up in outcome data within 12–24 months?
3. Will AI create enough new tasks to offset automation?
4. Where will tomorrow’s experts come from if the first rung disappears?
5. Who absorbs the cohorts who are not hired?

Thank You

npersily@stanford.edu · @Persily



Labor Economists: Tasks, Automation, and Augmentation

The "Task Approach" to Labor Markets: An Overview

David H. Autor
MIT, NBER, IZA and IAB

Characterizing tasks added and removed: Routine, Manual and Abstract

	Abstract	Routine	Manual
Tasks in use, 1977	30.4%	52.4%	17.2%
Tasks in use, 2018	52.0%	33.0%	15.0%
Tasks removed, '77-'18	22.1%	**64.5%**	13.4%
Tasks added, '77-'18	**75.6%**	17.9%	6.5%

When Autor classified tasks in his dataset as abstract, routine, or manual, he found that 64.5% of routine tasks were removed and 75.6% of abstract tasks were added between 1977 and 2018.

Task approach to the labor market effects of any new technology

 **Automate**	 **Augment**
Machine does the task *for* us	Machine does the task *with* us
Capital substitutes for labor	Capital complements or supplements labor








ATM











Is AI an Automating or Augmenting Technology?

Better Question:

How Should We Shape AI in Automating or Augmenting Directions?

TECH· AI

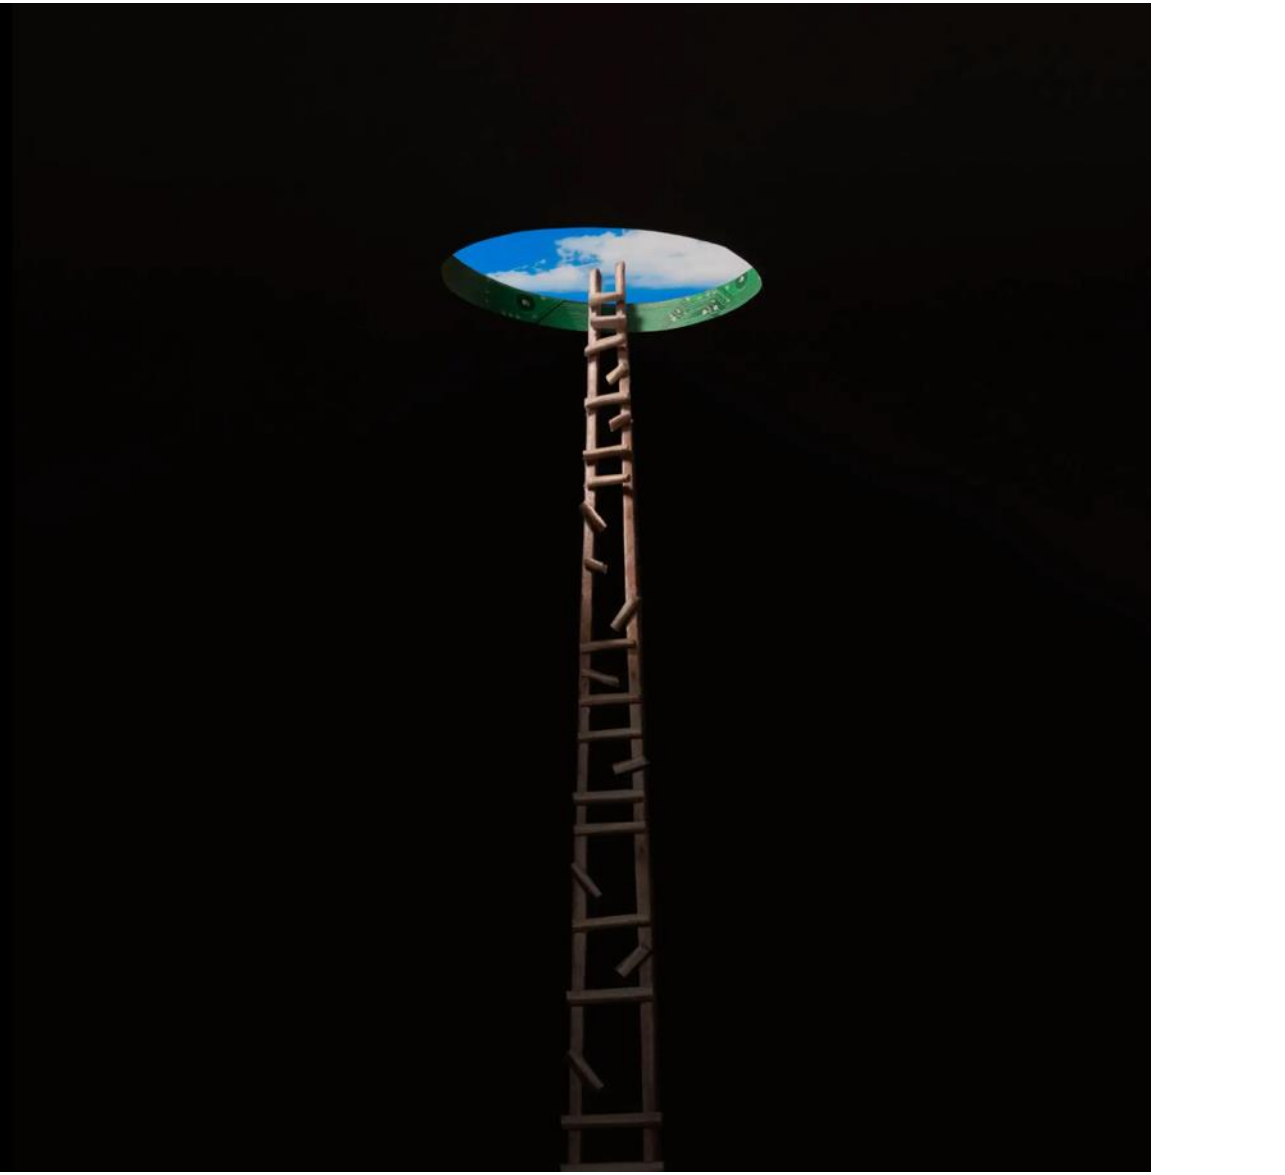
DeepMind co-founder Mustafa Suleyman warns AI is a ‘fundamentally labor replacing’ tool over the long term



