Who uses AI, and for what?

From potential exposure to observed use

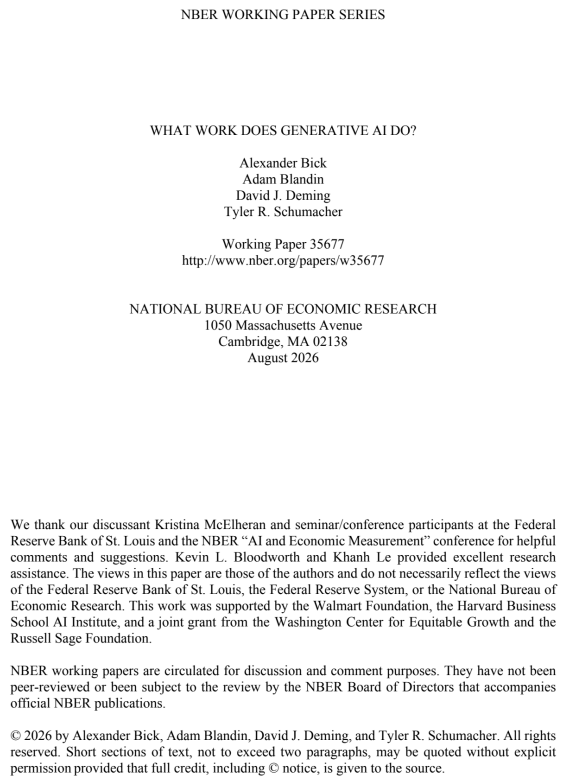
Sergio Ocampo
ECON 3399 · AI and Robots at Work

Three papers answering different measurement questions

Paper	Question
What Work Does Generative AI Do?	Which workers and tasks use GenAI?
The Rapid Adoption of Generative AI	How widespread and intensive is adoption?
How People Use ChatGPT	What happens within consumer ChatGPT conversations?

> Adoption, intensity & task use capture different dimensions of AI use.
> Together, they show how AI is entering work and everyday life.

What Work Does Generative AI Do?



NBER WORKING PAPER SERIES

WHAT WORK DOES GENERATIVE AI DO?

Alexander Bick
Adam Blandin
David J. Deming
Tyler R. Schumacher

Working Paper 35677
http://www.nber.org/papers/w35677

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2026

We thank our discussant Kristina McElheran and seminar/conference participants at the Federal Reserve Bank of St. Louis and the NBER "AI and Economic Measurement" conference for helpful comments and suggestions. Kevin L. Bloodworth and Khanh Le provided excellent research assistance. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the National Bureau of Economic Research. This work was supported by the Walmart Foundation, the Harvard Business School AI Institute, and a joint grant from the Washington Center for Equitable Growth and the Russell Sage Foundation.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Alexander Bick, Adam Blandin, David J. Deming, and Tyler R. Schumacher. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Working paper · First page

Alexander Bick, Adam Blandin,
David J. Deming & Tyler R. Schumacher

NBER Working Paper 35677 · August 2026

Open the paper

The Real-Time Population Survey (Bick & Blandin, 2023)

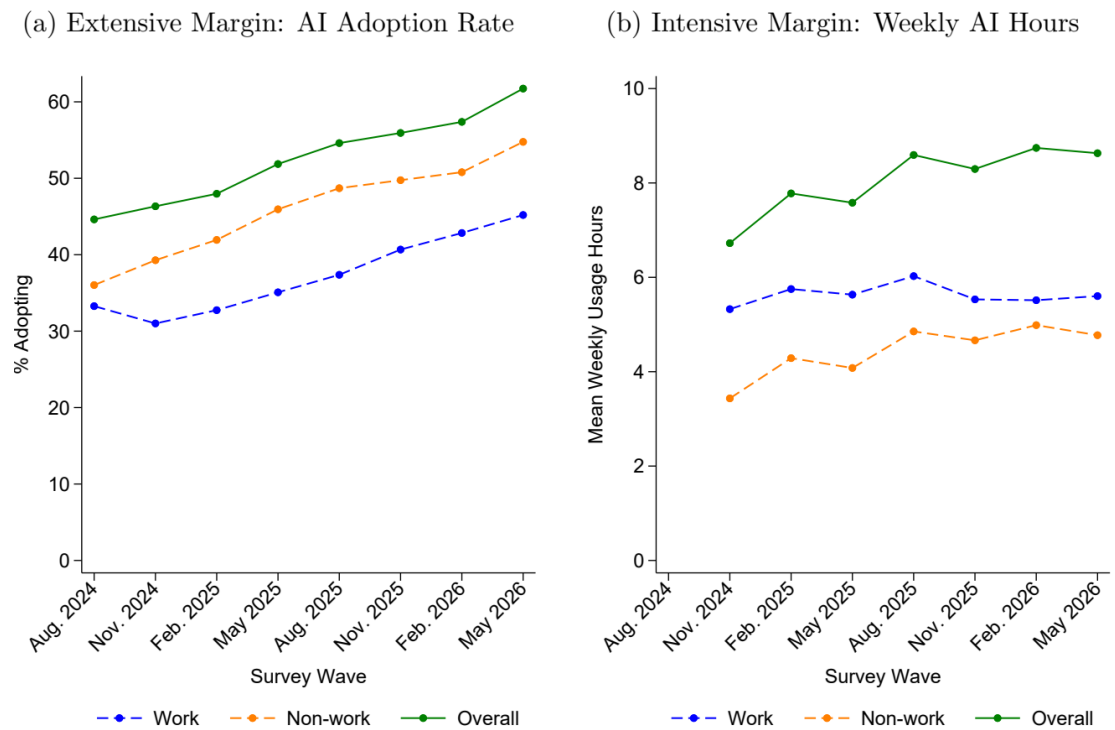
U.S. workers, ages 18–64 · Four waves, Aug. 2025–May 2026

- Worker's occupation ⟶ linked to O*NET
  - O*NET provides Detailed Work Activities (DWAs) for each occupation
  - Credit analysts: analyze business or financial data; prepare reports or budgets
- Workers select the tasks they perform from their occupation's top 10 DWAs
- GenAI users report which tasks they regularly use GenAI to help with

> Task-level adoption: reported GenAI use, not just potential exposure.

$$A_t = \frac{\sum_j w_j \mathbf{1}\{j \text{ performs } t \text{ and uses AI for } t\}}{\sum_j w_j \mathbf{1}\{j \text{ performs } t\}}$$

More people adopt, while users' work intensity changes less



Left: share using GenAI. Right: weekly hours conditional on use. Work rates refer to employed adults; overall/non-work rates to adults aged 18–64.

> By May 2026, 45% of workers report work use. Much of the recent growth comes from more users; weekly work-use hours per user remain around 5–6.

