Generative AI and the Informational Value of Educational Credentials

Evidence from the Master's Margin

Patricia Cortes · Chris Dellarocas · Qi Wang

Questrom School of Business, Boston University

Economics Conference · July 2026 · Preliminary draft

- Educational credentials help employers evaluate workers whose productivity is **hard to observe directly** (Spence 1973).
- But a credential's value depends on the **economic value of the skills it certifies**—not only on what it reveals about ability.
- When technological change alters the value of those skills, employers may reweight the credential.

This paper

Is **generative AI** reshaping employer reliance on education credentials in hiring?

One credential, two signals

Conceptual framework

Signal of general ability

Intelligence, persistence, conscientiousness, ambition—the capacity to complete a demanding program. Valued across jobs; its worth (γ) is unchanged, and may even *rise* in an AI-rich workplace.

Certification of specialized skill

Analysis, synthesis, forecasting, writing, coding taught in the program. Generative AI performs or augments these, so their value (λ) *falls*—especially in non-STEM fields.

- AI erodes the credential's value as a **skill certification**, not its value as a **signal of general ability**.
- The degree keeps its edge *only where employers lack an alternative signal of ability*—the source of the cross-margin pattern.

Undergraduate + experience

Signal *broad* capability: reasoning, communication, adaptability, work habits.

Master's programs

Certify *specialized* analytical skills: information synthesis, financial analysis, forecasting, writing, coding.

These are exactly the tasks recent generative AI performs well
(Brynjolfsson et al. 2023; Noy & Zhang 2023; Eloundou et al. 2023)

If AI devalues these skills, employers may rely less on the credential.

1. Higher-AI-exposure occupations show a **$\sim$10% relative decline** in the master's share of entry-level hires after 2022.
2. Event study: **flat pre-trends**, clean divergence *only* after ChatGPT, widening over time.
3. Concentrated in **non-STEM fields (esp. business)**; **no decline in STEM**.
4. Present for **experienced + internal hires**; **absent for first-job entrants**—the general-ability signal holds where no alternative exists.
5. Strongest at **AI-adopting firms**; **not a supply story**.

Section 1

Conceptual framework

Worker i's productivity in occupation o at time t:

$$y_i = \underbrace{\gamma_t \, \theta_i}_{\text{general capability}} + \underbrace{\lambda_{o,t} \, s_i}_{\text{specialized skill}}$$

- θ_i: general ability & character—judgment, adaptability, persistence. Valued across jobs.
- s_i: specialized analytical skill from graduate training.
- The degree *jointly* signals both; work history speaks mainly to θ, the credential mainly to s.

Generative AI $\equiv$ a fall in $\lambda_{o,t}$ (value of s) after 2022; γ_t (value of θ) is unchanged—or rising.

What the framework predicts

Conceptual framework

Hiring margin	Alternative signal of θ?	Predicted decline
First-job entrants	None (no work history)	**Little / none**
External experienced	Prior work history	**Largest**
Internal transfers	Observed performance	**Present**

- Where no alternative signal of θ exists (entrants), the degree's **general-ability signal** remains decisive—little decline.
- Where work history or performance already reveals θ (experienced, internal), only the devalued ***s*-component** remains—so hiring falls.
- The entrant null is *not* evidence against skill devaluation; it shows the general-ability signal **offsetting** it.

Section 2

Data & empirical strategy

Data: Revelio Labs worker inflows

- **103 million** U.S. position records, Jan 2018 – Nov 2025.
- Unit = a *position* (worker × firm × role); we study **inflows** (new starts = hiring decisions).
- Highest degree held *at the position's start date*.
- Sample restrictions isolate discretion over the master's margin:
  - **Entry-level** roles (seniority 1–2).
  - **Job Zones 3–5** (bachelor's vs. master's choice).
  - Exclude doctorate holders.

Main estimation sample	
Entry-level inflows	**48.9M**
Occupations	**586**
Master's share	**20.1%**
Missing education	43.4%

Master's includes the MBA. Outcome: master's share of inflows with non-missing education.

