Too Much of A(I) Good Thing?

Daron Acemoglu Nick Bloom Tessa Hall Melvin Liam John Van Reenen Zach Zaslavsky

LSE MIT Stanford

November 2026

Motivation: AI adoption is endogenous, so its effects are unknown

- Firms that adopt AI are a selected sample: larger, more productive, growing faster before adoption (McElheran et al., 2024)
- Many firm-level estimates of the effect of AI put that endogenous adoption on the right-hand side
  - Job postings and LinkedIn hiring (Goldfarb et al., 2023; Babina et al., 2024), surveys (Bick et al., 2024), occupational exposure (Eisfeldt et al., 2023; Humlum and Vestergaard, 2025)
  - A positive coefficient tells us which firms adopt, not what adoption does
- Timing is endogenous too: firms invest in AI when they expect it to pay
- So the open question is: does a firm's growth, cost, or value change when it invests in AI? And can firms over-invest, following the hype rather than the return?

This paper: finds a source of variation in AI investment that is not the firm's own decision, and follows outcomes for three years.

Two CEOs who died in office in 2021

Low-AI board



Arne Sorenson, CEO since 2012
Died Feb 2021

- No director posted about AI on LinkedIn before 2021
- Successor named eight days later: Tony Capuano, a 25-year Marriott insider who ran hotel development and had never posted about AI

High-AI board



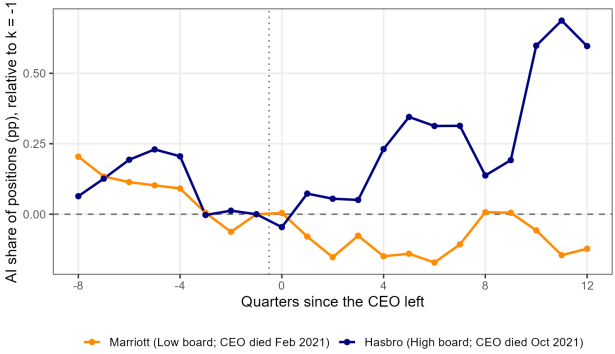
Brian Goldner, CEO since 2008
Died Oct 2021

- One AI poster on the board: independent director Rich Stoddart (8 AI posts, 2016–17)
- Stoddart becomes interim CEO, then chair; the board picks Chris Cocks, ex-Microsoft, head of Wizards of the Coast and digital gaming, who calls himself an "AI bull" on LinkedIn

Both CEOs died of illness in the same year. Neither board chose the timing.

Photos: Sorenson, World Travel & Tourism Council; Goldner, Hasbro/Business Wire. Departures: CEO Dismissal Database.

What happened next: Hasbro hires AI workers, Marriott does not



AI share of the firm's workforce (we define it later), each firm minus the mean of the other firms in its NAICS3 industry in the same quarter, relative to the quarter before the CEO left. Percentage points.

Hasbro: AI share up 0.6 pp relative to its industry by year three

Marriott: AI share flat relative to its industry

What we find

Design: Exogenous CEO departures (death, illness, retirement) as moments when the board's AI disposition can bind. Compare, within departure year and quarter and net of industry trends, firms whose boards posted about AI before the departure with firms whose boards did not.

Results, 122 AI-engaged vs 588 other boards with exogenous CEO turnover, three years after the departure

- Firms with AI-engaged boards have AI-engaged CEOs both before and after the departure ($+10$ and $+11$ pp, from a 2% base), and the firm's AI share of employment rises ($+0.04$ to $+0.29$ pp, significant on the cosine measures, not on job titles)
- **Employment:** no significant change (-1 to -4 log points, n.s. in both designs); junior share does not fall
- **No outcome improves.** Sales and market value point down; sales gains above 6% are ruled out from the second year on.

What we build to get there

1. **Leadership AI enthusiasm.** 2.7m LinkedIn posts by directors, executives, and blockholders of 2,796 US listed firms, classified as AI-related or not.
2. **Firm AI investment.** Share of the workforce in AI roles, from Revelio employment histories, quarterly, 2010Q1–2025Q4
   - Baseline: AI keywords in job titles (as in Babina et al. (2024))
   - Alternative: cosine similarity of a worker's time-stamped skills to workers at OpenAI, Anthropic, DeepMind
3. **Exogenous CEO departures.** Death, illness, retirement from the CEO Dismissal Database (Gentry et al., 2021), matched to BoardEx spells and Compustat outcomes, 2010Q1–2026Q2

Data 1: LinkedIn posts of corporate leaders

Universe and matching

- Directors, executives, individual blockholders of 2,796 US listed firms (ExecuComp, blockholder filings, DEF 14A proxies): 120,028 person–firm rows
- Web search on name and firm → LinkedIn URL for 74% of rows, 68,819 distinct profiles
- URLs cross-checked against Revelio: name and employer confirmed for 30,122 profiles (*strong match*)
- Strong-match rate, of rows searched: 50% executives, 36% directors, 16% blockholders

