UNIVERSITY OF OXFORD

Saïd Business School

Corporate Finance

AI, Data, and the Future of the Corporation

Shumiao Ouyang

2026

impact *from within*

Course Logistics

Module leader

- Shumiao Ouyang, Associate Professor of Finance
  shumiao.ouyang@sbs.ox.ac.uk

Structure

- 2 blocks $\times$ 12 weeks
- Hybrid: in-person, online, async

Assessment

- 3-hour open-book online examination (**25%**) due 9 December 2026
- Pre-recorded presentation (**25%**) due 17 December 2027
- 2-hour closed-book in-person examination (**50%**) due 10 January 2028

Corporate Finance — Module Overview

UNIVERSITY OF OXFORD Saïd Business School

Block 1

- Value, investment, financing
- Capital structure
- Valuation and modelling
- Risk and real options
- VC, PE, and M&A

Block 2

- Market efficiency, behaviour
- Raising capital and IPOs
- Payout policy
- Debt, credit risk, distress
- Governance and control

Today's agenda

1. AI adoption and firm value

2. AI, organisations, and workers

3. Data and the modern firm

Five questions for the next two hours

- **1 — Value.** What does a firm get back from an AI investment?
- **2 — Evidence.** How do we know AI *caused* the growth?
- **3 — Organisation.** What happens to jobs, skills, hierarchies?
- **4 — Distribution.** Who captures the value?
- **5 — Data.** Is data a new form of capital, and can we price it?

Table of Contents

AI adoption is already widespread

∼6,000 firms surveyed in the US, UK, Germany, and Australia

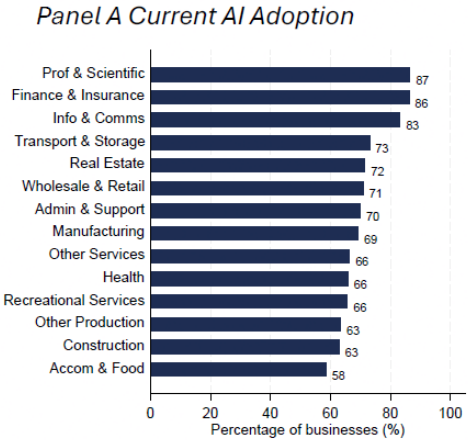

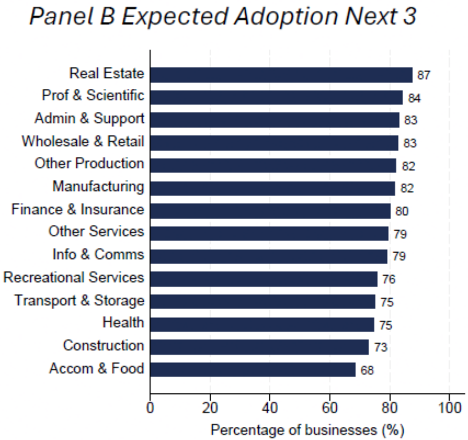

Source: Yotzov et al., 2026, Figure A6. Surveys by the Atlanta Fed, Bank of England, Bundesbank, and Macquarie University.

So far, AI has barely moved employment

Over 90% of firms report no impact; the average effect is essentially zero

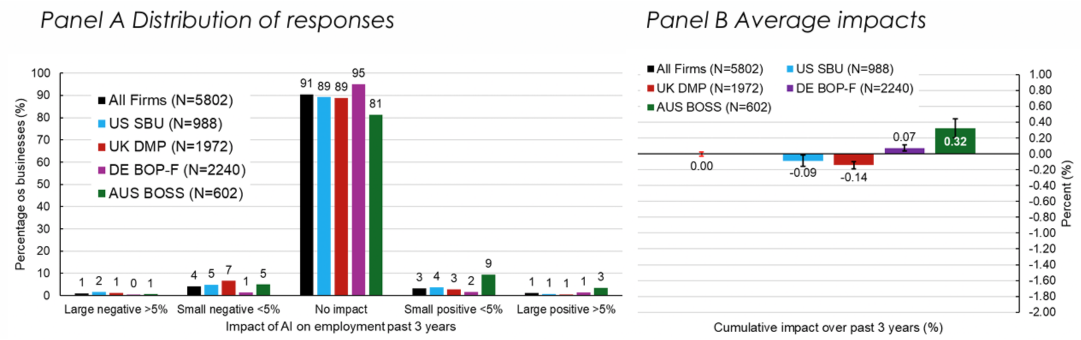

Source: Yotzov et al., 2026, Figure 8. Impacts quantified as ±7.5% (large), ±2.5% (small), 0% (none).