OPINION

GUEST ESSAY

Silicon Valley Is Bracing for a Permanent Underclass



Most people I know in the A.I. industry think the median person is screwed, and they have no idea what to do about it. I live in San Francisco, among the young researchers earning million-dollar salaries and the start-up founders competing to build the next unicorn. While Silicon Valley has long warned about the risk of rogue A.I., it has recently woken up to a more mundane nightmare: one in which many ordinary people lose their economic leverage as their jobs are automated away.

Whether you talk with engineers, venture capitalists, founders or managers, or with doomers, accelerationists, lefties or libertarians, the so-called San Francisco consensus on the impact of A.I. for workers is bleak. Many are convinced that advanced A.I. will soon surpass human capabilities. This would produce tremendous growth and scientific achievement, but it would also displace millions of jobs as fewer humans are needed to make the economy run. The technology will depress economic mobility and exacerbate inequality, while ferrying power and wealth to the A.I. companies and the existing owners of capital.



Anthropic Economic Index

Understanding AI's effects on the economy

Last updated: Sep 15, 2025

Directive automation is accelerating

As in previous reports, we also track not just *what* people use Claude for but *how* they collaborate with or delegate to Claude on Claude.ai.

At a high level, we distinguish between *automation* and *augmentation* modes of using Claude:

Automation encompasses interaction patterns focused on task completion:

- **Directive:** Users give Claude a task and it completes it with minimal back-and-forth
- **Feedback Loops:** Users automate tasks and provide feedback to Claude as needed

Augmentation focuses on collaborative interaction patterns:

- **Learning:** Users ask Claude for information or explanations about various topics
- **Task Iteration:** Users iterate on tasks collaboratively with Claude
- **Validation:** Users ask Claude for feedback on their work

The share of *directive* conversations sampled from Claude.ai conversations jumped from 27% in V1 in late 2024 to 39% in V3. This

How close is AI to automating remote work?

Remote Labor Index (RLI)