Corpus

- Every post on matched profiles, scraped via Apify: 2.7m posts, 29k profiles, 2013 – Oct 2026
- 57% original posts, 25% reposts, 18% quotes
- AI-related if text contains one of 24 AI/ML keywords (*AI, LLM, generative AI, ChatGPT, copilot, machine learning*, . . .): 6% of posts; over a third of active profiles have one

  ▸ Post keywords
- Sentiment of each post: Loughran–McDonald net sentiment
- Aggregated to firm-quarter: AI posts by directors, AI post share

What the posts look like: back to our example

Tony Capuano, Marriott CEO

Anthony Capuano • 3rd+
President and Chief Executive Officer at Marriott ...
Jun 16, 2025
+ Follow ···

Cannes Lions International Festival of Creativity is where the world's most powerful creative conversations happen - on stage, over coffee, and in every meeting room along the Croisette. It's an essential moment to tap into the energy of global culture, explore how brands are shaping behavior, and build partnerships that help us connect more meaningfully with travelers around the world.

I'm looking forward to inspiring conversations, new experiences, and exploring how travel shapes us all.

Rich Stoddart, Hasbro director

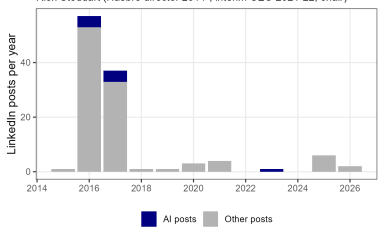
Rich Stoddart • 3rd+
Independent Director | Executive Coach | Trusted Advisor
Feb 21, 2017
+ Follow ···

With the increasing integration of artificial intelligence into the workplace, it has become vital to rethink our purpose and functions.

The Rise of AI Makes Emotional Intelligence More Important
hbr.org

16 1 6

Rich Stoddart (Hasbro director 2014-, interim CEO 2021-22, chair)



Chris Cocks, Hasbro CEO

Chris Cocks • 3rd+
Chief Executive Officer at Hasbro
18h
+ Follow ···

I had a great conversation with Liz and Rohan on Semafor's Compound Interest podcast about IP, play, AI and what comes next for Hasbro.

AI is on a lot of our minds. How do we use it responsibly? How can it improve an experience without losing its soul? And how do we make the economics work?

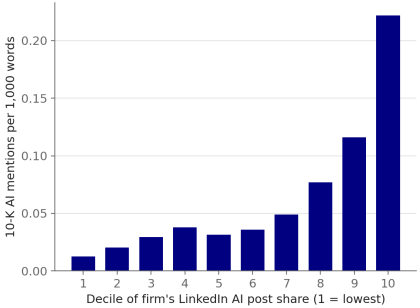
My take: AI is a powerful tool to supercharge our teams and unlock entirely new experiences for fans. The speed at which we can innovate and improve is remarkable, and I believe it will soon become the standard consumers expect. And the possibilities for fan-generated content, especially when paired with technologies like 3D and digital printing, are going to be incredible. We're just scratching the surface.

The posting measure tracks attention to AI

Share of leaders' own LinkedIn posts about AI, by month



10-K AI mentions by decile of a firm's AI post share



- AI share of leaders' own posts about 2% until ChatGPT (Nov 2022), then $\times 5$ on average since 2023 and $\times 8$ in the last year (12% and 18%)
- Validation: firms in the top decile of the AI post share mention AI in their 10-Ks 18$\times$ as often as the bottom decile (2,118 firms; rank correlation 0.35) ▸ Example posts

Data 2: firm AI investment from Revelio employment histories

Revelio

- Public LinkedIn profiles: positions with employer, dates, titles, SOC code, seniority; reported skills
- We use Revelio's sampling weights to scale profiles to employment; coverage $\approx$ half of US employment and rising
- Skills count for a position only if reported by its end date (+30 days): limits retrospective relabelling of past jobs (Bloom et al., 2026)
  - 30-day grace period given because skills used at a position may only be added after the position has completed and when applying for other jobs
- Workers assigned to listed firms through a timed ownership chain (CRSP, Exhibit 21)

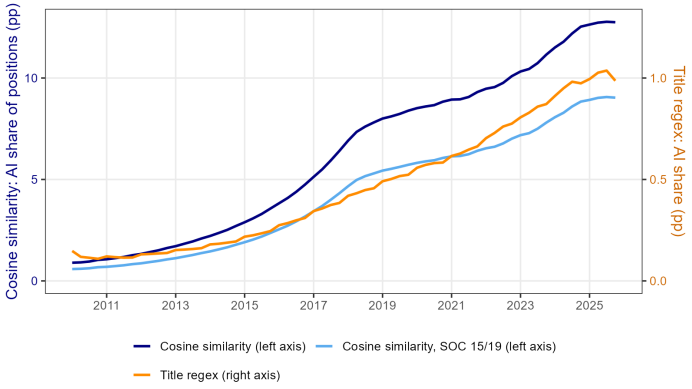
▸ Firms hiring AI workers are a selected sample

AI share of the workforce, following Babina et al. (2024): share of positions that are AI positions

1. **Title regex (baseline).** Title contains *artificial intelligence, machine learning, deep learning, NLP, computer vision* (and abbreviations). Narrow: 1% of positions in 2025. ▸ Title terms
2. **Skills similarity to AI labs (alternative).**
   - Every position is a binary vector over Revelio's skill vocabulary, using only skills reported by 30 days after the position's end date
   - Reference vector: mean skill vector of 16,008 positions at OpenAI, Anthropic, DeepMind. Alternative: 10,405 positions in SOC 15 or 19 (computing, science) at the AI labs.
   - Firm AI share = employment-weighted average cosine similarity: 0.12 in 2025 (0.085 for SOC 15/19).