Adoption reaches many occupations without saturating them

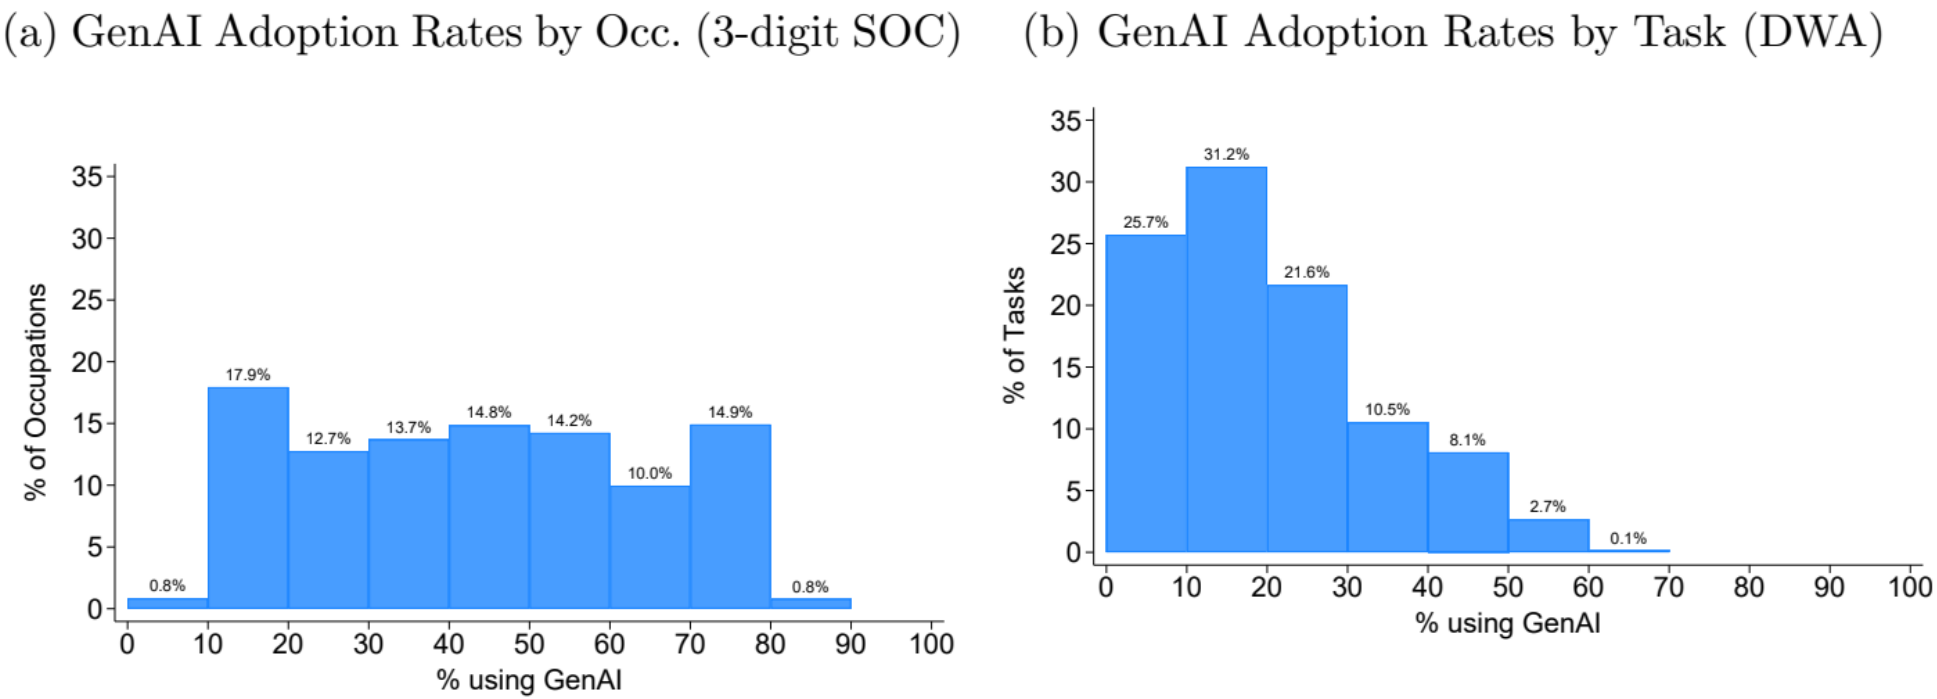
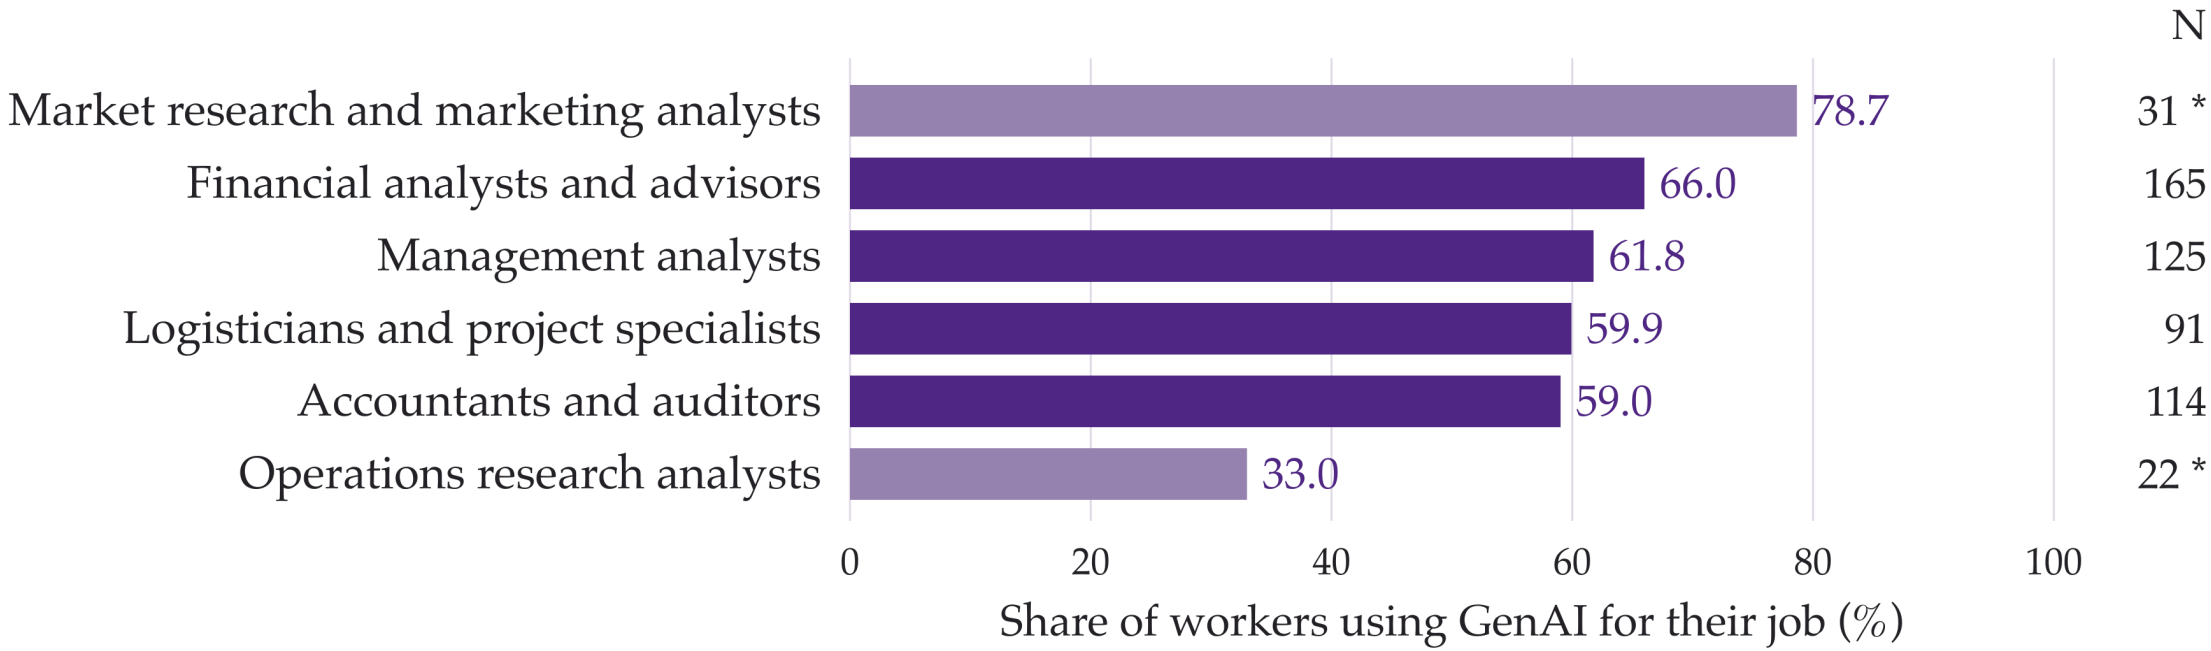


Figure 6: distribution of adoption rates across occupations (left) and tasks (right).

> Most occupations have some use, but few have adoption above 70%. No task reaches 70%: people doing the same activity often make different adoption choices.

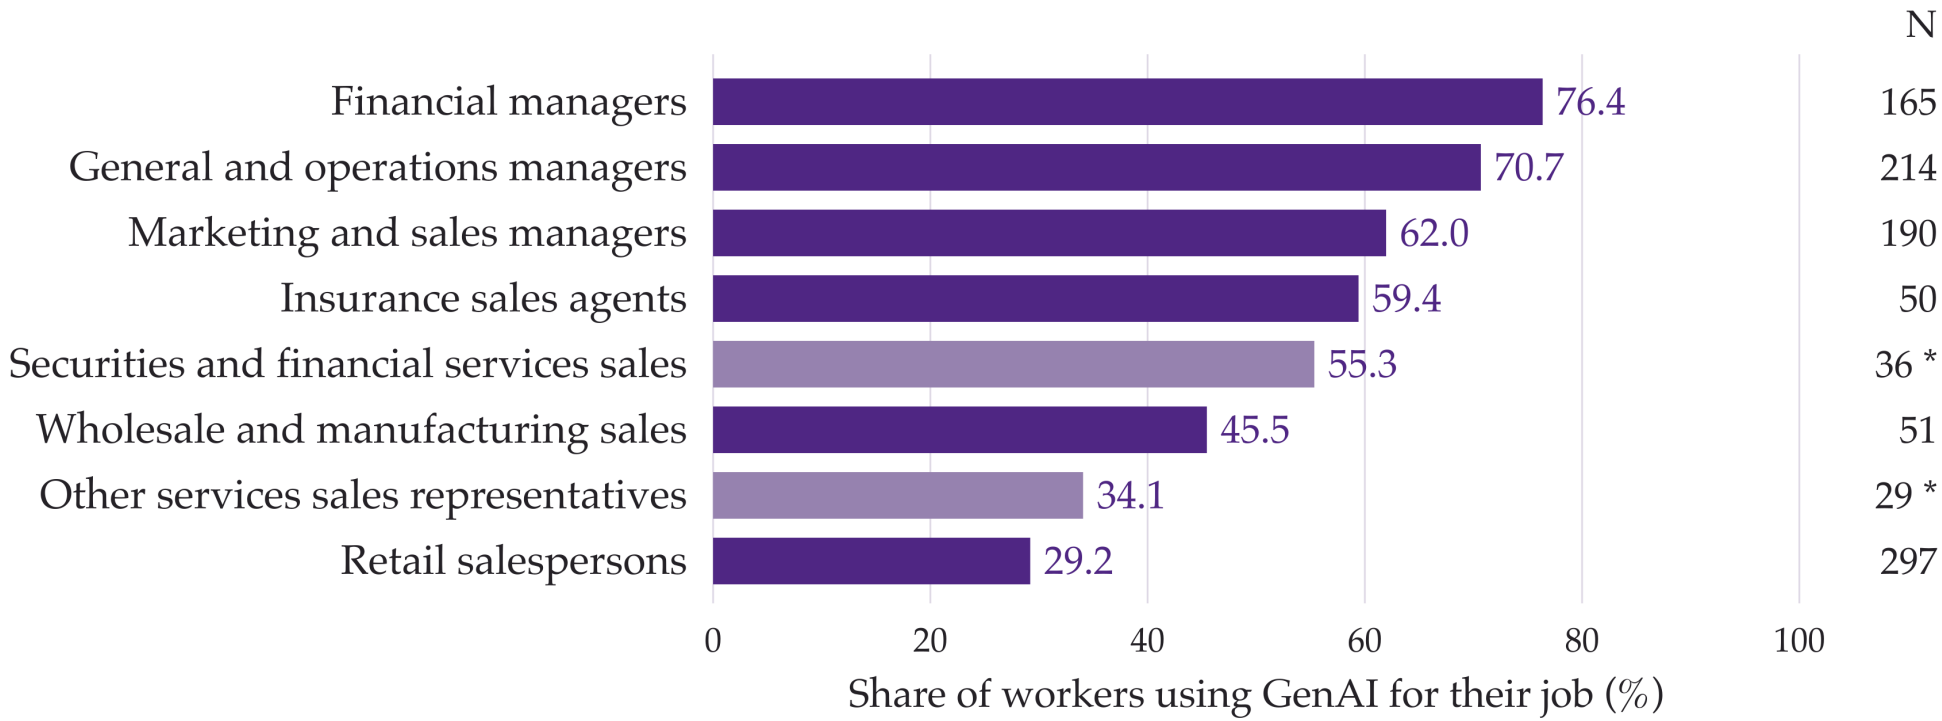
Adoption in analytical and financial careers



2018 SOC broad occupations. * Lighter bars: 20–49 observations; treat these estimates cautiously. N is unweighted.

> Financial analysts and advisors: 66.0%; management analysts: 61.8%; accountants: 59.0%. GenAI use is already common in these career areas.