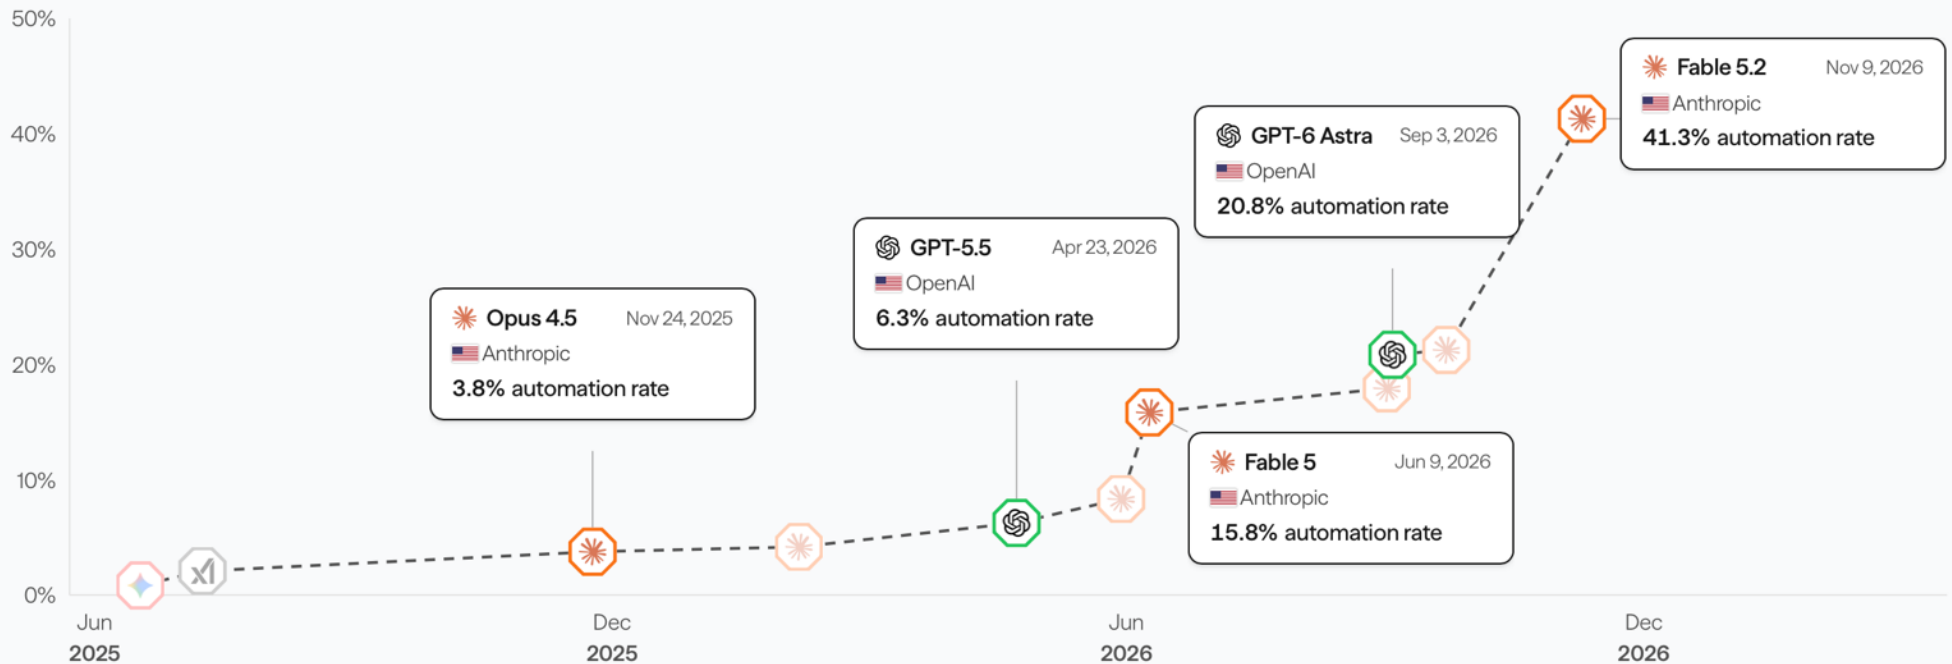
41.3%

full automation of remote projects

AUTOMATION PACE

×2 every ~3.3 months

over the last 12 months





Rival Visions of Computing

John McCarthy and

Douglas Engelbart



Letter from our Director

Welcome to the Stanford Artificial Intelligence Lab

The Stanford Artificial Intelligence Lab (SAIL) was founded by Prof. John McCarthy, one of the founding fathers of the field of AI. While the discipline of AI has transformed in many fundamental ways since its inception in the 1950s, SAIL remains a proud leading intellectual hub for scientists and engineers, an education Mecca for students, and a center of excellence for cutting edge research work.

Automation

John McCarthy, “What is AI?” (1997)

of the Turing test and the various partial Turing tests that have been implemented, i.e. with restrictions on the observer's knowledge of AI and the subject matter of questioning. It turns out that some people are easily led into believing that a rather dumb program is intelligent.

Q. Does AI aim at human-level intelligence?

A. Yes. The ultimate effort is to make computer programs that can solve problems and achieve goals in the world as well as humans. However, many people involved in particular research areas are much less ambitious.

Q. How far is AI from reaching human-level intelligence? When will it happen?

A. A few people think that human-level intelligence can be achieved by writing large numbers of programs of the kind people are now writing and assembling vast knowledge bases of facts in the languages now used for expressing knowledge.

However, most AI researchers believe that new fundamental ideas are required, and therefore it cannot be predicted when human-level intelligence will be achieved.