▸ Correlation across measures

The three measures over time



- Our event-study panel: 2,140 firms, 2010Q1–2025Q4; AI positions over all positions, by quarter
- Title regex 0.1% $\rightarrow$ 1.0%
- Cosine similarity rises earliest and fastest, 2015–18, and keeps rising
- All three accelerate after 2022
- Before 2022 these may pick up programming and ML rather than AI as we now mean it. Departures from 2022 on give noisier estimates in the same direction; too few years for an AI-only sample, so read results as a mixture of AI and earlier ML ▸ 2022 onwards

Data 3: Firm outcomes and CEO departures

Outcomes (firm-quarter panel, 2010Q1–2026Q2, 2,140 firms)

- Compustat quarterly, to fiscal 2026Q2: log sales, log COGS, log market value; Compustat log employees, annual to 2025
- Revelio: junior share = weighted share of positions at the three lowest of seven seniority levels

CEO departures

- CEO spells from BoardEx: 6,742 exits 2010–2025
- Matched to the CEO Dismissal Database (Gentry et al., 2021), which codes the reason
- Keep **death, illness, retirement** (exogenous); drop dismissals, new opportunities, other
- Among matched exits: 50% retirement, 22% dismissal, 2% illness, 1% death

Event Study Design

The problem. Boards that post about AI sit at firms already moving into AI, in AI-heavy industries, in later years.

The idea, from Acemoglu et al. (2022), a board acts on its preferences most visibly when it chooses and briefs a new CEO. If the departure is unrelated to AI strategy, the board's AI disposition at that moment is a pre-determined trait.

What makes the comparison credible

1. **Exogenous departure.** Death, illness, retirement only. The board's posts were written before the CEO leaves.
2. **Departing firms only.** Every firm in the regression has a CEO leaving for the same kind of reason, so whatever an unexpected departure does on its own (sales dip, mark-down) is differenced out.
3. **Within departure year and quarter, net of industry trends.** A high-exposure departure in 2021 is compared only with low-exposure departures in 2021, after removing each NAICS3 industry's own quarterly path.

Exposure. $\mathrm{High}_i = 1$ if a director seated at the departure posted about AI before it. Directors are validated LinkedIn profiles seated through a BoardEx spell or a DEF 14A listing. 122 High vs 591 Low exogenous episodes.

▸ Falsification: does the board move AI hiring before the CEO leaves?

Sample: departure episodes by year

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	Total
Low board, exogenous	60	80	85	55	60	68	52	48	45	25	12	1	591
High board, exogenous	1	0	1	1	3	10	11	23	24	20	22	6	122
Low board, all departures	158	158	170	143	136	139	139	131	120	93	84	70	1,541
High board, all departures	1	0	2	4	9	20	32	54	61	68	105	116	472
No CEO departure (reference firms)													327

- 87% of High exogenous episodes are in 2020 or later: the comparison rests on the last five years
- Within a departure year, High and Low boards are compared quarter by quarter
- A firm enters at its first exogenous exit; a later exit starts a new episode only after 16 quarters

▸ All CEO departures

Mechanism: do AI-engaged boards appoint AI-engaged CEOs?

	Departing	Successor	Succ. − Dep.
High board	0.095**	0.106**	0.011
	(0.038)	(0.045)	(0.040)
Low-board mean	0.016	0.022	
Episodes	418	418	418

Outcome: 1 if the CEO posted about AI before taking office. Exogenous episodes with a successor within 8 quarters. NAICS3 and departure-year effects; SE clustered by firm.

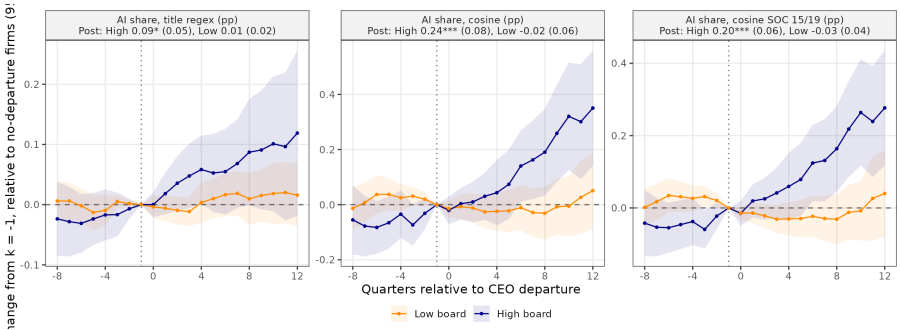
- The departing CEO is already 10 pp more likely to be an AI poster at High boards
- The successor is 11 pp more likely to be an AI poster, from a Low base of 2%
- Successor minus departing is 0.011 (s.e. 0.040): no evidence that the board shifts the CEO's AI engagement when it gets to choose