The next three years look different: firms expect job cuts

Expected impact: -0.7% on average — UK -1.4%, US -1.2%

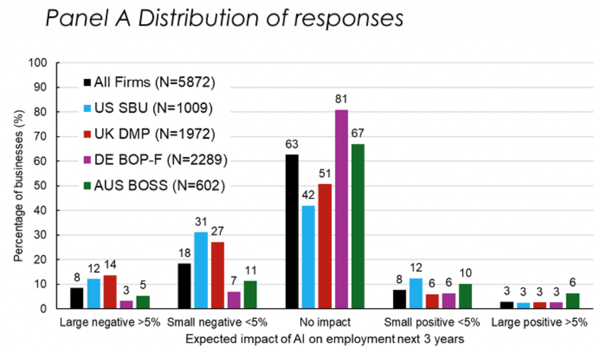

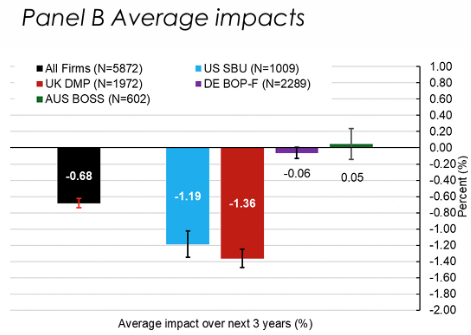

Source: Yotzov et al., 2026, Figure 9.

Productivity so far: a small but positive +0.29%

Unlike employment, realised responses already skew positive

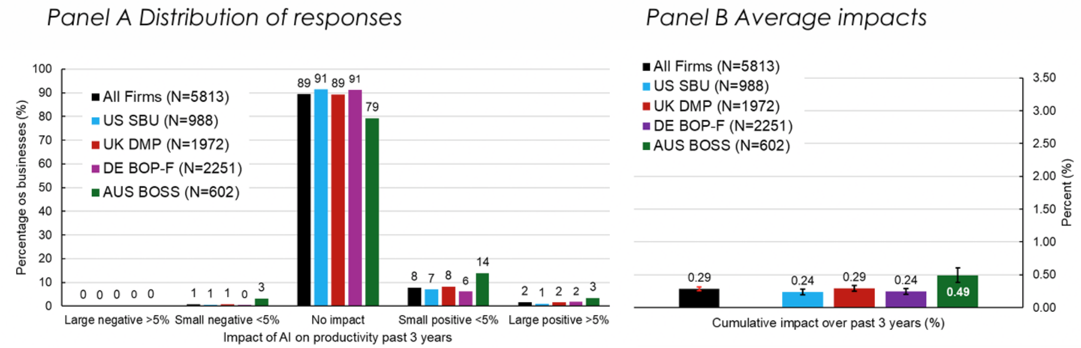

Source: Yotzov et al., 2026, Figure 10. Productivity = sales per employee.

Expected productivity: +1.4% — four times the realised gain

US +2.3%, UK +1.9%; 37% of firms expect gains ahead

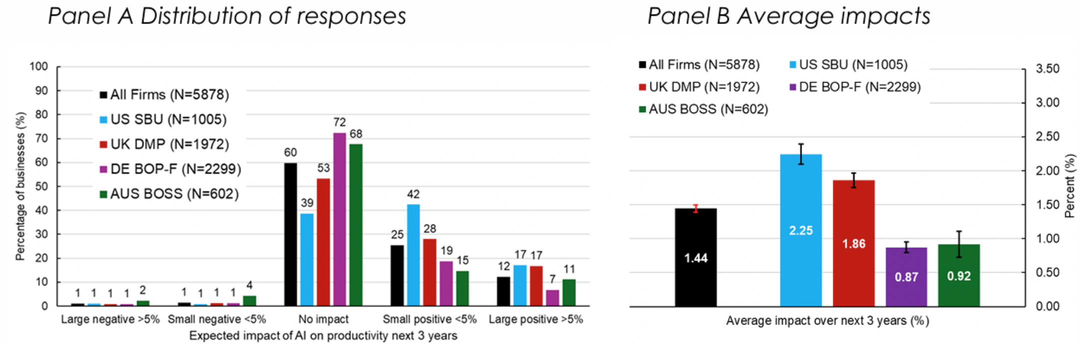

Source: Yotzov et al., 2026, Figure 11.

The adoption puzzle

Everyone is adopting; almost nothing has happened — yet

>90%
report no employment impact

+0.29%
realised productivity gain

+1.4%
expected productivity gain

- Hype — or the early phase of a GPT (Brynjolfsson, Rock, and Syverson, 2021)?
- Next: what firms *do*, and what markets *price*.

What kind of technology is AI, economically?

Four features that drive everything else today

1 — Prediction
Feeds every decision.

2 — General purpose
Like electricity.

3 — Human-embodied
No balance-sheet line.

4 — Data is the moat
Models free; data scarce.

Two channels from AI to firm value

Product innovation

Cheaper experimentation ⇒ **new products**.

Moderna used AI to search vaccine designs computationally — and had a COVID-19 vaccine designed in 65 days.

Footprint: more products and patents; employment *rises*.

Process innovation

Better forecasts, automation ⇒ **lower cost**.

UPS uses AI to optimise delivery routes — same parcels delivered, ∼100 million fewer miles a year.

Footprint: higher sales per worker; no new products.

Different footprints — the data can tell us **which channel dominates**.

Even if AI creates value, who captures it?

Three fault lines to watch all lecture

- **Firms vs. workers** — substitute or complement?
- **Large vs. small** — superstar firms?
- **Incumbents vs. entrants** — are rents defensible?