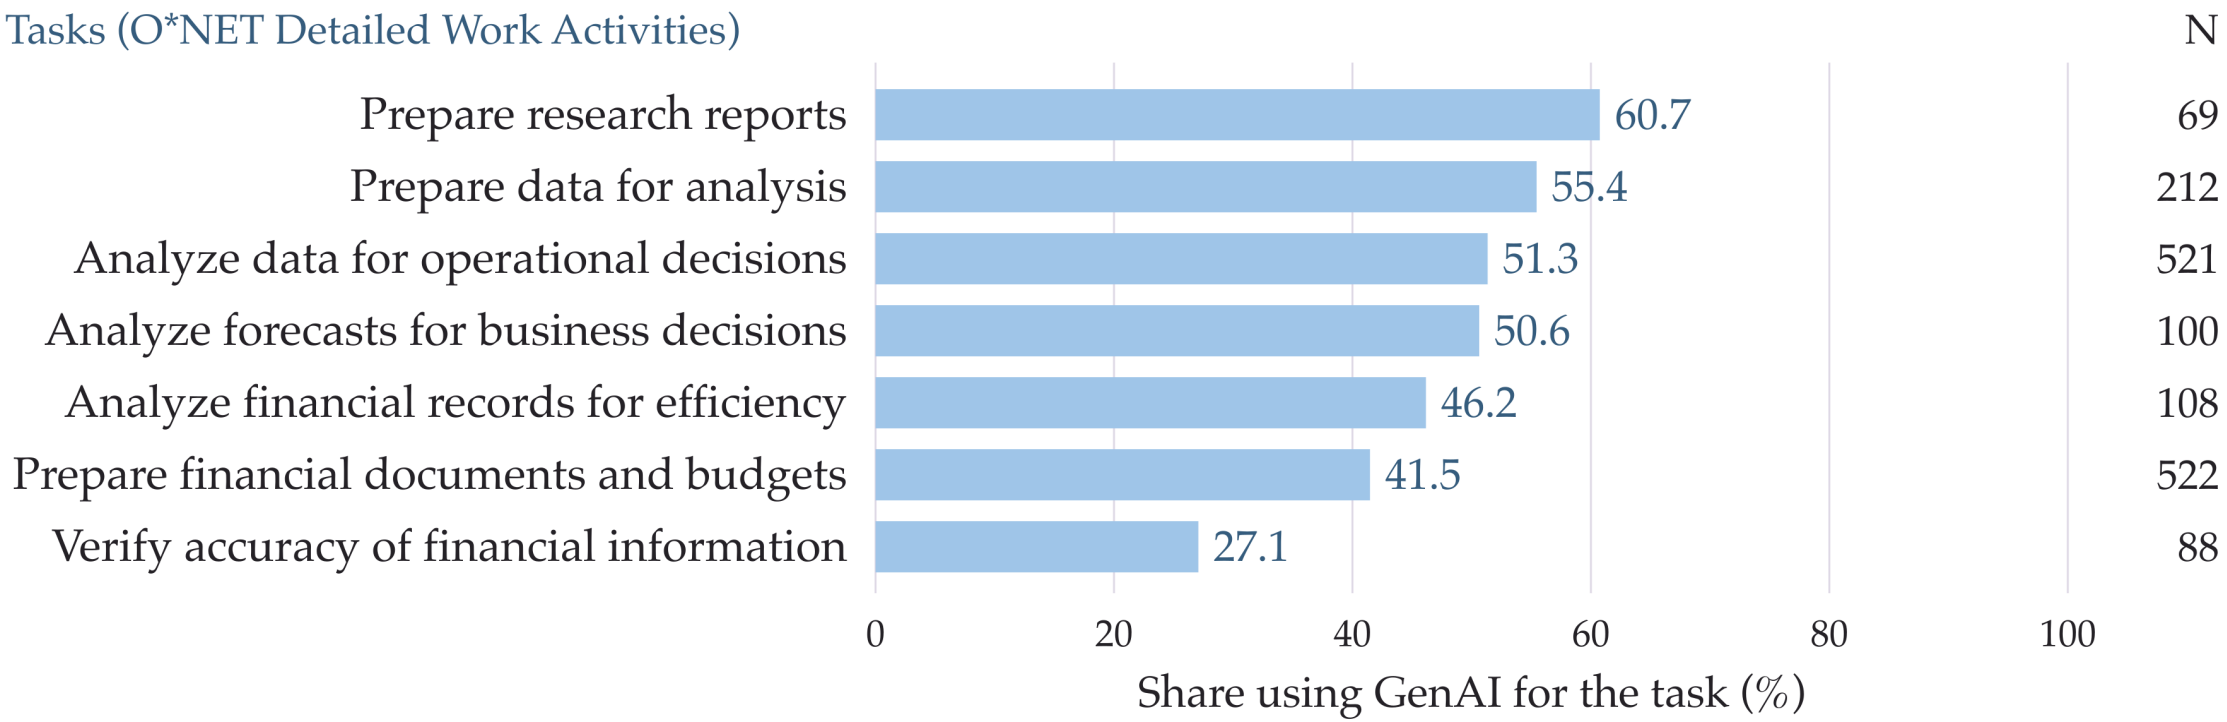
Sales roles and management roles have different adoption rates



Use anywhere in the job. * Lighter bars: 20–49 observations. Download includes official occupation codes.

> Financial managers: 76.4%; wholesale/manufacturing sales: 45.5%; retail sales: 29.2%. A single "sales" average would hide these differences.

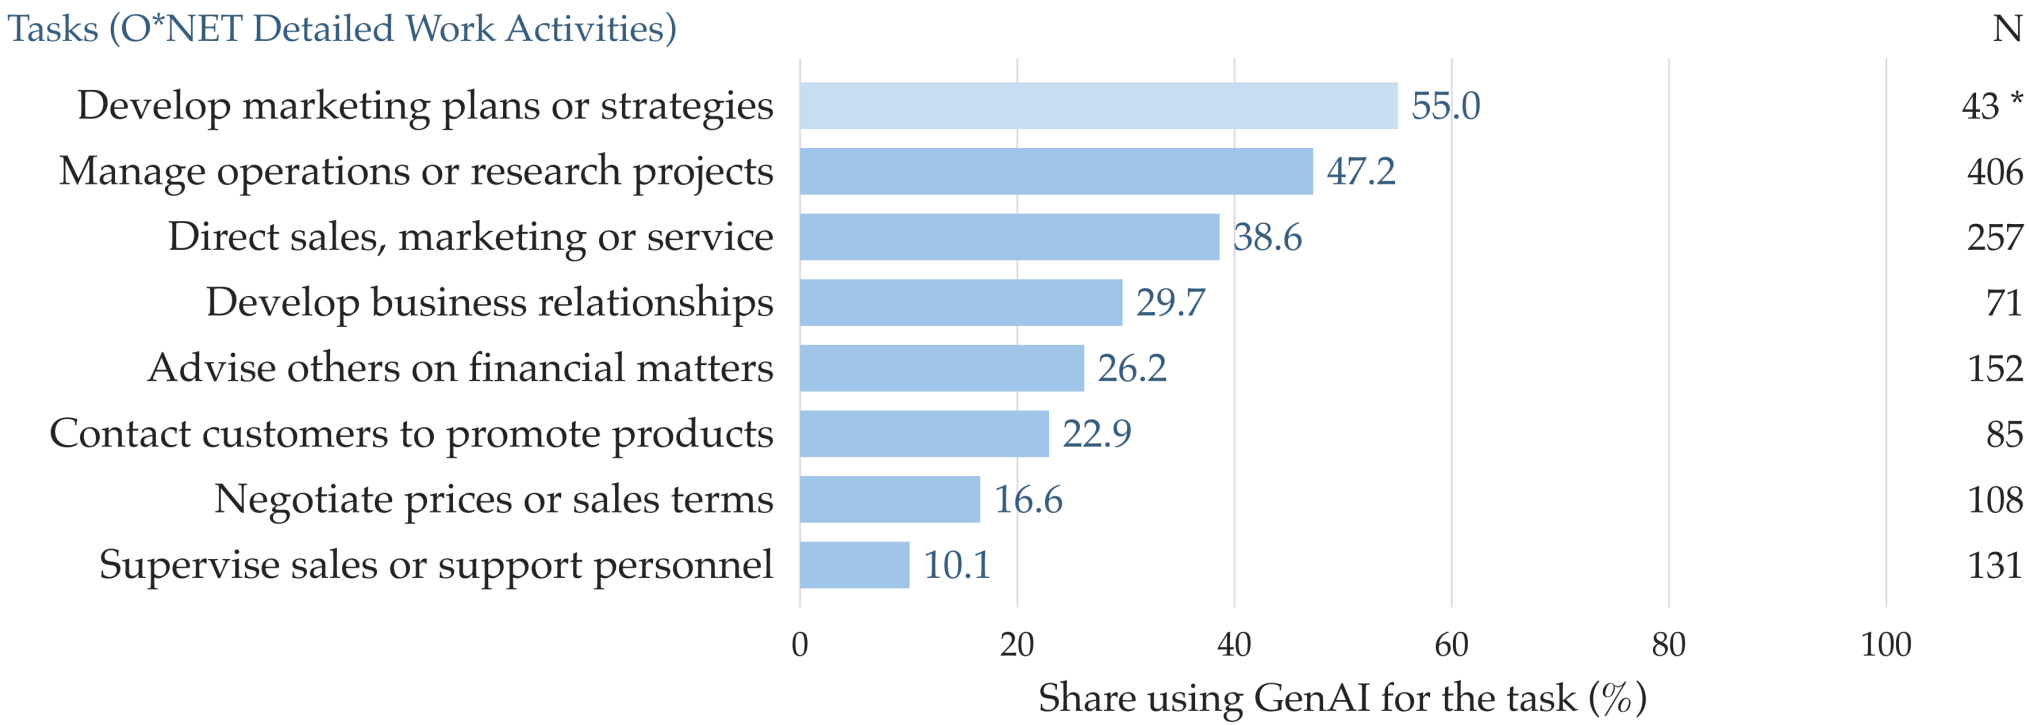
Task adoption: research, data preparation & analysis



Use among those performing each task, pooled across occupations. These are not rates within finance alone.