High vs Low AI boards: AI hiring

Episode-quarters $k = -8, \ldots, 12$ plus the 1,198 firms with no exogenous departure, held at $k = -1$:

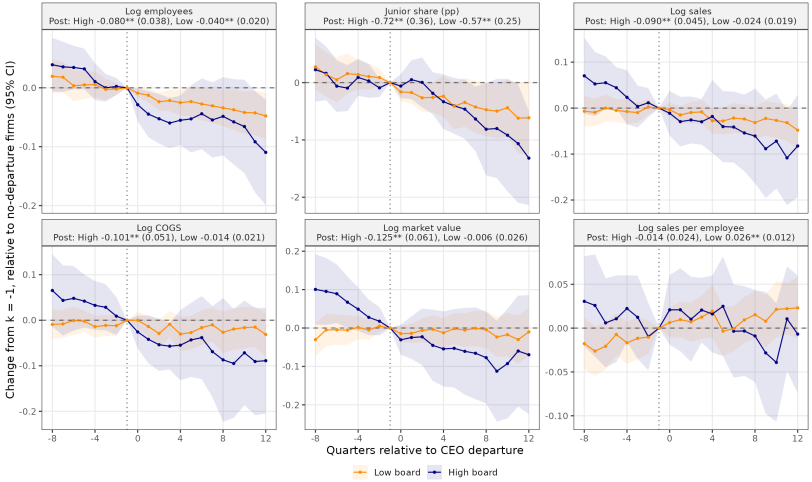
$$y_{it} = \alpha_i + \eta_{n(i),t} + \sum_{k \neq -1} \beta_k \, \mathbf{1}[k_{it} = k] + \sum_{k \neq -1} \delta_k \, \mathbf{1}[k_{it} = k] \times \mathrm{High}_i + \varepsilon_{it}$$

α_i episode effects; η NAICS3 $\times$ calendar quarter, identified off the no-departure firms. β_k is the Low-board path, $\beta_k + \delta_k$ the High-board path. SE clustered by firm.



122 High / 588 Low episodes. Strip: pooled Post ($k \geq 4$) coefficient for each path. Shares in pp. Low boards: no change in AI hiring after a departure. High boards: AI share rises on every measure (+0.09 to +0.24 pp).

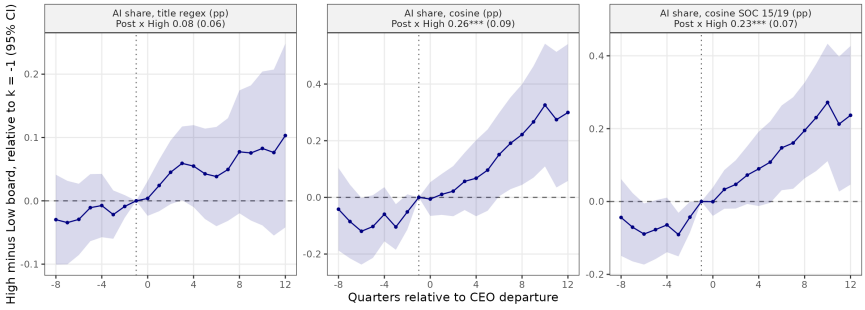
High vs Low AI boards: other outcomes



Same regression and sample. Low boards after a departure: employment -4 log points**, junior share -0.6 pp**, sales -2 log points. High boards: employment -8 log points**, sales -9^{**}, market value -12^{**}. Sales per employee -1.4 (High) and $+2.6$ (Low) log points.

High minus Low: AI hiring

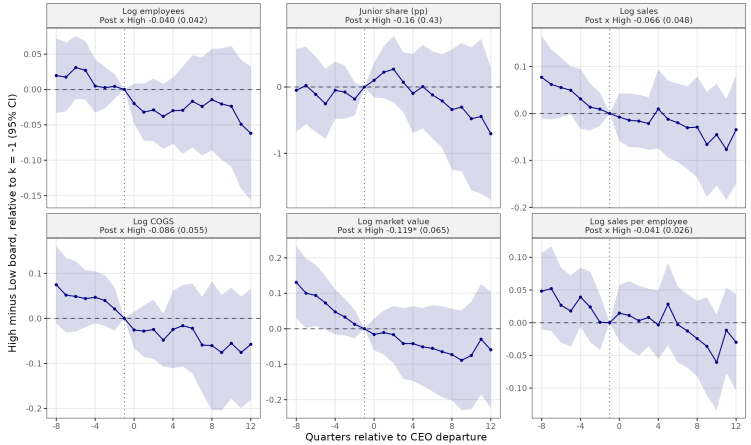
δ_k from the same regression: the gap between the two lines. Under each panel: the pooled coefficient on $\text{Post}_{it} \times \text{High}_i$, $\text{Post} = \mathbf{1}[k \geq 4]$ with a separate transition term for $0 \leq k \leq 3$ (SE), the same regression as the paths slide. SE clustered by firm.