Value *creation* and value *capture* are different questions.

The measurement problem: AI never appears on the balance sheet

Babina, Fedyk, He, and Hodson, 2024, Journal of Financial Economics

- AI investment is **intangible and embodied in people** — it hides in wage bills, not capex.
- R&D and SG&A are too coarse; adoption surveys are sparse and self-reported.

The measurement problem: AI never appears on the balance sheet

Babina, Fedyk, He, and Hodson, 2024, Journal of Financial Economics

- AI investment is **intangible and embodied in people** — it hides in wage bills, not capex.
- R&D and SG&A are too coarse; adoption surveys are sparse and self-reported.
- The fix: **measure AI from who the firm hires** — watch the people who build the asset.

535m	**180m**
resume s: 64% of the US workforce (Cognism)	job postings, full text (Burning Glass)

Which jobs are "AI jobs"? Let the data decide

Skills scored by co-occurrence with core AI skills

- Score $\sim$15,000 skills by how often they co-occur with ML, NLP, and computer vision.
- A job is AI-related if its average skill score is high enough.
- Firm-level AI investment $=$ **share of employees in AI roles**, each year.

0.90

AI-relatedness of "TensorFlow"

0.003

AI-relatedness of "Microsoft Office"

AI investment grew sevenfold in a decade

The signature of a general purpose technology

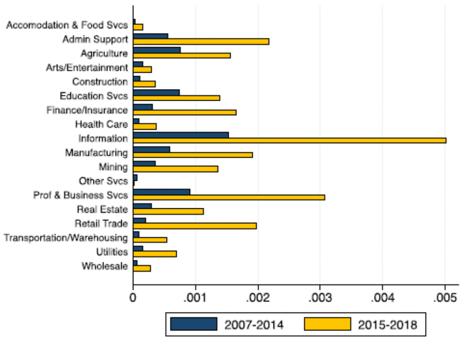

(a) Resume data (Cognism)

Stock: share of AI workers

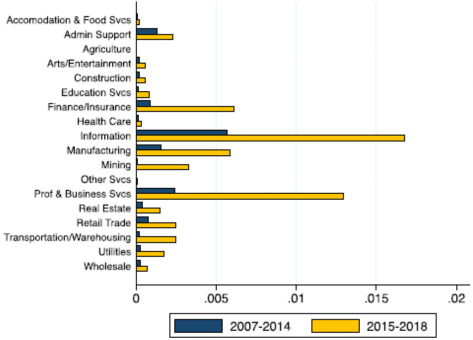

(b) Job posting data (Burning Glass)

Flow: share of AI postings

Source: Babina, Fedyk, He, and Hodson, 2024. Information leads (0.15% → 0.50%); larger, cash-richer firms adopt most.

Correlation is not causation — so use history as the experiment

IV: pre-2010 hiring ties to universities that later became AI powerhouses

- **Worry:** good firms hire AI talent *and* grow.
- **Instrument:** 2010 STEM hiring shares × pre-2010 university AI strength.
- **Relevance:** AI-strong schools tripled AI graduates after 2012 ($F \approx 13$–19).
- **Exclusion:** ties predate commercial AI; no pre-trends.

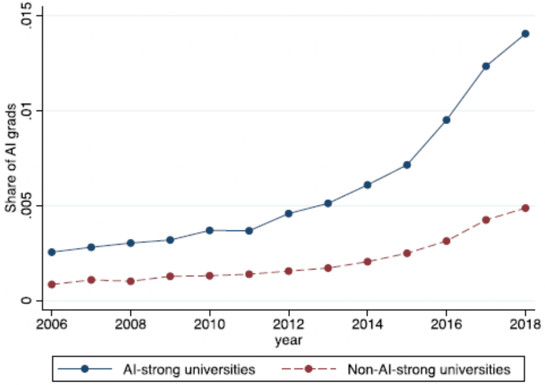

Source: Babina, Fedyk, He, and Hodson, 2024. A firm that recruited from a future AI powerhouse in 2010 got cheap AI talent later — for reasons unrelated to its own prospects. Controls for CS-strong and top-10 schools.

AI-investing firms grow — through product innovation

Effects of one SD more AI hiring over eight years

- **+19.5%** sales, **+18.1%** employment, **+22.3%** market value.
- The footprint is **product innovation**: trademarks, patents, product variety — and employment *rises*.
- **Not** process innovation: costs and productivity barely move.

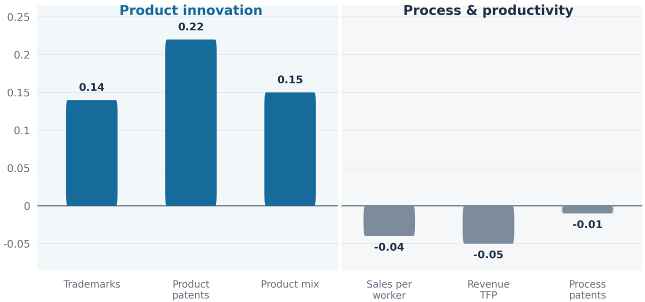

Source: Babina, Fedyk, He, and Hodson, 2024.