> Research reports: 60.7%; data preparation: 55.4%; checking financial accuracy: 27.1%. Even analytical workflows contain activities with very different use.

Task adoption: planning, client work & sales



Use among those performing each task, across occupations. * Marketing strategy: N = 43; other cells: N ≥ 71.

> Marketing planning: 55.0%; negotiating prices: 16.6%; supervising sales/support staff: 10.1%. Using AI somewhere in a job does not mean using it for every responsibility.

Explore these career areas and their activities

Management Analysts

Adoption rate	Unweighted observations
61.8%	125

RPS: August/November 2025 and February/May 2026. Weighted rate; respondent observations. Denominator: workers in this occupation with valid survey information. Use the HTML presentation to select other occupations and activities.

Exposure predicts some of adoption

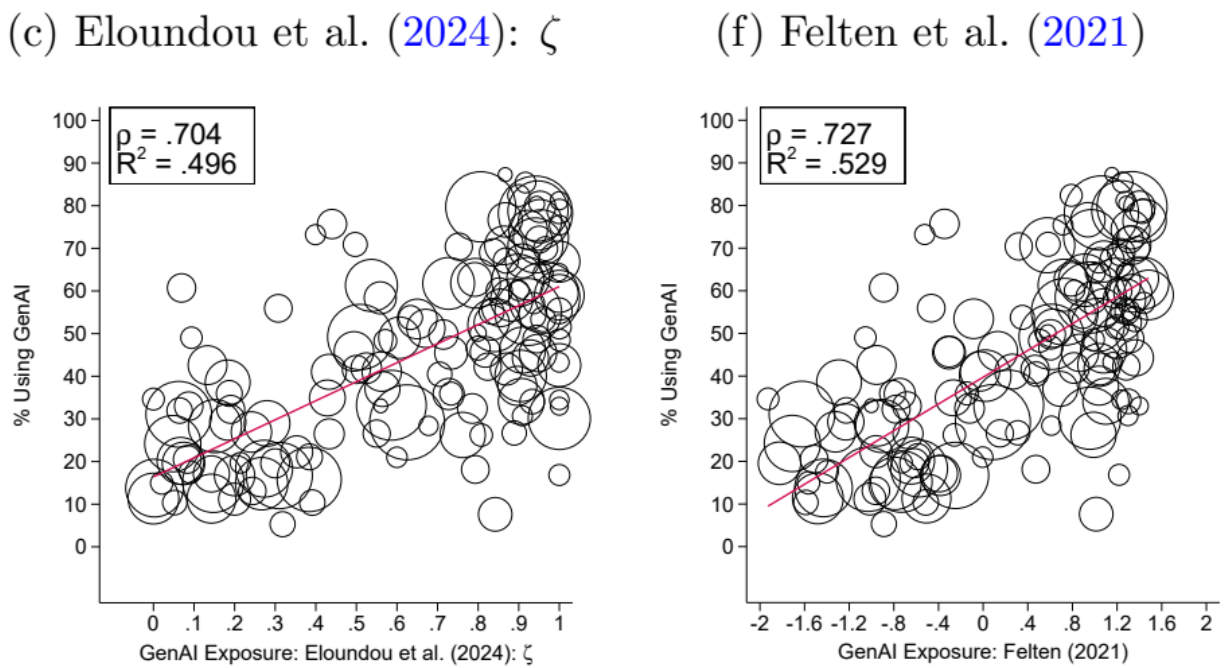


Figure 7C,F: adoption versus Eloundou high-capability (ζ) + Felten exposure; occupation-size weights.

> Some measures rank occupational adoption fairly well. That does not mean they identify which worker within an occupation uses AI.

How much individual adoption does exposure explain?

Adjusted R^2

Predictor set (separate regressions)	**Worker use**	**Worker–task use**
Occupation indicators	0.176	0.111
Task (DWA) indicators	—	0.107
Eloundou et al. ζ exposure	0.081	0.028
Felten et al. exposure	0.086	—

Adjusted R^2, not coefficients. Outcomes are use indicators; pooled four waves, with wave effects. Table 2.

> An occupational exposure score explains about 8% of variation in worker use. The task exposure score explains about 3% of worker–task use.

Exposure explains only part of adoption

Exposure scores explain 5–53% of adoption variation across occupations and 7–44% across tasks.

Predicted adoption from Eloundou et al. (2024) ζ exposure

Occupation / task	**Actual (%)**	**Predicted (%)**
Tellers	18.1	51.8
Customer-service representatives	30.1	61.0
Chief executives	79.7	52.4
Develop business/marketing plans (task)	52.1	27.8

Predicted adoption is fitted from Eloundou ζ exposure, not the raw exposure score. Table 3.

> High predicted exposure does not guarantee high adoption.

Why exposure and adoption can diverge

- Worker skills & incentives
  - Computer skills, experience & expected gains affect willingness to adopt
- Task & workplace barriers
  - Confidential data, compliance & rigid workflows can impede clerical uses
- Autonomy + easier integration
  - Freedom to experiment with coding, business plans or organizational problems
- Exposure scores may also misjudge actual productivity gains

> Potential gains lead to adoption only when they outweigh the barriers.

Who the worker is adds substantial predictive information

Predictor set	Adjusted R^2
Task (DWA) indicators	0.112
Task indicators + demographics	0.127
Task indicators + worker indicators	0.440

55,634 worker–task observations; people performing at least two tasks. Table 4.

> Demographics add little. A worker's general propensity to use AI adds much more: adoption differs greatly among people doing similar work.

The Rapid Adoption of Generative AI

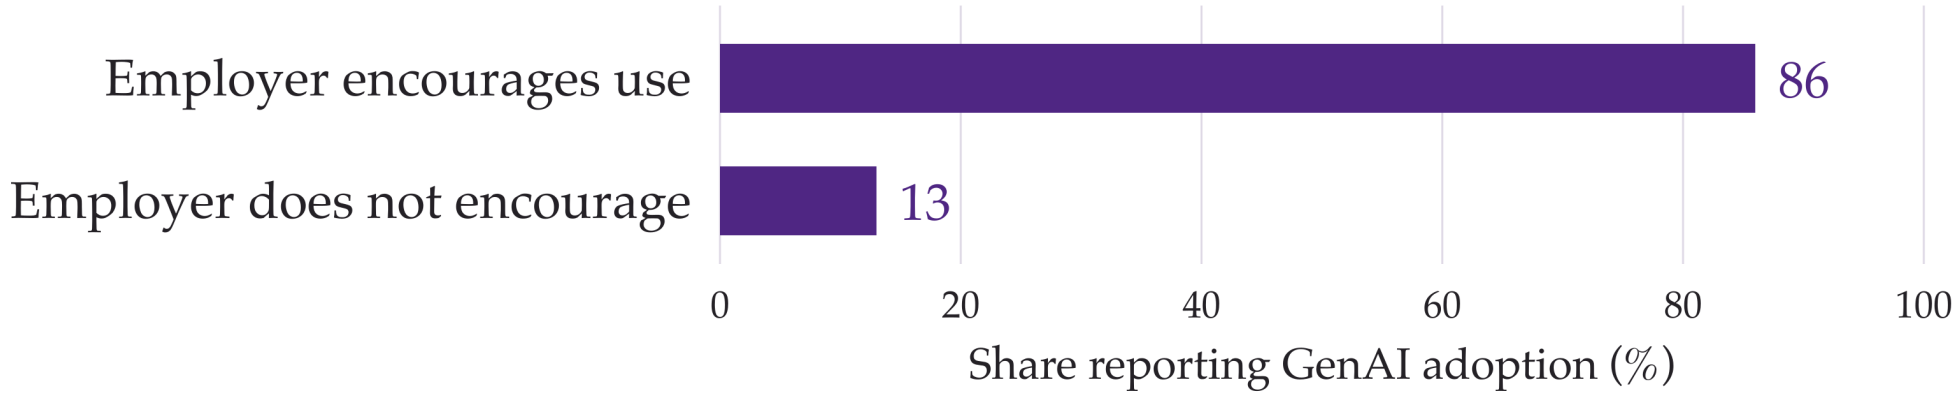


Publisher's article page · RPS, August/November 2024

Management Science 72(10), 8440–8456 · 2026 · Open the paper

> AI adoption is higher among younger and more educated workers, and higher among men (35%) than women (29%). See Figure 3 ⟶

Employer encouragement is strongly associated with adoption



Published descriptive comparison from Rapid Adoption, August/November 2024 evidence. Employer encouragement is reported by the worker.

> Workers with similar qualifications can face very different workplace environments. To interpret the 86% versus 13% gap, we need to examine what else differs.

Education and workplace environment in the adoption regression

Outcome: used GenAI at work last week	**(1)**	**(2)**	**(3)**
Bachelor's degree	0.193 (0.014)	0.135 (0.015)	0.045 (0.012)
Graduate degree	0.236 (0.017)	0.157 (0.019)	0.044 (0.015)
Employer encourages AI use	—	—	0.651 (0.014)
Occupation and exposure controls	No	Yes	Yes
Adjusted R^2	0.08	0.12	0.51

N = 6,898 per column; age + sex controls; reference: below bachelor's. SEs in parentheses.