Pooled Post × High significant at 1% on both cosine measures, n.s. on title regex. Joint test of all post coefficients ($k = 0, \ldots, 12$): n.s. on all three measures. Joint pre-trend test rejects at 10% on 2 of 3 measures.

▸ Pooled effects table ▸ Industry controls ▸ Continuous exposure

High minus Low: other outcomes



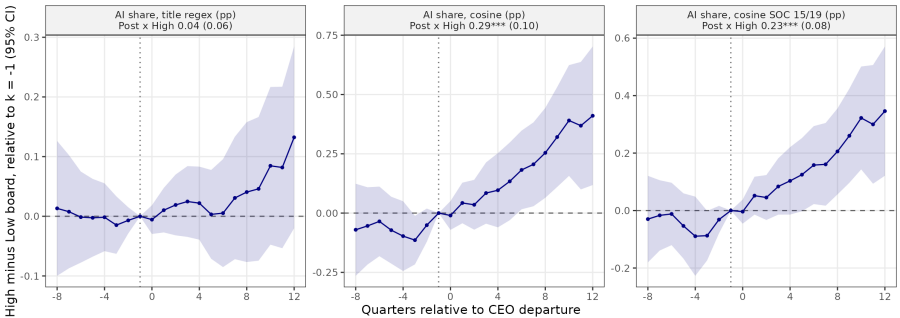
Employment −4.0 log points (pooled, n.s.); sales −6.6 and market value −11.9* log points (joint test of the post coefficients: sales n.s., market value n.s.); junior share −0.2 pp; COGS −8.6 and sales per employee −4.1 log points. Sales, COGS, market value and sales per employee were already sloping down before the departure.

Refinement: departing firms only, within departure year

Drop the reference firms; add departure-year cohort × quarter effects $\mu_{c(i),t}$ in place of the common path β_k:

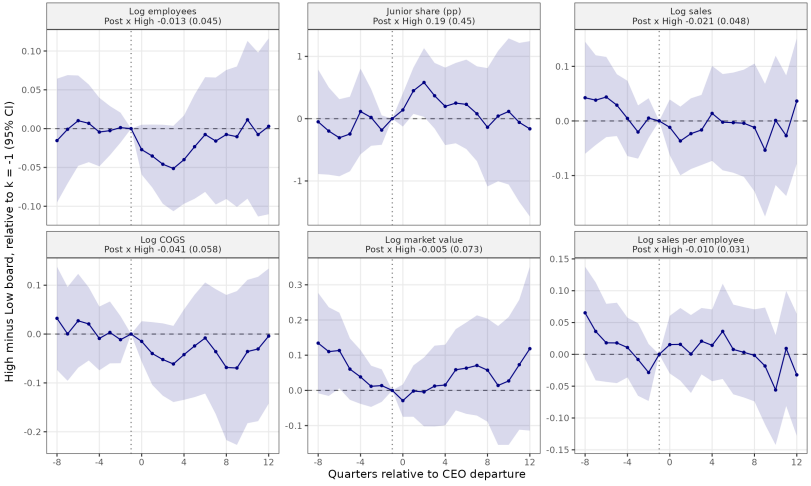
$$y_{it} = \alpha_i + \eta_{n(i),t} + \mu_{c(i),t} + \sum_{k \neq -1} \delta_k \, \mathbf{1}[k_{it} = k] \times \mathrm{High}_i + \varepsilon_{it}$$

A High departure in 2021 is compared only with Low departures in 2021, net of industry × quarter. Removes any difference in what a departure does across years; costs power (Low boards are the only reference).



Same statistics under each panel as before. AI share +0.04 pp (title regex), +0.29 (cosine), +0.23 (SOC 15/19 cosine); joint test of all post coefficients: title regex n.s., cosine n.s., SOC 15/19 cosine 10%; joint pre-trend test does not reject on any measure. First-stage F for the cosine share 5.5.

Refinement: other outcomes



Same regression and statistics. Employment −1.3 log points (pooled, n.s.) with no pre-trend; junior share 0.2 pp, sales −2.1, COGS −4.1, market value −0.5 log points; sales per employee −1.0. No outcome is significant; market value slopes down before the departure.

Conclusions

- Two new measures: leadership AI enthusiasm from 2.7m LinkedIn posts; firm AI workforce from job titles and from skills similarity to frontier-lab workers
- One design: exogenous CEO departures let the board's AI disposition bind at a moment it did not choose
- Firms with AI-engaged boards have AI-engaged CEOs before and after the departure; AI hiring rises on the cosine measures, not on job titles
- On the real side no outcome improves: no effect is significant in the refinement, and one-sided bounds rule out gains above about 6% in sales and 12% in market value; productivity per worker is flat
- The evidence is consistent with over-investment or slow payoffs, and shows no sign of AI-led growth within three years; pre-trends on the real outcomes and small samples limit what we can say

Appendix: firms with at least one AI worker are a selected sample

Sentiment-sample firms, $\geq$ 50 workers, 2025	AI Using	Other
Firms	1,116	800
Mean workers	8,044	907
Median workers	2,698	438
Mean predicted wage ($)	117,916	104,975
Share graduate degree at hire	0.18	0.14
Share STEM at hire	0.26	0.18

AI-hiring = at least one worker with an AI role title (title regex). 58% of firms. Workers are each person's latest position running into 2025. Firms are Compustat gvkeys in the sentiment sample.