The gains go disproportionately to firms that were already big

Concentration and superstar dynamics rise

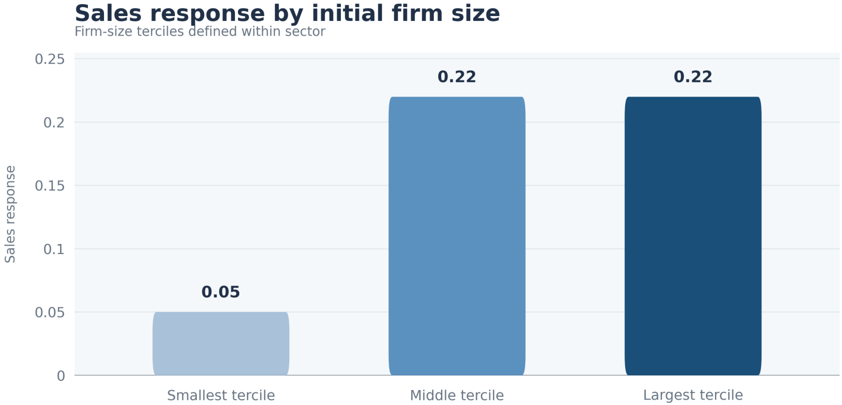

Source: Babina, Fedyk, He, and Hodson, 2024. Industries still grow (+17% sales per SD), but HHI and top-firm shares rise.

30 November 2022: the market meets ChatGPT

Eisfeldt, Schubert, Taska, and Zhang, 2026, Journal of Finance forthcoming

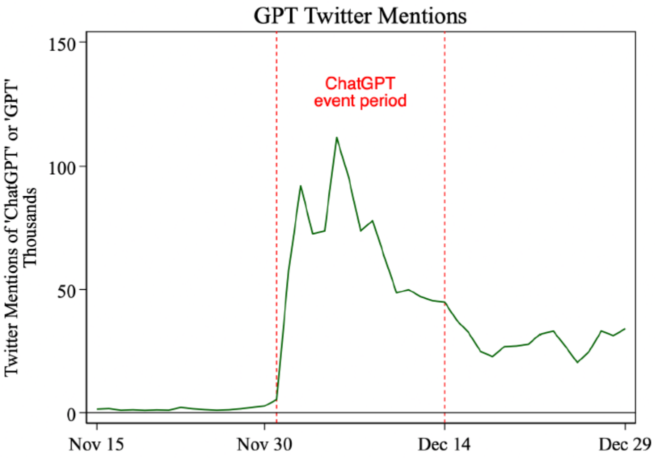

Social media attention to the release of ChatGPT.

"Artificial minus Human": exposed firms gained ∼5% in two weeks

+0.44%/day CAPM alpha; no pre-trend, no reversal; encore at GPT-4

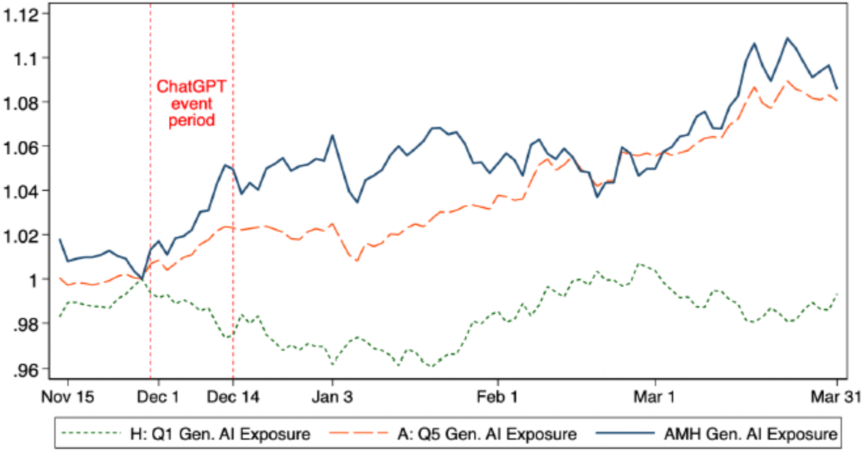

Source: Eisfeldt, Schubert, Taska, and Zhang, 2026, Figure 1. Value-weighted quintile portfolios; CARs from 29 Nov 2022.

Table of Contents

Firms that invest in AI grow — but what happens inside them?

Babina, Fedyk, He, and Hodson, 2025 — same firms, same data, new outcomes

Workers

Does AI displace highly skilled workers — or increase demand for them?

Skills

How does demand shift across education, degrees, and technical skills?

Hierarchy

Does AI reduce the need for managers and layers of approval?

Data

Same resume s and postings — now tracking *who sits where* in the organisation.

Theory: AI is cheap access to knowledge

Knowledge hierarchies (Garicano, 2000)

- Hierarchies **economise on scarce knowledge**: routine problems solved below, hard ones escalate to experts.
- AI cuts the cost of accessing knowledge ⇒ the front line solves **more itself**.
- Prediction: **wider spans of control, thinner middle** — flatter firms.

Before	→	With AI
senior		senior
middle		middle
front line		front line + AI

Rival hypotheses: IT can also *centralise* control — and AI may *automate* skilled work outright. Which wins is an empirical question.