> The bachelor's gap falls from 19.3 to 4.5 percentage points as controls are added. Workplace encouragement carries substantial predictive information.

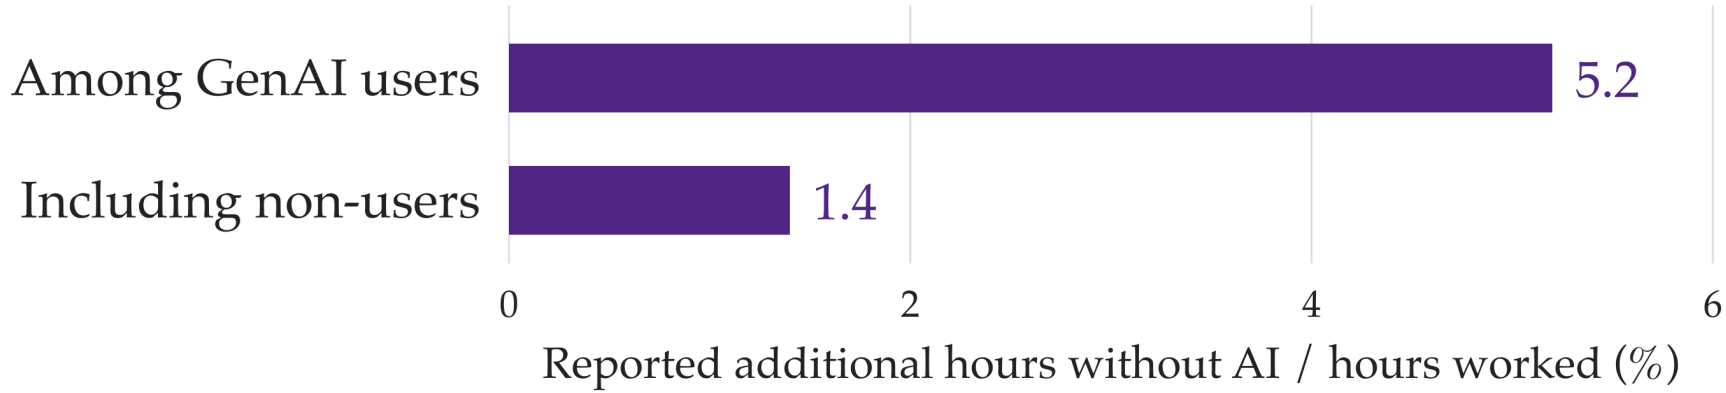
Experienced users apply AI to more tasks

Outcome: number of AI-assisted tasks	Estimate
Used GenAI at least six months earlier	0.429 (0.113)
Observations	6,565
Adjusted R^2	0.441

Current users; occupation, wave, age, sex, and number of tasks controlled for. Robust standard error in parentheses.

> Earlier users report about 0.43 more AI-assisted tasks. Experience is associated with broader use, even holding their task count fixed.

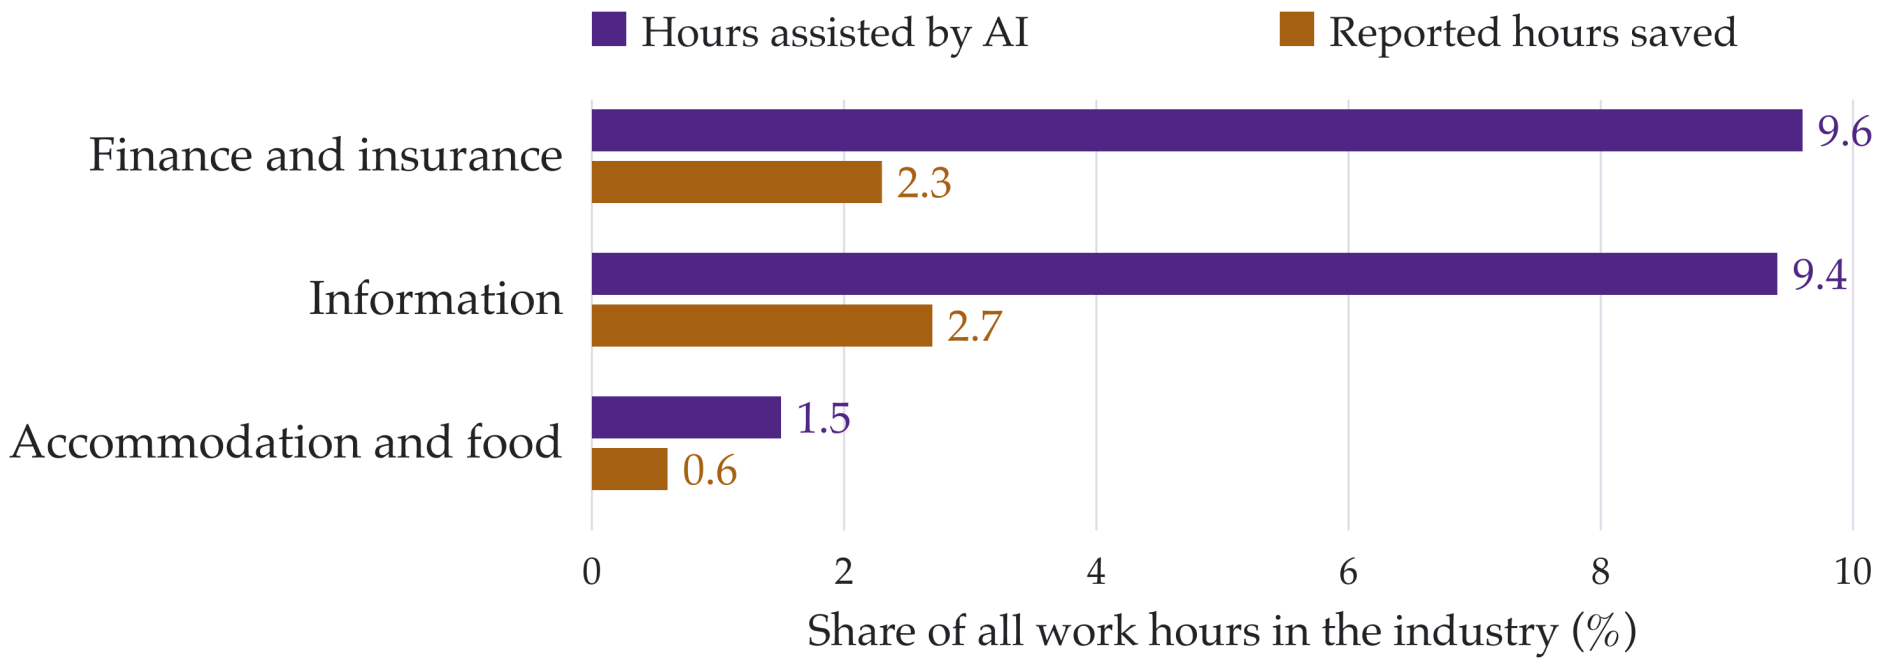
Reported time savings depend on who is included



November 2024. Self-reported extra hours needed to complete the same work without GenAI, relative to hours worked.

> For a 40-hour week, 5.2% corresponds to about 2.1 hours. Including workers who do not use GenAI reduces the average to 1.4%.

Finance has relatively intensive use and reported time savings



November 2024; shares of total industry work hours, including non-users. Published values for three selected industries.

> In finance, 9.6% of work hours are assisted by AI and reported savings equal 2.3% of hours. Time using a tool and time saved by it are different measures.