Q. Are computers the right kind of machine to be made intelligent?

NEWS

AI Benchmarks Hit Saturation

DATE APRIL 03, 2023

TOPICS MACHINE LEARNING

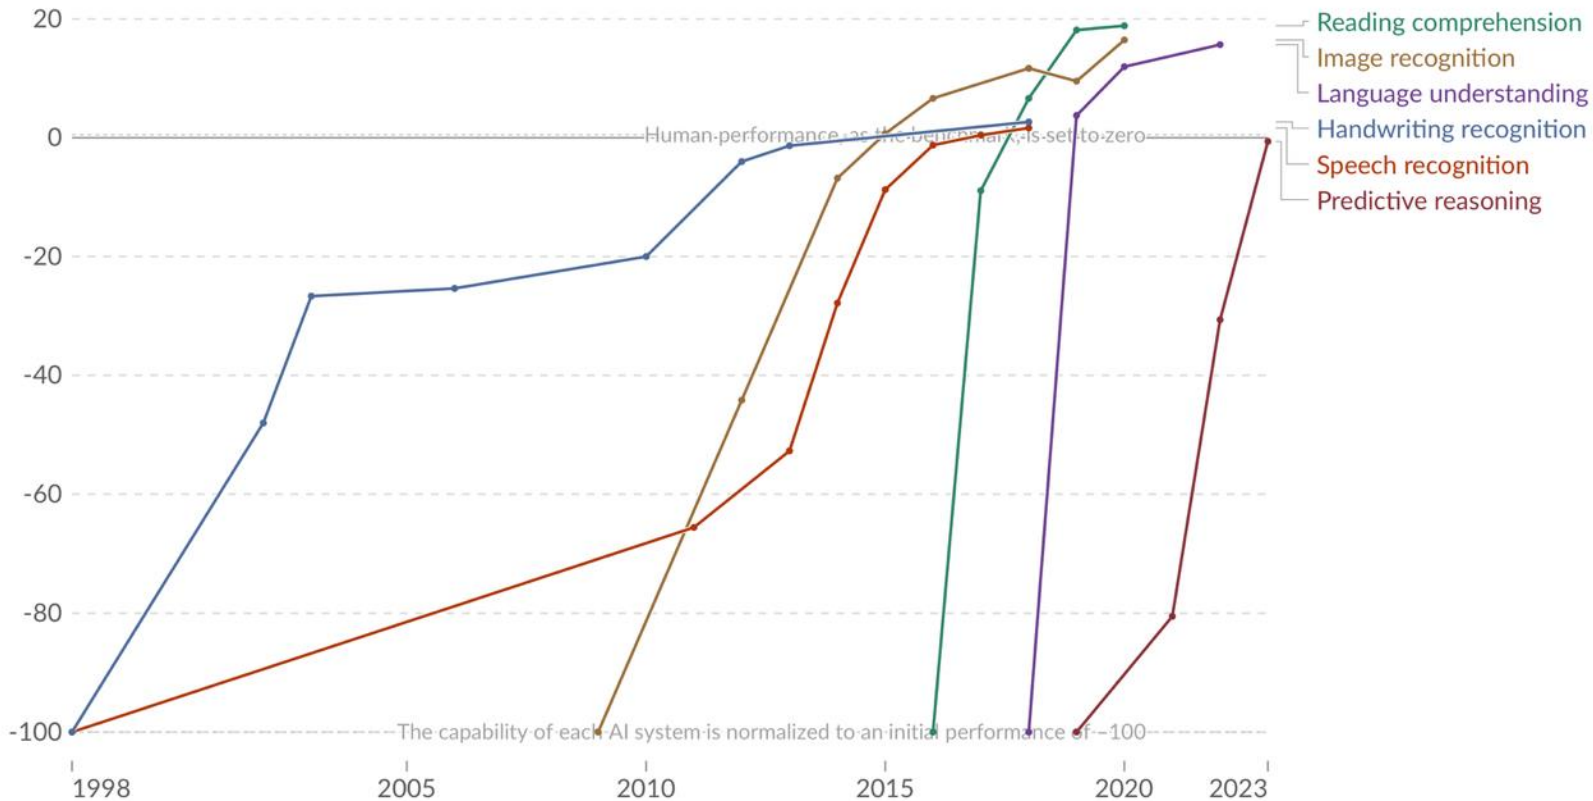


AI continues to surpass human performance; it's time to reevaluate our tests.

Test scores of AI systems on various capabilities relative to human performance

Our World in Data

Within each domain, the initial performance of the AI is set to −100. Human performance is used as a baseline, set to zero. When the AI's performance crosses the zero line, it scored more points than humans.