Finding 1: AI-investing firms become flatter

Change in workforce shares per one SD more AI investment

+1.6%	−0.8%	−0.7%
junior / individual contributors	middle management	senior management

- Exactly the knowledge-hierarchy prediction: **wider spans, less management**.
- The squeezed layer is **the managers in between** — the front line *grows*.
- Empowered juniors escalate less, so each manager can oversee more people.

Finding 2: AI firms upskill

Change per one SD more AI investment

Education (resumes)

+3.7% BA   +2.9% MA   +0.6% PhD

−7.2% **workers without college**

Skills (postings)

↑ STEM degrees, data analysis, IT

↓ social science, traditional operations

- Job postings corroborate: required education rises by **0.5 years**.
- AI workers are **excluded** from these outcomes — the *rest* of the firm changes.

Timing: PhD/STEM-rich firms **adopt first**; flat structures do not predict adoption. Technical human capital is a **precondition** — flattening is a **consequence**.

Measuring exposure: an LLM reads every task in the US economy

Eisfeldt, Schubert, Taska, and Zhang, 2026 — from tasks to occupations to firms

1 — Task
GPT scores all 19,265 O*NET tasks

→

2 — Occupation
average over its tasks; 923 occupations

→

3 — Firm
weight by occupation mix; 2,518 firms

- Score = 1: ChatGPT alone halves the task time; = 0.5: needs added software; = 0: untouched.
- Telemarketer, "adjust sales scripts" = 1; manager, "review transactions" = 0.5; installer, "connect AC" = 0.
- Firm exposure varies a lot: **mean 35%, SD 8%** — that spread drives the event study.

LLM task labels validated against human labels, following Eloundou et al. (2023). Occupation mix from Revelio Labs.

Generative AI targets cognitive work

40% of white-collar tasks exposed vs. 9% of blue-collar and service jobs

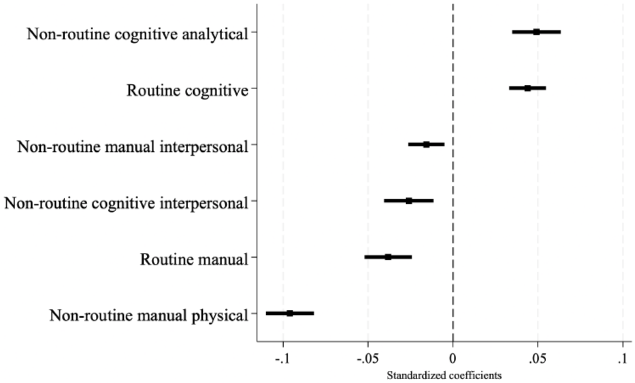

Source: Eisfeldt, Schubert, Taska, and Zhang, 2026. Loads on non-routine analytical skills; negative on manual work.

Exposure rises with pay: high earners are in the crosshairs

Most exposed: legal, computer & mathematical, business & financial operations

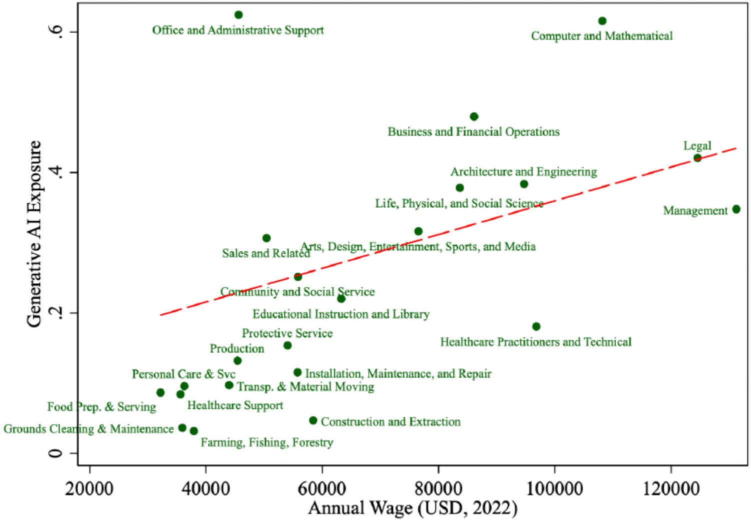

In the field, AI delivers large productivity gains

Randomised and natural experiments, 2023–2025

+15%
customer-service productivity

−40%
writing time, with *higher* quality

+26–56%
software task completion

- **Novices gain most:** +30% for the least-experienced agents; same for taxi drivers and junior developers.
- **The tool matters:** *agentic* coding AI delivers its largest gains to experienced workers.
- **Matching improves:** AI-led interviews raise offers +12%, starts +18%, retention +16–18%.

Sources: Brynjolfsson, Li, and Raymond, 2025 (customer service); Noy and Zhang, 2023 (writing); Peng, Kalliamvakou, Cihon, and Demirer, 2023; Cui et al., 2024 (software); Kanazawa, Kawaguchi, Shigeoka, and Watanabe, 2022 (taxi drivers); Sarkar, 2025 (agentic coding); Jabarian and Henkel, 2025 (interviews).