Reported savings do not establish measured productivity gains

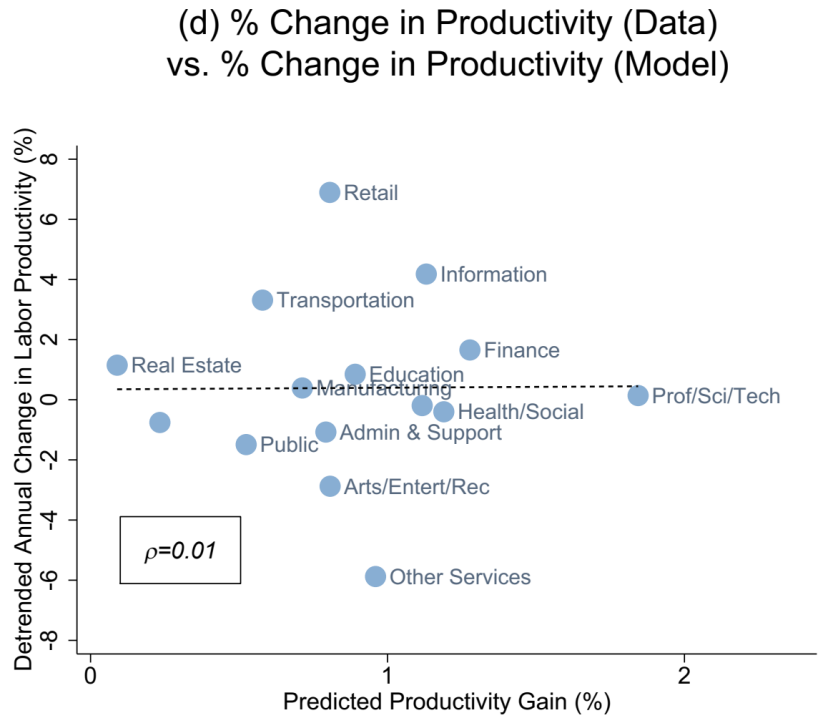
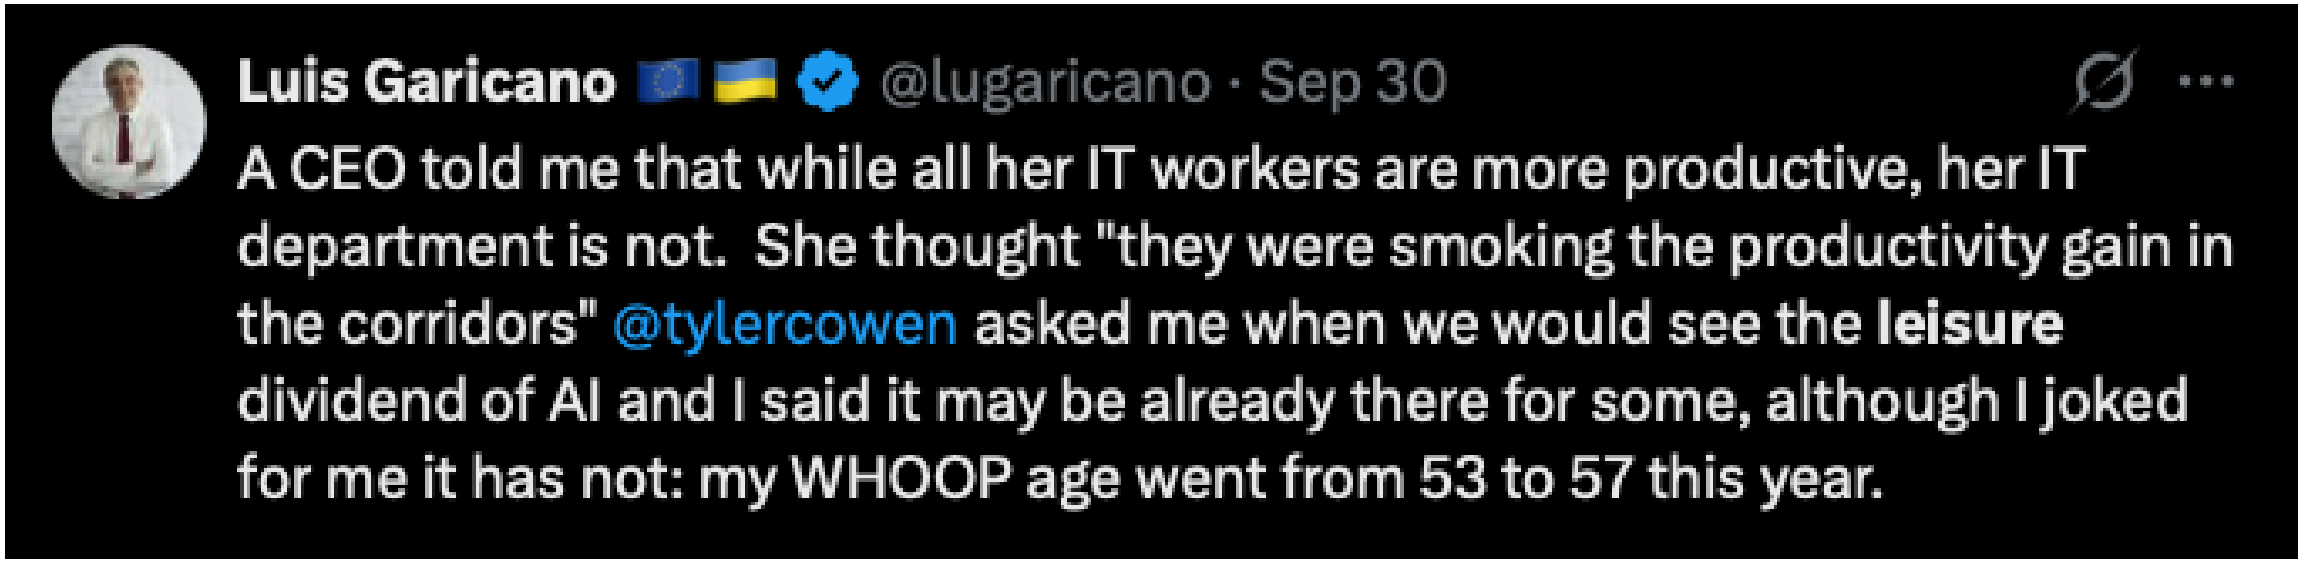


Figure 9d. Horizontal: model-implied gain from reported savings. Vertical: detrended annual productivity growth, 2022–24.

> Across these industries, the correlation is 0.01—essentially zero. The survey suggests potential gains, but this comparison does not validate them in measured productivity.

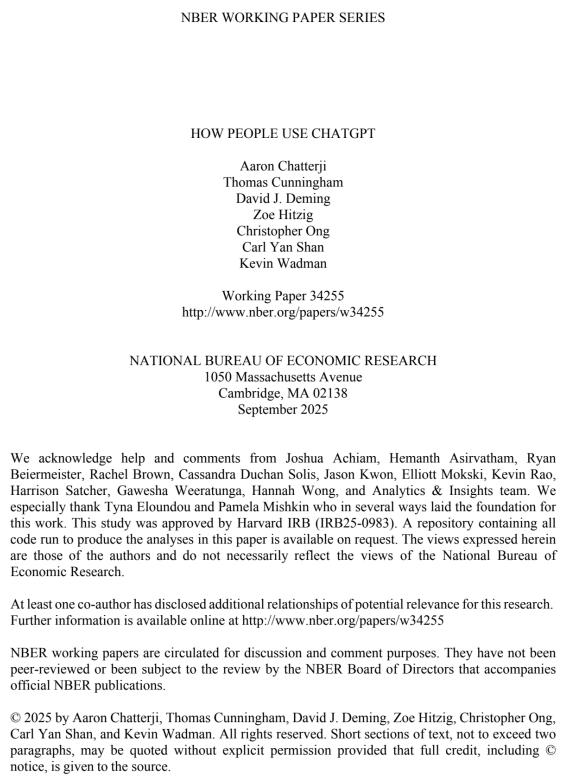
Time savings do not automatically become cost savings



- Reassign tasks & remove bottlenecks ⟶ more output or fewer paid inputs

> Saving time on individual tasks need not reduce the firm's costs.
> The payoff depends on how tasks, staffing and workflows are reorganized.

How People Use ChatGPT



NBER WORKING PAPER SERIES

HOW PEOPLE USE CHATGPT

Aaron Chatterji
Thomas Cunningham
David J. Deming
Zoe Hitzig
Christopher Ong
Carl Yan Shan
Kevin Wadman

Working Paper 34255
http://www.nber.org/papers/w34255

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2025

We acknowledge help and comments from Joshua Achiam, Hemanth Asirvatham, Ryan Beiermeister, Rachel Brown, Cassandra Duchan Solis, Jason Kwon, Elliott Mokski, Kevin Rao, Harrison Satcher, Gawesha Weeratunga, Hannah Wong, and Analytics & Insights team. We especially thank Tyna Eloundou and Pamela Mishkin who in several ways laid the foundation for this work. This study was approved by Harvard IRB (IRB25-0983). A repository containing all code run to produce the analyses in this paper is available on request. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w34255

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Aaron Chatterji, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Aaron Chatterji, Thomas Cunningham, David J. Deming
Zoe Hitzig, Christopher Ong
Carl Yan Shan & Kevin Wadman

NBER Working Paper 34255 · September 2025

Open the paper

Working paper · First page

From ChatGPT messages to measures of use

- Consumer ChatGPT only: Free, Plus & Pro; May 2024–June 2025
- Sample ~1.1 million conversations ⟶ one user message from each
- Remove identifying details ⟶ LLM classification using prior messages
- Labels: work/non-work; topic; Asking, Doing or Expressing
- Validate against human labels on public WildChat conversations
- Occupation comparisons: ~40,000 matched users; aggregate outputs only

> Researchers see classifications + aggregate statistics, not raw conversations.
> The results describe message shares among users, not workforce adoption.