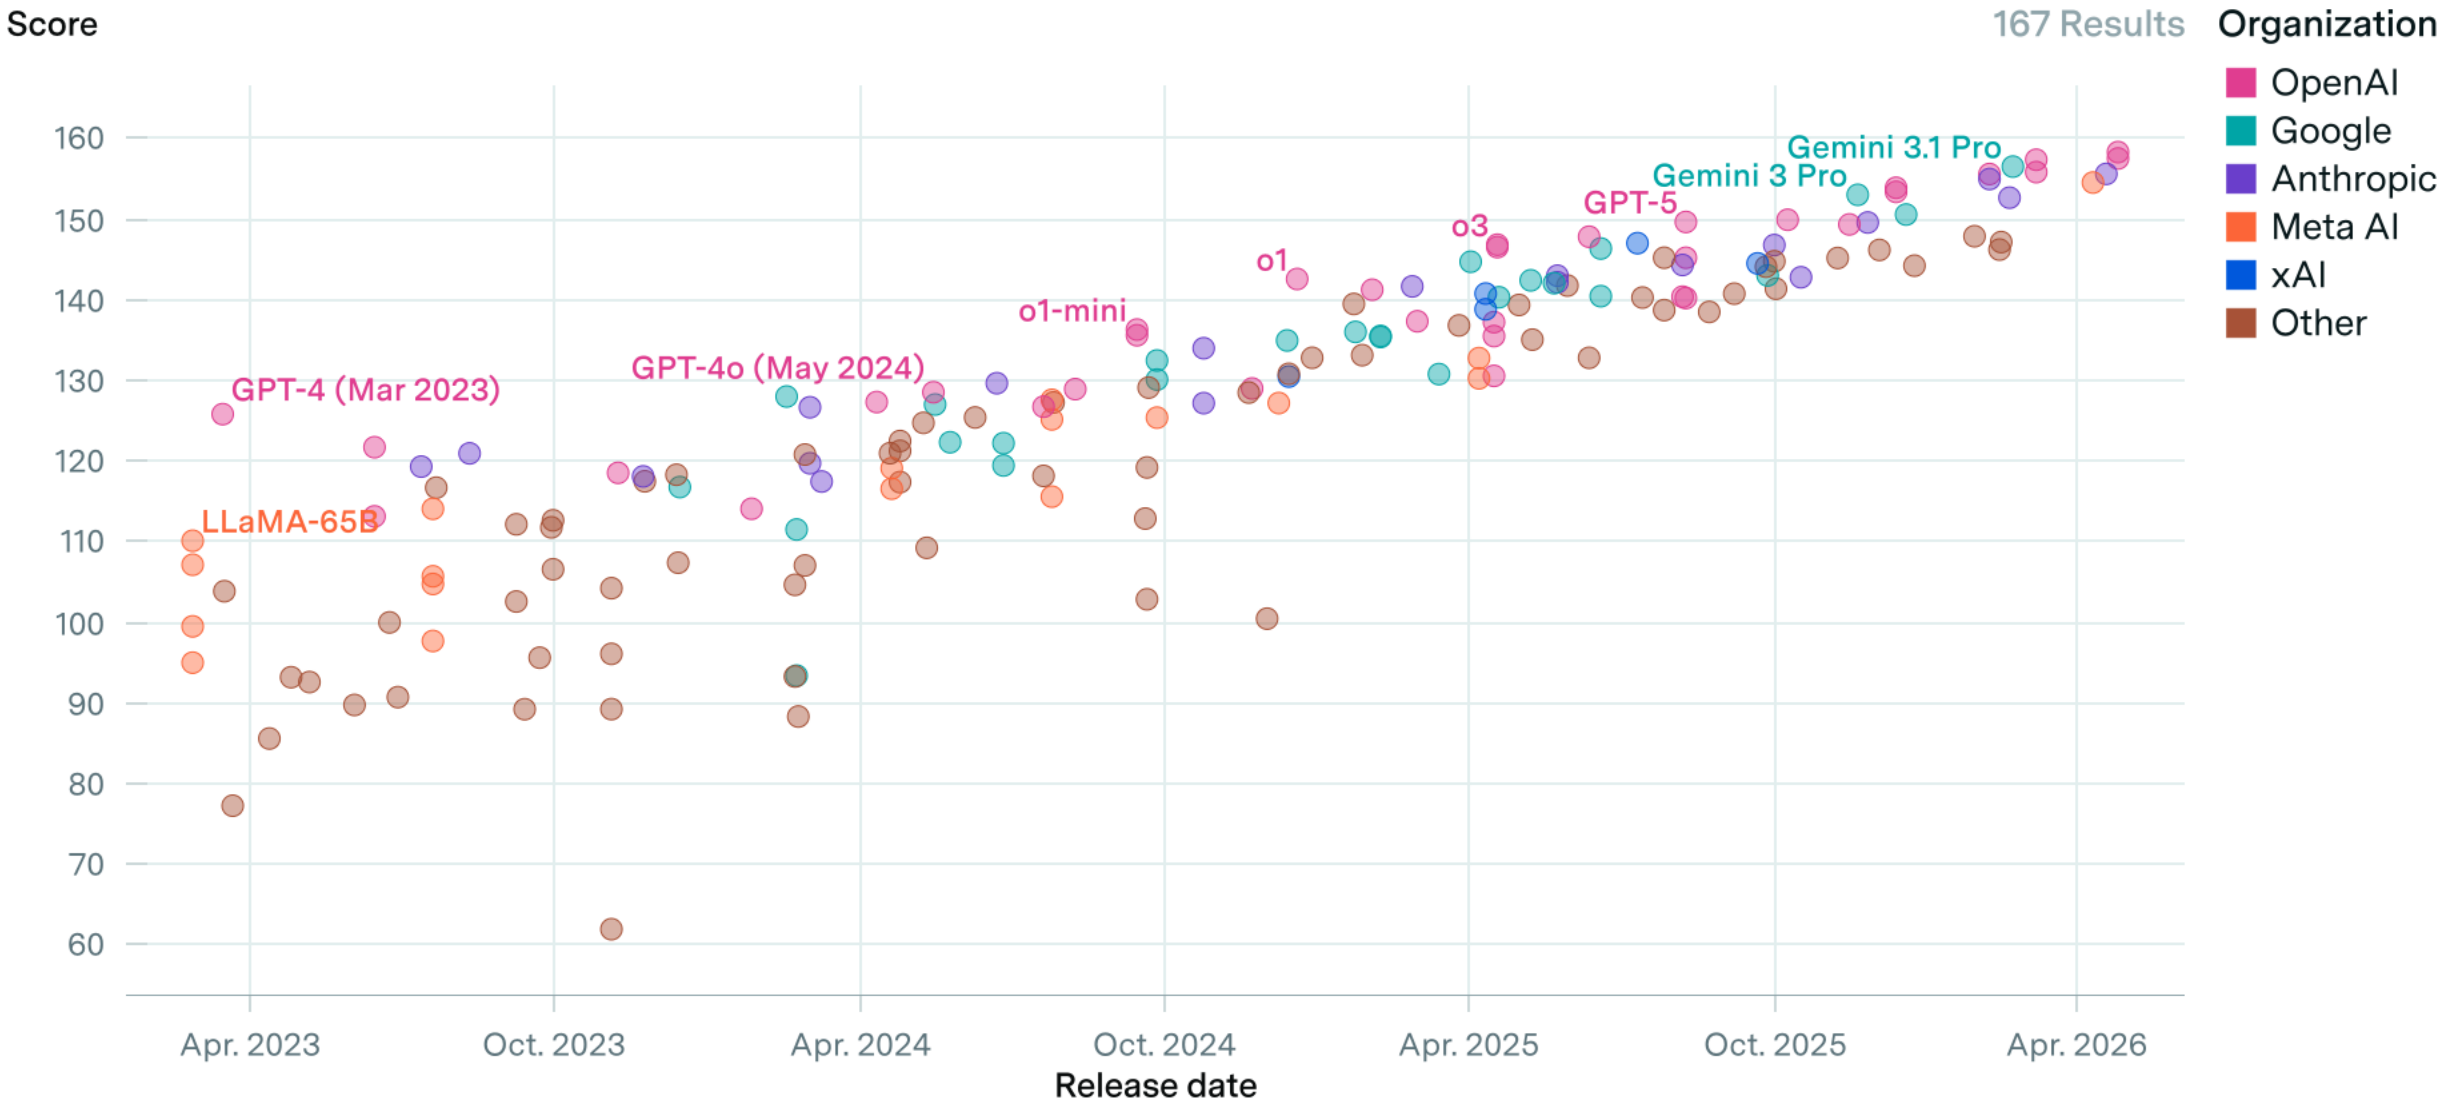


Data source: Kiela et al. (2023)

OurWorldinData.org/artificial-intelligence | CC BY

Note: For each capability, the first