Finance is not exempt: two cautionary tales

Field evidence from insurance and accounting

Insurance sales (Liu, 2024)

RCT with 11,000 agents: AI demand-prediction **raises sales**…

…but **adverse selection worsens** — agents stop doing their own risk assessment.

Accounting (Choi and Xie, 2026)

AI-enabled accountants serve **55% more clients**; data entry falls 3.5 h/week.

Faster monthly closes, **better** reporting quality — time shifts to judgment.

Same lesson twice: AI shifts **what the human does**.

Audit: the cleanest laboratory

Fedyk, Hodson, Khimich, and Fedyk, 2022, Review of Accounting Studies

Why audit?

- One product, 36 firms
- Quality: restatements
- Price: audit fees
- Highly AI-exposed

Quality — better audits?

Price — cheaper audits?

Labour — who does the work?

Again, measure AI from hiring: audit AI teams grew fivefold

310,000+ resume s; AI share 0.08% (2010) → 0.39% (2018)

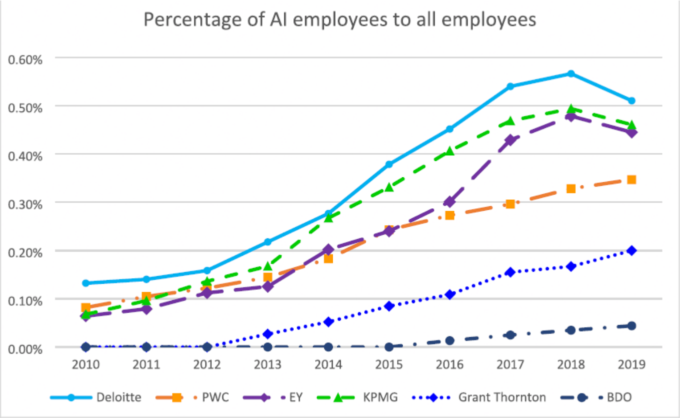

Fig. 1 AI Employees in the Six Largest Audit Firms

Source: Fedyk, Hodson, Khimich, and Fedyk, 2022

One SD more AI: better audits, cheaper audits, fewer juniors

Quality ↑ (fewer restatements)   Fees ↓   Juniors ↓ (junior-only, slow)

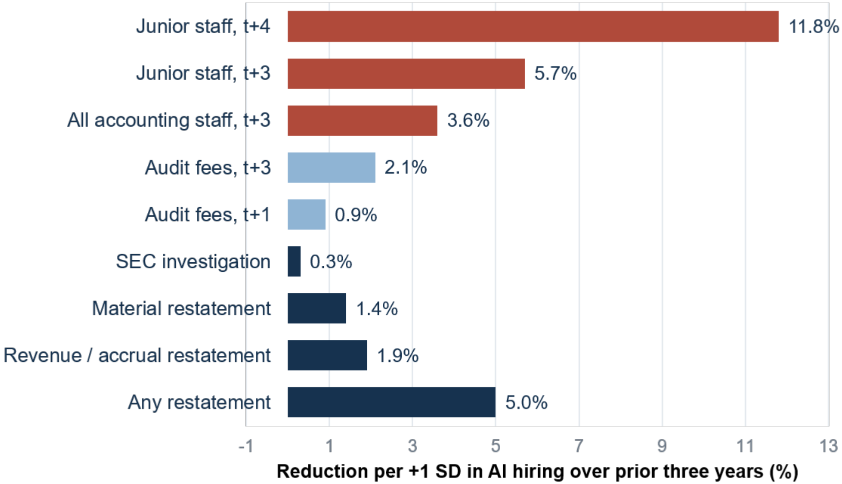

Reading the audit results

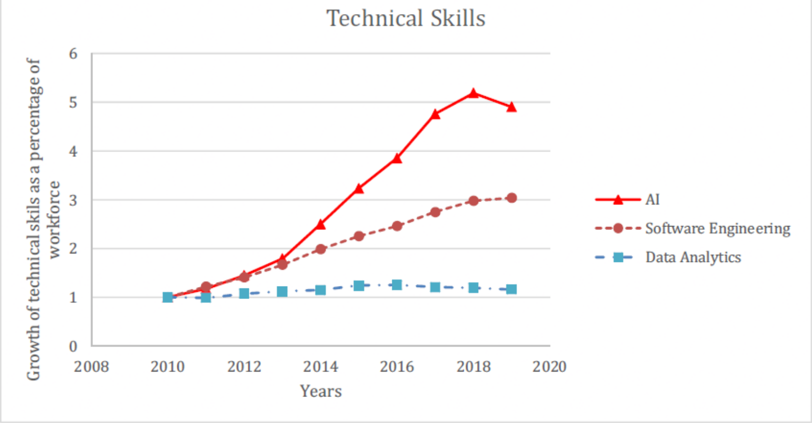

Fig. 3 Graph of Growth in Technical Skills

If AI does the junior work, where do the next partners come from?

Table of Contents

The industrialisation of knowledge work?

The analogy is precise enough to measure

Industrial revolution

Machines let each worker command **more capital** — output per worker soared.