Back

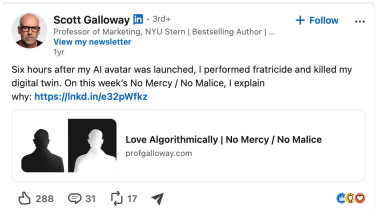
Appendix: AI keywords in LinkedIn posts

A post is an AI post if its text contains any of these 24 keywords.

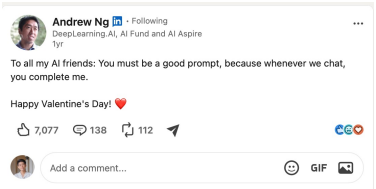
Core AI	artificial intelligence, ai, generative ai, gen ai, llm, large language model
Products	chatgpt, gpt, claude, gemini, copilot
Data Science/ Machine Learning	data science, predictive analytics, automation, algorithm, machine learning, ml, deep learning, neural network
Applications	nlp, natural language processing, computer vision, robotics, autonomous

Back

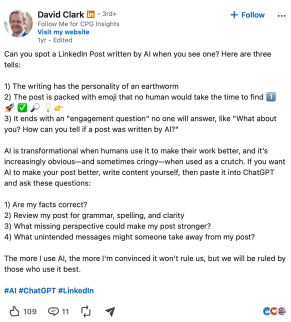
Appendix: entertaining AI posts



Scott Galloway, New York Times Director, 10 Oct 2025



Andrew Ng, Amazon Blockholder, 14 Feb 2025



David Clark, O-I Glass Director, 16 Mar 2025

Own AI posts by directors, executives and blockholders of listed firms in our sample; names confirmed against our search records.

Back

Appendix: AI terms in the title regex

A position is an AI position if its job title contains any of these terms.

- AI
- A.I.
- ML
- M.L.
- LLM
- artificial intelligence
- machine learning
- deep learn
- neural network
- large language model
- foundation model
- natural language processing
- NLP
- computer vision
- reinforcement learning
- cognitive engineer
- generative AI
- GenAI
- chatbot
- conversational AI
- edge AI
- autonomous system
- research data scientist
- quantitative data scientist
- data science architect
- MLOps
- ML operations
- ML ops
- AI engineer
- AI developer
- AI researcher
- AI architect
- algorithm developer
- AI product
- AI strategy
- AI solutions

Back

Appendix: correlation across AI workforce measures

Firm-year correlation	Title regex	Cosine (pooled)	Cosine (SOC 15/19)
Title regex	1.00		
Cosine (pooled 3 labs)	0.25	1.00	
Cosine (SOC 15/19 reference)	0.31	0.97	1.00

Pearson correlations across 33,785 firm-year observations, 2010–2025. All three measures are observed for the same firm and year.

- The keyword and cosine measures are only weakly related at the firm level
- The two cosine variants are highly correlated

Back

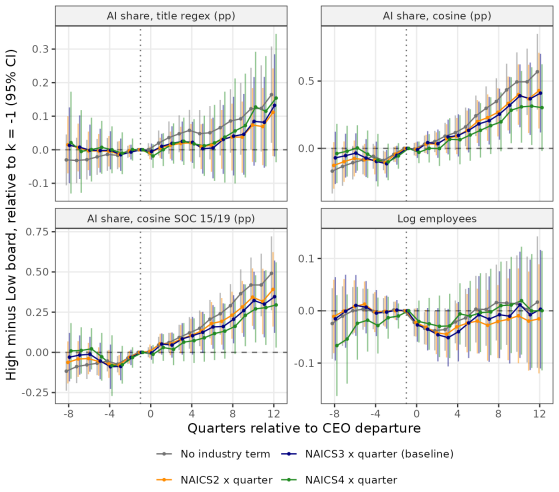
Appendix: pooled effects, Post × High from year two

	Calendar reference			Refinement: within cohort, departing firms only
	High path	Low path	High − Low	
AI share (title regex), pp	0.095^{*}	0.015	0.080	0.040
AI share (cosine), pp	0.237^{***}	-0.025	0.262^{***}	0.289^{***}
AI share (cosine, SOC 15/19), pp	0.197^{***}	-0.032	0.229^{***}	0.234^{***}
Log employees (Compustat)	-0.080^{**}	-0.040^{**}	-0.040	-0.013
Junior share, pp	-0.722^{**}	-0.566^{**}	-0.156	0.186
Log sales	-0.090^{**}	-0.024	-0.066	-0.021
Log COGS	-0.101^{**}	-0.014	-0.086	-0.041
Log market value	-0.125^{**}	-0.006	-0.119^{*}	-0.005
Log sales per employee	-0.014	0.026^{**}	-0.041	-0.010
Episodes (High / Low)	122 / 588	122 / 588	122 / 588	122 / 588
Reference firms	1,198	1,198	1,198	0
First-stage F, cosine share			1.6	5.5

Calendar-reference columns from one regression with Trans, Post, Trans × High and Post × High: Low path = Post, High path = Post + Post × High, High − Low = Post × High (exactly the difference). Refinement: Post × High with cohort × quarter effects. Transition $0 \leq k \leq 3$ has its own term. SE clustered by firm; *, **, ***: 10, 5, 1%. Employment and junior share are significant only in the two paths; sales, COGS and market value only in the High path (market value also at 10% in the gap); sales per employee only in the Low path. No real outcome is significant in the refinement.