year always shows a baseline of −100, even if better performance was recorded later that year.

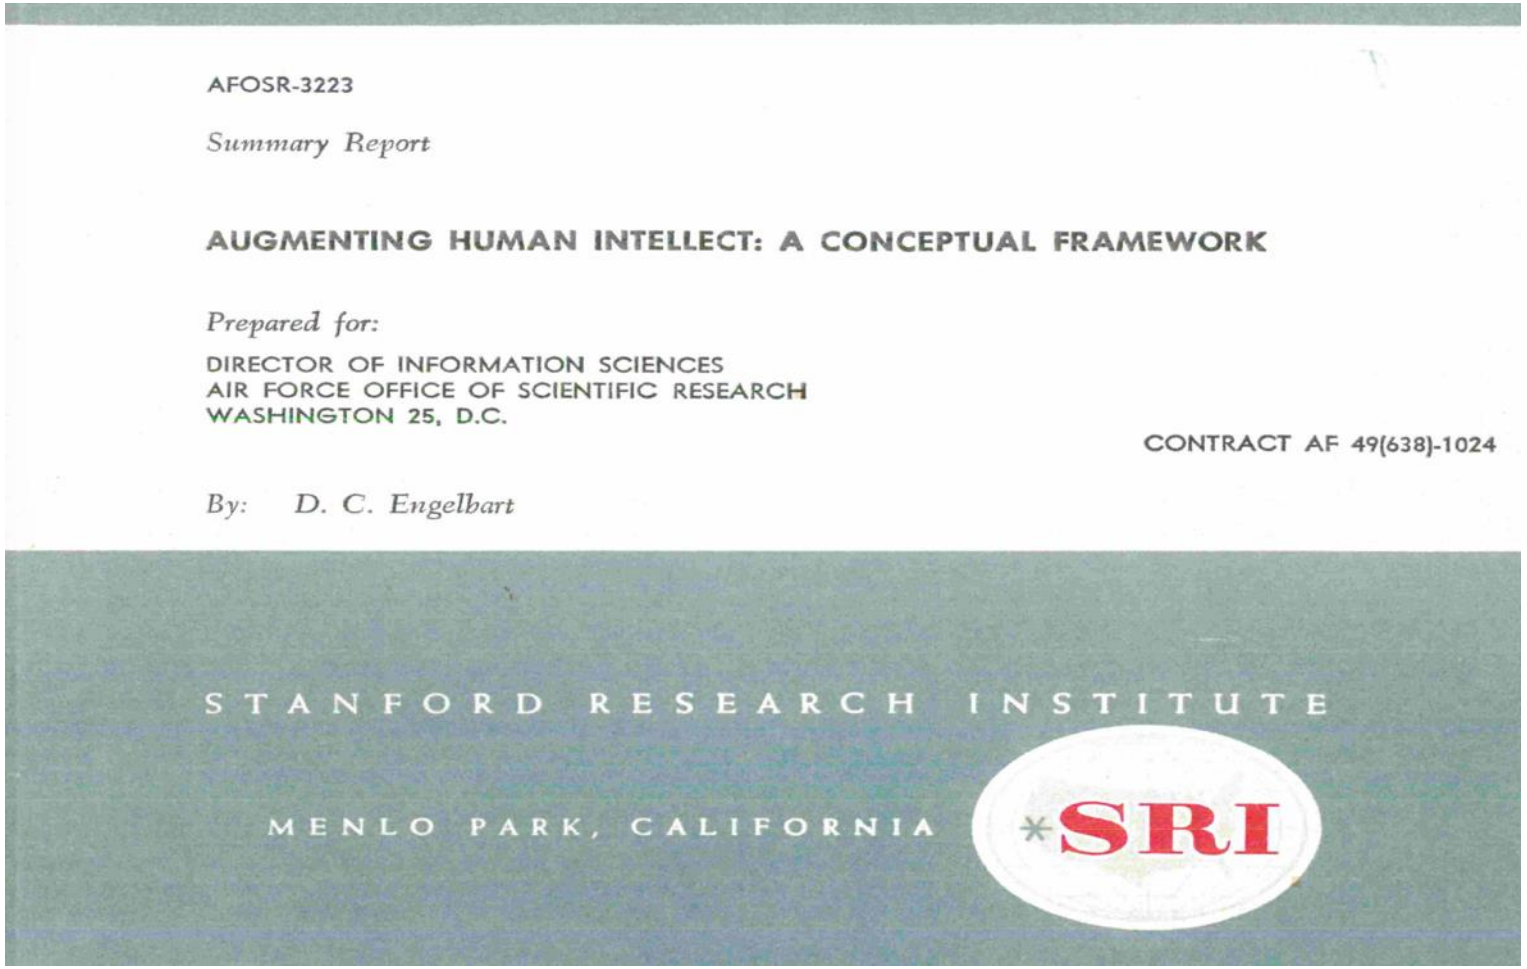
Epoch Capabilities Index (ECI)