AI revolution?

Machine learning lets each knowledge worker command **more data**.

- Then, income shifted from labour to capital owners — now to **data owners**?
- Incremental upgrade or transformation? The question is quantitative: **one exponent**.

Source: Abis and Veldkamp, 2024, Review of Financial Studies.

Economists summarise a revolution as a change in one exponent

Production functions

$$Y = A K^{\alpha} L^{1-\alpha}$$

- α does double duty: **curvature of returns *and* capital's income share**.
- Higher $\alpha \Rightarrow$ weaker diminishing returns $\Rightarrow$ more capital per worker.
- And mechanically, a **smaller labour share** — it is exactly $1 - \alpha$.

Historians put industrialisation at α up **5–20pp** — our yardstick for AI.

Historical estimates by Clark, Allen, and Klein–Kosobud, as compiled in Abis and Veldkamp, 2024.

Knowledge production in investment management 2010-2018

Two technologies, one data stock — estimated from job postings and wages

raw data → **data managers** build stock D → **analysts** L use D → knowledge

$$K^{\mathrm{AI}} = A\, D^{\alpha} L^{1-\alpha} \qquad K^{\mathrm{OT}} = A\, D^{\gamma} L^{1-\gamma}$$

- The whole question: **is $\alpha > \gamma$?** Does ML weaken diminishing returns to data?
- Nothing is directly observed — but under Cobb–Douglas, **wage bills reveal the exponents**.

Source: Abis and Veldkamp, 2024. Investment management, 2010–2018; AI analysts out-earn traditional ones by $26,333. Data accumulates like capital: $D_{t+1} = (1-\delta)D_t +$ data managers' output.

Diminishing returns to data are weakening

Firms substitute data for people — at the low end of industrial-revolution estimates

$0.82 \rightarrow 0.87$

data exponent, OT vs. ML

-5%

of knowledge income shifts to data owners

$810bn

one sector's data stock (2018)

- Each extra dataset stays useful longer $\Rightarrow$ optimal **data per analyst rises**.
- The $+0.05$ is **conservative**: lower depreciation δ implies $+0.17$.
- A falling labour *share* is not falling employment: **headcounts grew**.

Source: Abis and Veldkamp, 2024, Table 2 and Figure 6. Standard errors 0.001–0.002; all differences significant.

Accounting expenses the data; economics capitalises it

Why your ratios are quietly wrong

Accounting view

A data engineer's salary = **expense**. Nothing reaches the balance sheet.

Economic view

The same salary = **investment** in a durable, depreciating stock.

Every ratio built on book values inherits the gap:

- **Tobin's q** — numerator prices the data, denominator omits it.
- **ROA** — understated asset base, spuriously high returns.
- **Market-to-book** — rises with no change in profitability.
- **Measured productivity** — investment booked as cost, growth understated.

On intangibles and measurement more broadly: Crouzet, Eberly, Eisfeldt, and Papanikolaou, 2022.

Data is unlike machines: non-rival, networked across firms

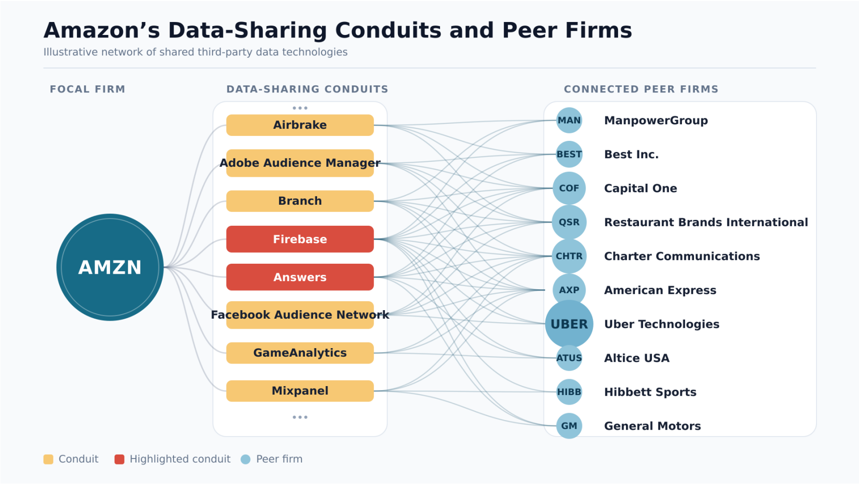

Amazon's data-sharing conduits and peers (Bian, Huang, Li, and Tang, 2024).

If data is the asset, textbook corporate finance changes

Three places where the standard answers move

Capital structure

Hard to verify, firm-specific, barely redeployable $\Rightarrow$ poor collateral $\Rightarrow$ **equity and cash**.

M&A

Data procurement: build (hire) vs. buy (acquire). Non-rivalry creates real synergy — if the value survives the deal.

Firm size

Weaker diminishing returns + non-rivalry $\Rightarrow$ **larger firms** — the superstar result again.

Underneath it all: data is generated by customers, built by employees, financed by shareholders — who *should* own it?

So far AI has been a copilot. What happens when it works alone?