Back

Appendix: industry controls

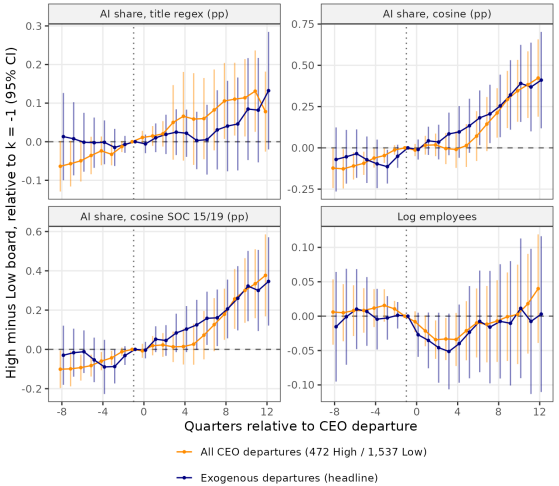


Refinement specification with the industry term varied. Pooled cosine effect: none 0.40***, NAICS2 0.32***, NAICS3 0.29***, NAICS4 0.22**. First-stage F (cosine): none 15.1, NAICS2 7.3, NAICS3 5.5, NAICS4 1.9.

- AI-share pre-period unchanged by the industry term; the title-regex path is nearly the same under every term
- Employment gap: 0.01 without an industry term, −0.02 to 0.02 with

Back

Appendix: all CEO departures



Refinement specification on every CEO departure, 472 High / 1,537 Low episodes, against the exogenous sample.

- 2.8 times the sample, tighter bands; AI-share effects survive (0.11–0.23 pp)
- The joint pre-trend test does not reject on the cosine shares ($p = 0.19$)
- Real contrasts do not shrink (employment -0.015, s.e. 0.028)

Back

Appendix: continuous board exposure

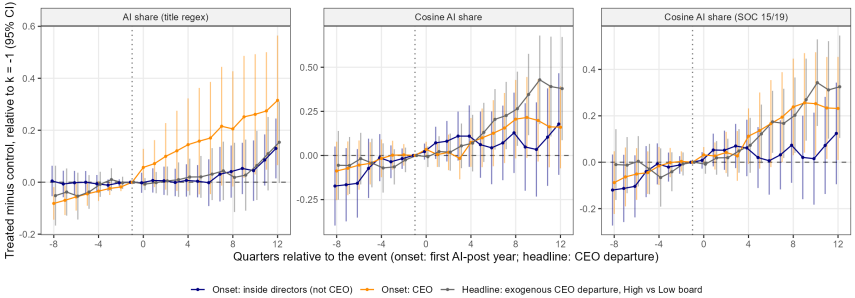
Post × exposure (per s.d.)	Refinement FE	+ reference firms
AI share (title regex), pp	0.017	0.028
AI share (cosine), pp	0.141***	0.117***
AI share (cosine, SOC 15/19), pp	0.122***	0.103***
Log employees	−0.010	−0.013
Junior share, pp	0.012	−0.092
Log sales	−0.014	−0.024
Log COGS	−0.022	−0.035*
Log sales per employee	−0.003	−0.015
Log market value	0.031	−0.016

Exposure = log(1 + board AI posts before the departure), scaled by its s.d. across episodes (0.78). Zero is a Low board. 710 episodes, 122 with nonzero posts.

- High is an any-post indicator; the count is zero for most boards
- The continuous version matches the binary results: one s.d. of exposure raises the cosine share 0.14 pp and lowers employment 1.0 log points
- The average High board sits 2.25 s.d. up the scale, so implied effects are of the same size as the binary ones
- Value outcomes flat here

Back

Appendix: falsification test, AI hiring

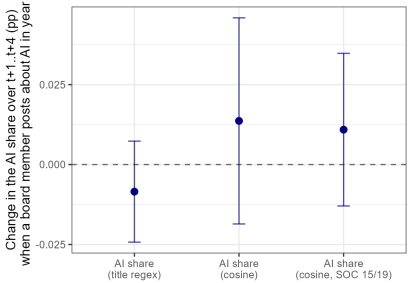


Event studies around the first year the sitting CEO or a board member other than the CEO posts about AI, onsets 2015–21; stack + NAICS3 × quarter + cohort × quarter. Grey: refinement specification.