What people do inside ChatGPT

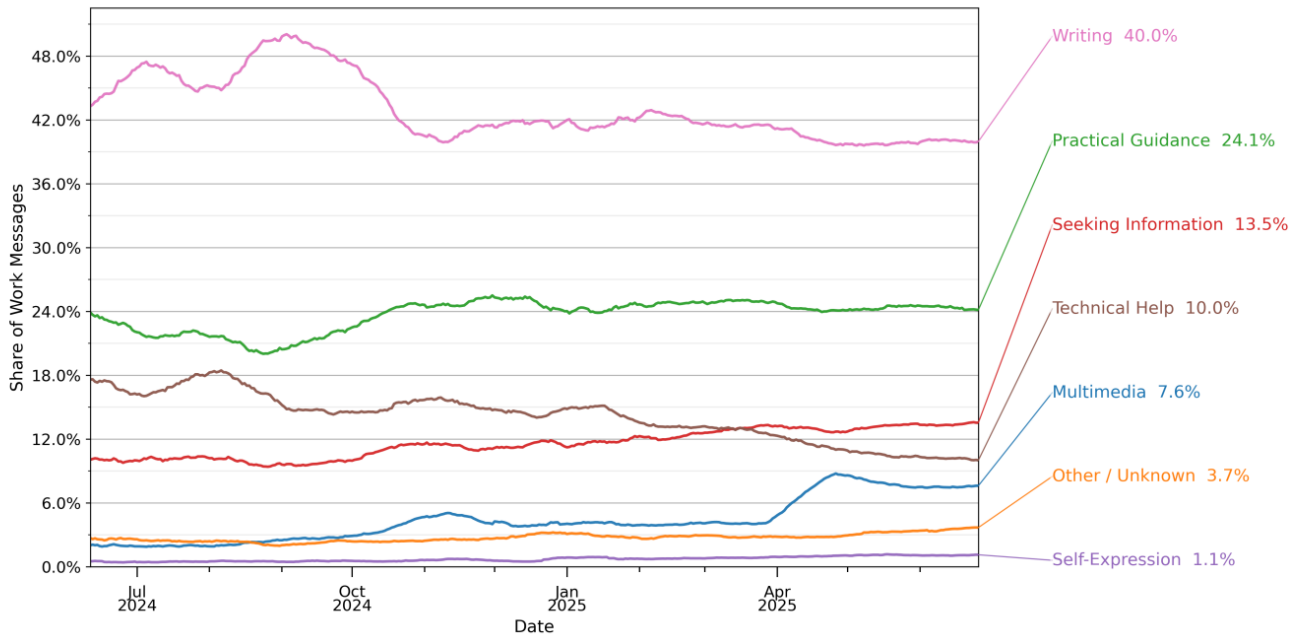
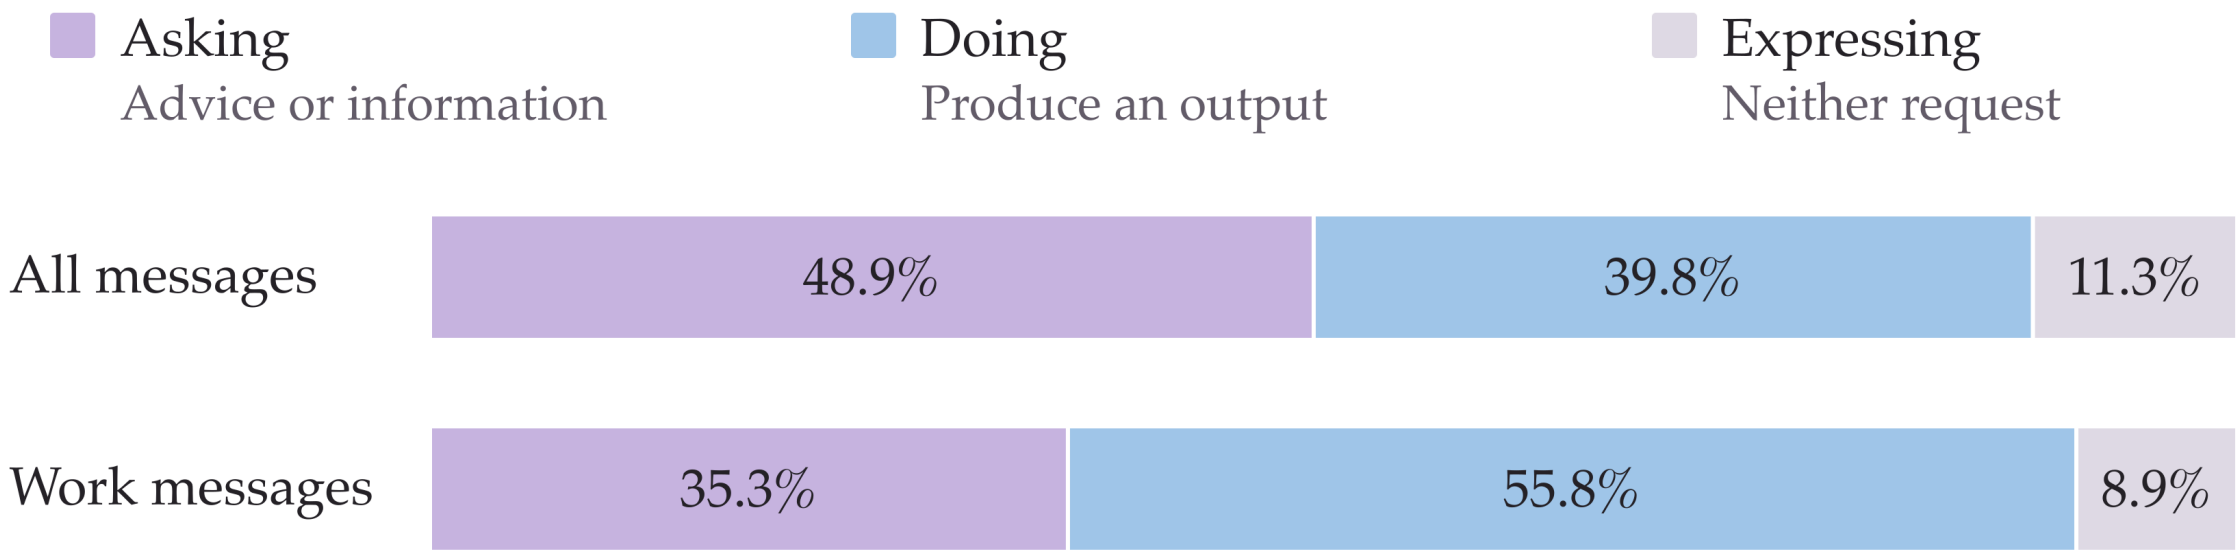


Figure 8: topics within work-related consumer ChatGPT messages, through July 2025.

> Writing is about 40% of work messages; practical guidance is about 24%. Much workplace use helps produce or revise an output.

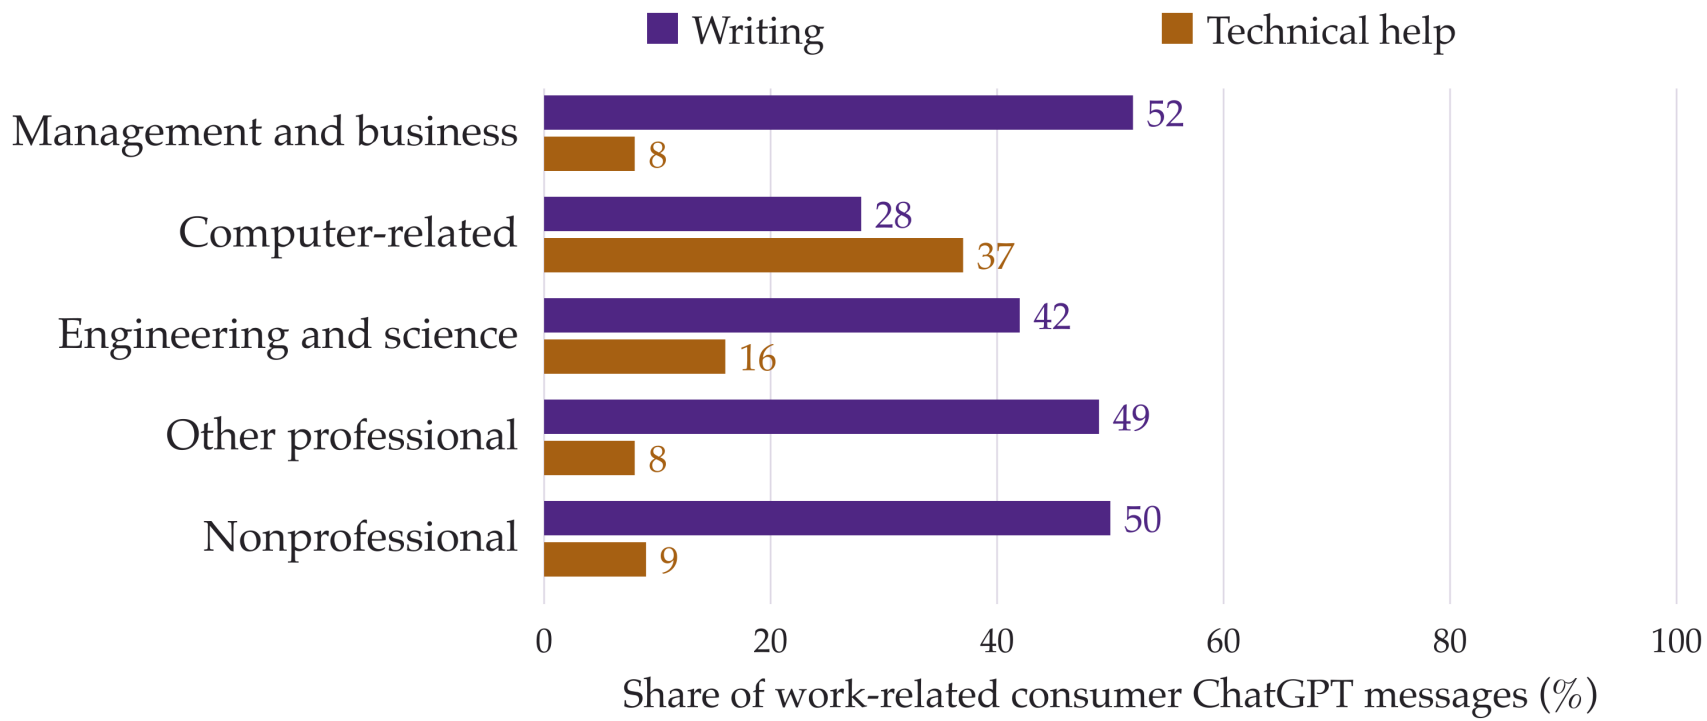
ChatGPT supports decisions as well as producing outputs



Message shares, May 2024–June 2025. Redrawn from Figures 10–11; each row sums to 100%.

> Over one-third of work messages seek advice or information.
> AI can support decisions as well as produce a draft, calculation or other output.

Business users' work conversations concentrate on writing



Published unadjusted shares, Figure 23 C1–C2. Two topic categories shown; the remaining topics account for the rest.

> Writing accounts for 52% of management/business work messages. Computer-related users devote 37% to technical help: the use case varies with the work.

What the results tell us about AI diffusion

- Substantial adoption in analytical, finance & management jobs; uneven task use
- Worker + workplace differences matter beyond technical exposure
- Reported savings + chat logs: use patterns, not causal wage/productivity effects

> The next lecture asks what changes when otherwise comparable workers gain access to AI.

Appendix: Accessing the adoption data

RPS public data

Open in a browser tab

1. Open occupation + task adoption workbooks
2. Read the README before the estimates
3. Choose an aggregation level; keep its identifier
4. Check rates, sample counts & suppression rules

Appendix: which career estimates can the public file support?

Broad occupation	Unweighted N	Released adoption rate
Economists	14	Suppressed
Credit analysts	9	Suppressed
Data scientists	8	Suppressed
Operations research analysts	22	33.0%
Market research / marketing analysts	31	78.7%

> A missing estimate does not mean no adoption. Small released cells also require caution; use broader groups when the question allows it.

Appendix: A regression can organize the comparison

$$\text{Adoption}_o = a + b\,\text{Exposure}_o + u_o$$

A positive slope says that more exposed occupations tend to use AI more.

The residual contains measurement differences and other determinants of adoption.

It does not, by itself, identify the effect of exposure or adoption on productivity.

Appendix: Does experience carry over to other activities?

Same focal activity; different exposure to AI use through other activities.

The instrument averages **other workers' adoption rates** on that person's other tasks. It predicts their own use outside the focal activity.

> For a causal learning interpretation, the task portfolio must affect focal-task adoption through that other use, rather than through unobserved worker or workplace differences.