Extending the knowledge hierarchy: AI that holds tacit knowledge and scales on compute

Copilot (non-autonomous)

Can only **advise** a human, who stays in the loop — it *assists* rather than competes.

Coworker (autonomous)

Performs tasks independently — effectively a new worker, *competing* for work.

- Everything in Part 2 — call centres, writing, coding — is **copilot** evidence.
- The theory asks what changes when AI graduates to **coworker**.

Source: Ide and Talamàs, 2025, Journal of Political Economy.

Who wins depends on how good and how autonomous AI is

Equilibrium predictions, not estimates

- **Basic autonomous AI** replaces routine work: **experts gain**, the least knowledgeable can lose.
- **Advanced autonomous AI** solves hard problems for everyone: **both ends gain** — the middle is squeezed.
- **Copilots** favour novices (they assist, not compete); **autonomy** favours experts — and raises **total output**.

The policy trade-off: restricting AI autonomy may compress inequality — **at the price of aggregate output**.

Source: Ide and Talamàs, 2025. Copilot results match Part 2's field evidence; the middle-squeeze echoes the flattening in Babina, Fedyk, He, and Hodson, 2025.

What you should take away

Five answers to the five opening questions

- **1 — Value:** $\sim$20% growth per SD of AI hiring, via product innovation.
- **2 — Evidence:** measure AI from hiring; instrument, because adopters differ.
- **3 — Organisation:** flatter, more skilled firms; juniors hit in audit.
- **4 — Distribution:** big firms, clients, and data owners win.
- **5 — Data:** the exponent moved; the asset is real — and unrecorded.

References I

- Abis, Simona and Laura Veldkamp (2024). "The Changing Economics of Knowledge Production". In: *The Review of Financial Studies* 37.1, pp. 89–118.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson (2024). "Artificial Intelligence, Firm Growth, and Product Innovation". In: *Journal of Financial Economics* 151, p. 103745.
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson (2025). "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition". In: *Technology, Productivity, and Economic Growth*. Ed. by Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner. University of Chicago Press, pp. 75–118.
- Bian, Bo, Qiushi Huang, Ye Li, and Huan Tang (2024). *Data as a Networked Asset*. SSRN Working Paper 5183829. Revise and resubmit, Journal of Political Economy.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025). "Generative AI at Work". In: *The Quarterly Journal of Economics* 140.2.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2021). "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies". In: *American Economic Journal: Macroeconomics* 13.1, pp. 333–372.

References II

- Choi, Jung Ho and Chloe Xie (2026). "Human + AI in Accounting: Early Evidence from the Field". In: *Journal of Accounting Research*. Forthcoming.
- Crouzet, Nicolas, Janice C. Eberly, Andrea L. Eisfeldt, and Dimitris Papanikolaou (2022). "The Economics of Intangible Capital". In: *Journal of Economic Perspectives* 36.3, pp. 29–52.
- Cui, Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz (2024). *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*. Working paper, MIT and Microsoft.
- Eisfeldt, Andrea L., Gregor Schubert, Bledi Taska, and Miao Ben Zhang (2026). *Generative AI and Firm Values*. Journal of Finance, forthcoming. NBER Working Paper 31222.
- Falato, Antonio, Dalida Kadyrzhanova, Jae Sim, and Roberto Steri (2022). "Rising Intangible Capital, Shrinking Debt Capacity, and the US Corporate Savings Glut". In: *The Journal of Finance* 77.5, pp. 2799–2852.
- Fedyk, Anastassia, James Hodson, Natalya Khimich, and Tatiana Fedyk (2022). "Is Artificial Intelligence Improving the Audit Process?" In: *Review of Accounting Studies* 27.3, pp. 938–985.

References III

- Garicano, Luis (2000). "Hierarchies and the Organization of Knowledge in Production". In: *Journal of Political Economy* 108.5, pp. 874–904.
- Ide, Enrique and Eduard Talamàs (2025). "Artificial Intelligence in the Knowledge Economy". In: *Journal of Political Economy* 133.12, pp. 3762–3800.
- Jabarian, Brian and Luca Henkel (2025). *Voice AI in Firms: A Natural Field Experiment on Automated Job Interviews*. SSRN Working Paper 5395709.
- Jones, Charles I. and Christopher Tonetti (2020). "Nonrivalry and the Economics of Data". In: *American Economic Review* 110.9, pp. 2819–2858.
- Kanazawa, Kyogo, Daiji Kawaguchi, Hitoshi Shigeoka, and Yasutora Watanabe (2022). *AI, Skill, and Productivity: The Case of Taxi Drivers*. NBER Working Paper 30612.
- Liu, Xing (2024). *Artificial Intelligence and Information Production in Selection Markets: Evidence from Insurance*. Working paper.
- Noy, Shakked and Whitney Zhang (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence". In: *Science* 381.6654, pp. 187–192.

References IV

- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer (2023). *The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.* arXiv:2302.06590.
- Sarkar, Suproteem K. (2025). *AI Agents and Higher-Order Work.* Working paper, University of Chicago.
- Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang (2026). *Firm Data on AI.* NBER Working Paper 34836. Four-country firm survey: US, UK, Germany, Australia.