EPOCH AI | CC-BY

Augmentation

Douglas Engelbart, 1962



AFOSR-3223

Summary Report

AUGMENTING HUMAN INTELLECT: A CONCEPTUAL FRAMEWORK

Prepared for:

DIRECTOR OF INFORMATION SCIENCES
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
WASHINGTON 25, D.C.

CONTRACT AF 49(638)-1024

By: D. C. Engelbart

STANFORD RESEARCH INSTITUTE

MENLO PARK, CALIFORNIA

*SRI

I. INTRODUCTION 1

A. GENERAL 1a

By "augmenting human intellect" we mean increasing the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems. Increased capability in this respect is taken to mean a mixture of the following: more-rapid comprehension, better comprehension, the possibility of gaining a useful degree of

25 years of the free encyclopedia

Search Wikipedia

Search

The Mother of All Demos

17 languages

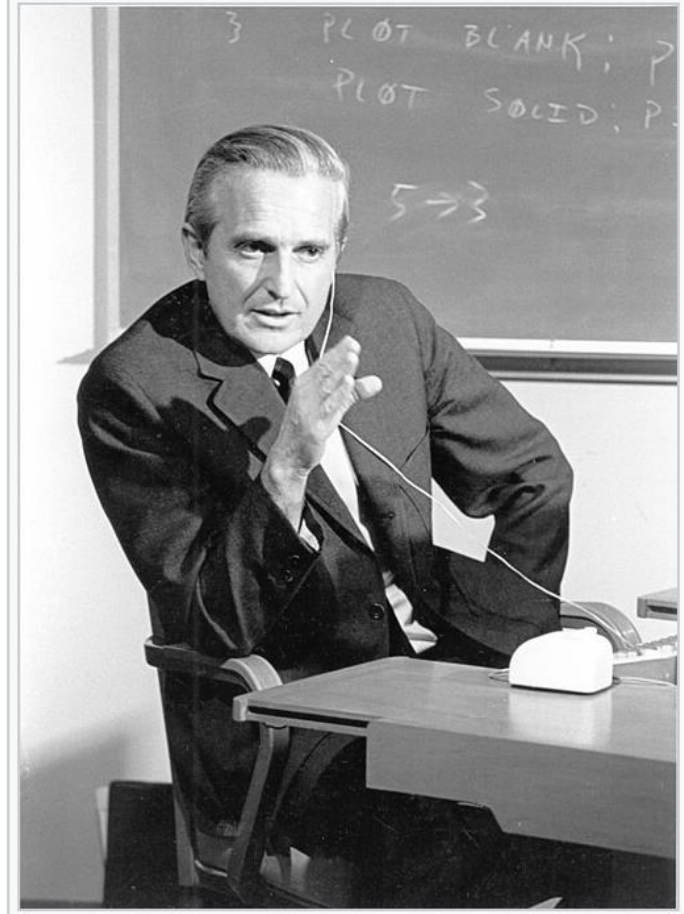
nts hide

ound
stration
ce
nces
s
ons
raphy
al links

From Wikipedia, the free encyclopedia

"**The Mother of All Demos**" was a landmark computer demonstration, named retroactively, of developments by Stanford Research Institute's Augmentation Research Center. It was presented at the Association for Computing Machinery / Institute of Electrical and Electronics Engineers (ACM/IEEE)—Computer Society's Fall Joint Computer Conference in San Francisco, by Douglas Engelbart, on December 9, 1968.[1]

The 90-minute live demonstration featured the introduction of a complete computer hardware and software system called the oN-Line System or, more commonly, NLS, which demonstrated for the first time many of the fundamental elements of modern personal computing, including windows, hypertext, graphics, efficient navigation and command input, video conferencing, the computer mouse, word processing, dynamic file linking, revision control, and a collaborative real-time editor.