Appendix: The spillover estimates are positive

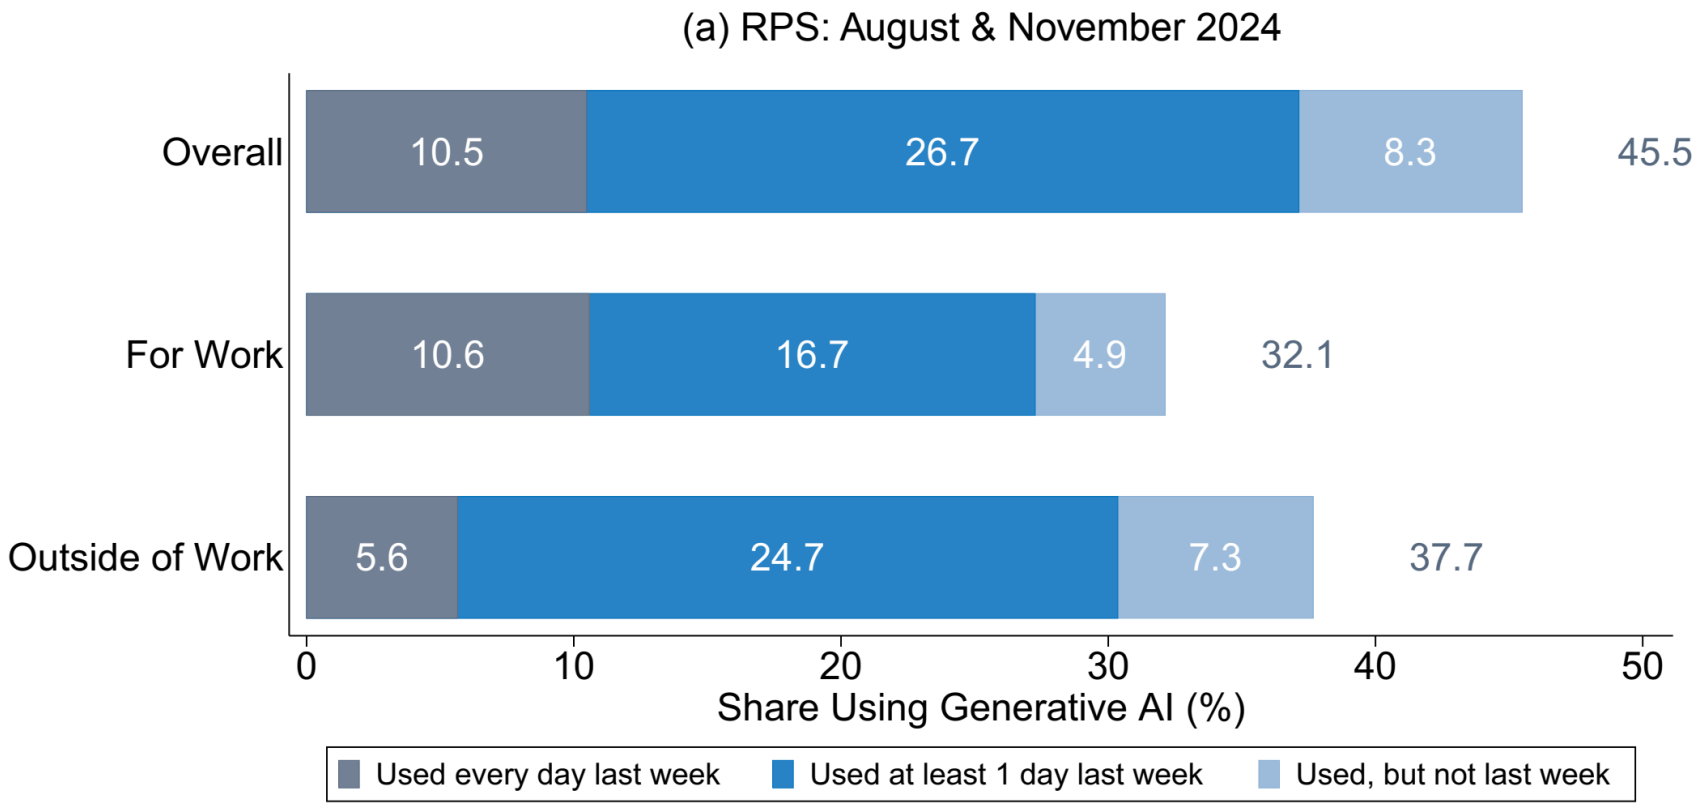
Outcome / explanatory variable	OLS	IV
Focal-task use / other-task use share	0.621 (0.011)	0.545 (0.067)
Home use / work-use indicator	0.371 (0.013)	0.504 (0.129)

Table 6, col. 2: 55,091 worker–IWA observations. Table 7: 11,609 workers. Reported standard errors in parentheses.

A 10 percentage-point increase in other-task use corresponds to a **5.45-point** increase in focal-task use in the IV specification.

> Evidence of learning across activities + work-to-home spillovers, under the IV assumptions.

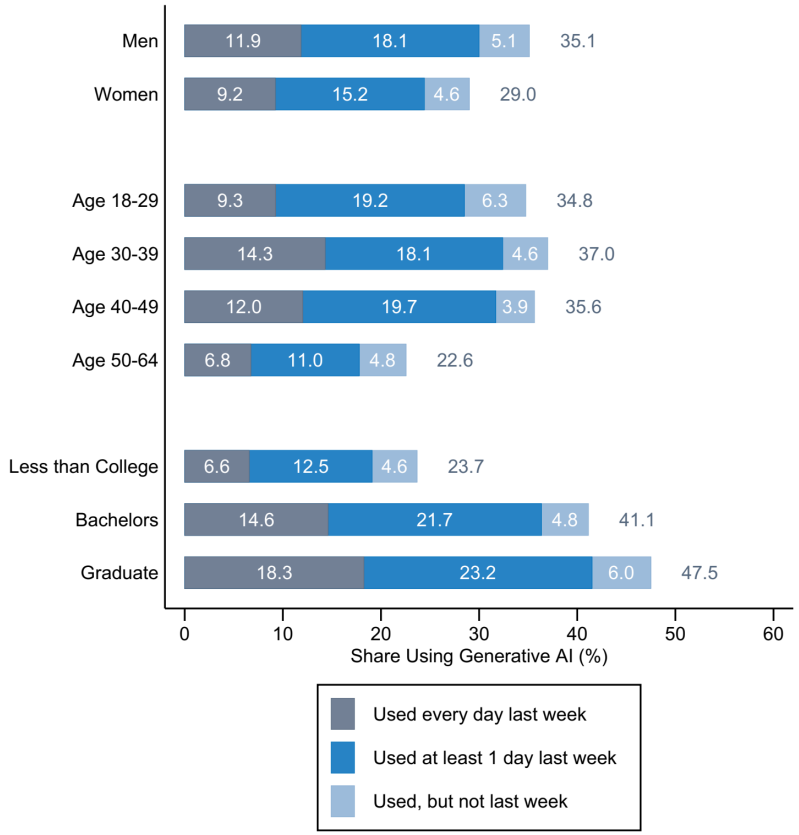
Appendix: How widespread is adoption?



RPS, August and November 2024. Bars separate daily, some-day, and earlier use.

The answer changes when we ask about using AI at all, using it last week, or using it every day.

Appendix: Who adopts AI at work? (Figure 3)



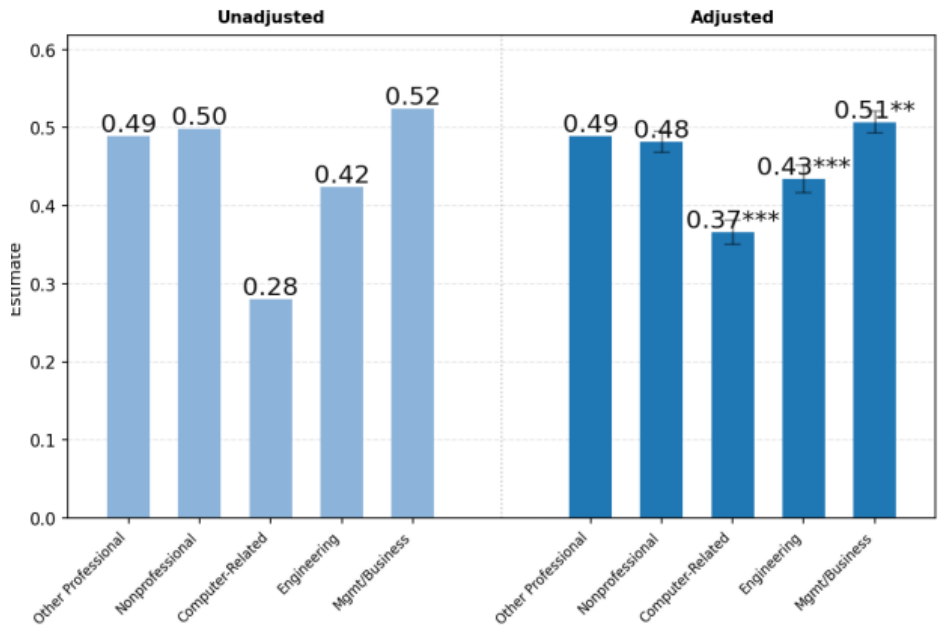
Employed adults aged 18–64; N = 6,935.

Work adoption

- Men: 35.1%; women: 29.0%
- Ages 18–49: 34.8%–37.0%
- Ages 50–64: 22.6%
- Below bachelor's: 23.7%
- Bachelor's: 41.1%; graduate: 47.5%

← Back to slide 17

Appendix: the business-writing result survives adjustment



Panel C1. Writing.

Original writing panel. Adjustments include education, age, name-based gender category, seniority, firm size, and industry.

> Management/business writing is about 52% before adjustment and 51% after. The result describes message composition among sampled users.

Appendix: The same task can have different adoption rates

Two workers can perform similar activities and face the same technology, yet only one uses AI.

- Employer access + training
- Suitable workflow + incentives to experiment
- Ability to share task information with an external tool