- CEO onset: title-regex share +0.26 pp** from year two, above the headline (+0.04); cosine as in the headline
- Other board members: title-regex share flat (+0.05, n.s.); cosine smaller, imprecise
- The AI posting that goes with AI hiring is the CEO's: the board's posts do not move hiring while the old CEO is in office

Back

Appendix: falsification test, board posting under a fixed CEO

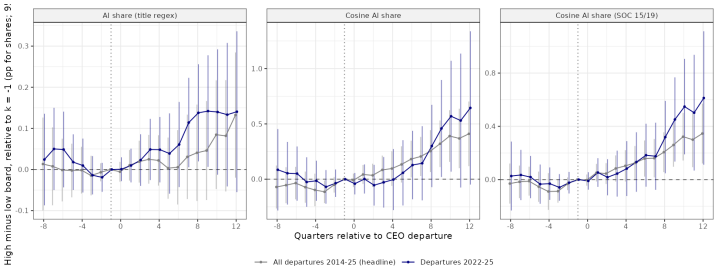


Change in the AI share over the next four quarters, pp, on whether a board member other than the CEO posts about AI in year t. Firm-quarters 2013–2024 under the incumbent CEO, cut off four quarters before any CEO change (58k firm-quarters, 1,858 firms, 3,338 CEO spells). Firm + NAICS3 × quarter FE; 95% CI, clustered by firm.

- While the same CEO is in post, a board posting about AI does not raise AI hiring
- It is the CEO, not the board, who changes hiring behaviour: board AI sentiment binds only when the CEO changes

Back

Appendix: departures from 2022 onwards, AI hiring

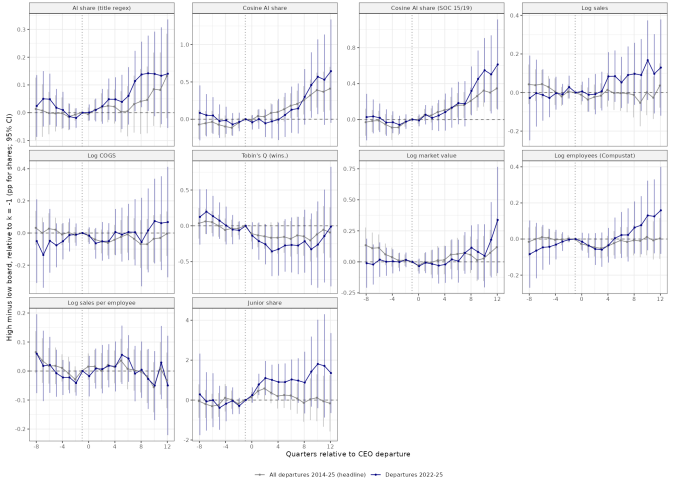


Refinement specification (episode + NAICS3 × quarter + departure-year × quarter FE) on all exogenous departures 2014–25 (grey, 122 High / 588 Low) and on departures from 2022 only (navy, 72 High / 83 Low). 95% CI, clustered by firm.

- Before 2022 "AI" in posts and job titles mostly meant ML, data science and automation; from 2022 it means generative AI
- Departures 2022–25: cosine shares +0.21 (n.s.) and +0.26** pp from year two (headline 0.29, 0.23), flat pre-trends; title regex +0.08 (n.s.); after ChatGPT alone, cosine +0.11 (n.s.). Noisier, with the post period cut off by 2025Q4, but the same direction as the headline
- Even here a large share of the boards' posts are pre-2022 (27 of the 72 High boards): not enough years for an AI-only sample, so our estimates mix generative AI with earlier ML and automation

Back

Appendix: departures from 2022 onwards, all outcomes



Same regressions, all outcomes. Employment: -7.1 log points for the 2014–21 departures and $+6.3$ (n.s.) from 2022. From 2022: Tobin's Q -0.28, market value $+4.8$ log points.

Back

References I

Daron Acemoglu, Alex He, and Daniel le Maire. Eclipse of rent-sharing: The effects of managers' business education on wages and the labor share in the US and Denmark. Working paper 29874, National Bureau of Economic Research, 2022. [Check for a published version].

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024.

Alexander Bick, Adam Blandin, and David J. Deming. The rapid adoption of generative AI. Working paper 32966, National Bureau of Economic Research, 2024.

Nicholas Bloom, Gideon Moore, Lisa K. Simon, and Caelan Wilkie-Rogers. Time travel on professional profiles. Working paper 35546, National Bureau of Economic Research, July 2026.

Andrea L. Eisfeldt, Gregor Schubert, and Miao Ben Zhang. Generative AI and firm values. Working paper 31222, National Bureau of Economic Research, 2023.

Richard J. Gentry, Joseph S. Harrison, Timothy J. Quigley, and Steven Boivie. A database of CEO turnover and dismissal in S&P 1500 firms, 2000–2018. *Strategic Management Journal*, 42(5): 968–991, 2021.

Avi Goldfarb, Bledi Taska, and Florenta Teodoridis. Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings. *Research Policy*, 52(1):104653, 2023.

Anders Humlum and Emilie Vestergaard. Large language models, small labor market effects. Working paper 2025-56, Becker Friedman Institute, 2025.

Kristina McElheran, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2):375–415, 2024.