Engelbart practicing for the demo

A Bicycle for the Mind



"What a computer is to me is the most remarkable tool that we have ever come up with, and it's the equivalent of a bicycle for our minds."

– Steve Jobs



The Turing Trap:

The Promise & Peril of Human-Like Artificial Intelligence

Erik Brynjolfsson

Stanford Digital Economy Lab

http://brynjolfsson.com

Forthcoming in *Dædalus*, Spring 2022

IN 1950, ALAN Turing famously created what's now known as the Turing Test, a way of deciding whether a computer is intelligent. If the computer could converse so fluently that it passed as a human? Presto: That's artificial intelligence.

Turing's test became the north star for generations of AI pioneers. For decades, they've labored mightily to mimic basic human skills, with wild success: We've now got AI that can hold conversations, draw pictures, or play expert rounds of chess, Go, and fast-paced video games.

But now some AI thinkers wonder whether we've succeeded a little *too* well—at the wrong task. Mimicking human abilities, they believe, has led to direct economic competition between people and machines. Maybe Turing led us astray and copying humans was the wrong goal.

That's the argument of Erik Brynjolfsson, who is the director of Stanford's Digital Economy Lab and has long written about AI's effect on labor. In a recent paper, he argued that Turing-chasing scientists have driven up wage inequalities.

Centaur Evaluations

ECONOMICS OF TRANSFORMATIVE AI

NEW MEASURES OF THE ECONOMY



AI should not be an imitation game. It's time to evaluate **human-AI collaboration**.

Also by Roy Bahat



Learning about the future of work

Follow publication

Can we build better AI? Consider the case of the loom vs. the crane.



Roy Bahat Follow 6 min read · Sep 27, 2023

451 6

Not all artificial intelligence is created equal. Founders of companies who create AI will need to build their services differently depending on whether it expands human abilities (“cranes”) instead of replacing them (“looms”).

Happy Tuesday, Rob

How can I help you today?

Opus 4.5

Write | Learn | Code | Life stuff | Claude's choice

Set up Claude for your classes, career, and research



Study project
Learn and master any subject



Career project
Find the next step for your career



Research project
Analyze and organize research

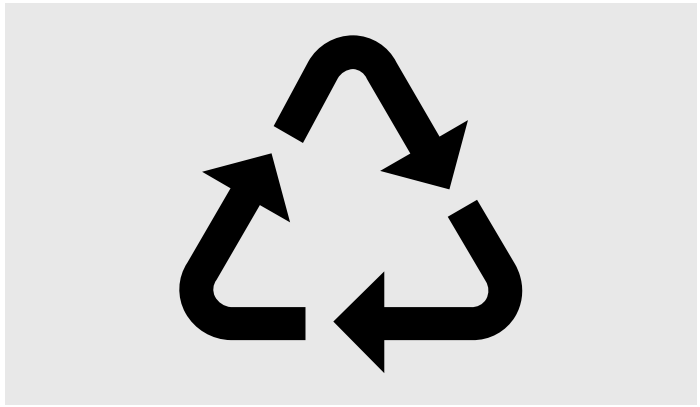
Is Augmentation > Automation?



Towards a normative framework: automate and augment

Reasons to **Automate**	Reasons to **Augment**
Save time	Improve performance, amplify human capabilities
Increase quality	Retain human accountability
Substitute machine labor for drudgery or dangerous tasks	Create intrinsic satisfaction from work; integrity

A normative framework: the three As of AI adoption



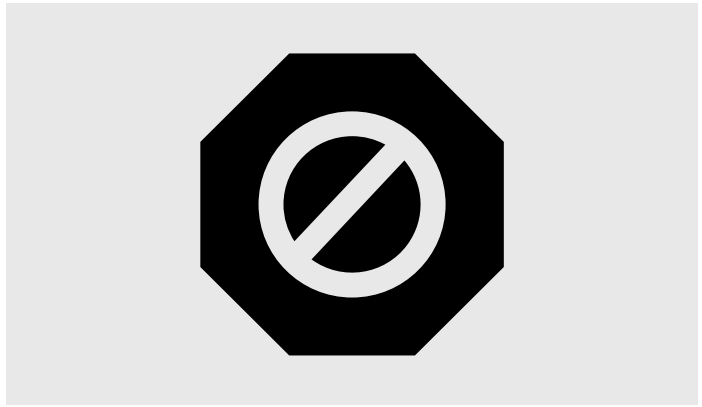
Automate

- **Improve** quality, time efficiency, or scope of coverage
- **Reduce** dangerous or drudgery tasks



Augment

- **Improve** quality
- **Retain** human accountability



Avoid

- **Offloading** cognition
- **Creating** brainrot
- **Diminishing** expertise, integrity, authenticity, or purpose