- Occupation-level, importance-weighted LLM-exposure index β_o^{imp} (Eloundou et al. 2023).
- **Predetermined**: built from a *pre-ChatGPT* assessment of LLM capabilities applied to O*NET tasks.
- Baseline compares the **top (Q5)** vs. **bottom (Q1)** exposure quintiles.

	Q1 (low)	**Q5 (high)**
Mean β_o^{imp}	0.363	0.710
Representative roles	care, support, technical	analysts, accountants, paralegals, writers

Continuous specification uses the full distribution as a complementary check.

Event study (preferred):

$$Y_{ot} = \alpha_o + \gamma_{j(o)t} + \sum_{s \neq 2022} \delta_s \, 1[\text{Exposure}_o = 5] \times 1[t = s] + \varepsilon_{ot}$$

Difference-in-differences (summary):

$$Y_{ot} = \alpha_o + \gamma_{j(o)t} + \delta \, \text{Post}_t \times 1[\text{Exposure}_o = 5] + \varepsilon_{ot}$$

- Y_{ot} = master's share of inflows; α_o occupation FE; $\gamma_{j(o)t}$ **Job Zone** $\times$ **Year** FE. SEs clustered by occupation.
- Identifying assumption: **parallel trends** absent AI. $\text{Post}_t = 1[t \geq 2023]$.

Section 3

Results

Main result: the master's share diverges after ChatGPT

Results — event study

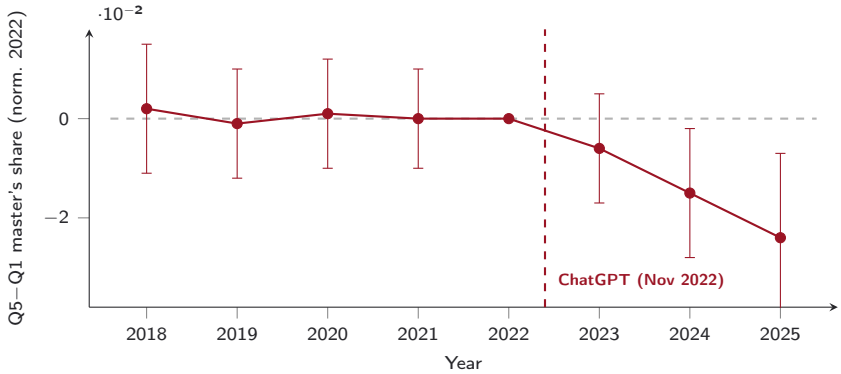

Flat, insignificant pre-trends; divergence emerges only after 2022, reaching −0.024 by 2025.

Reading the event study

Results — event study

- Pre-2018–2021 coefficients **near zero and insignificant** ⇒ parallel-trends assumption supported.
- Post-2022 coefficients turn negative and **grow steadily**: -0.015 (2024), -0.024 (2025).
- Gradual widening is consistent with the **diffusion** of generative AI through workplaces.
- **Not** a general retreat from credentials: bachelor's and bachelor's-or-higher shares show **no discrete post-2022 break** once pre-trends are accounted for—the adjustment is specific to the master's margin.

Heterogeneity by field: the decline is entirely non-STEM

Results — coefficient plot (Table 1)

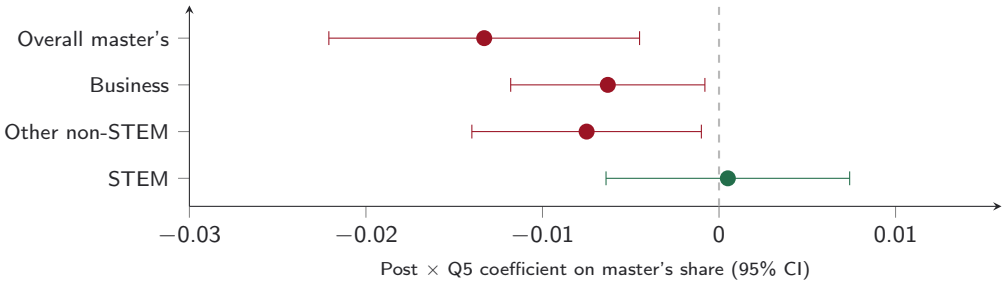

- Decline is **entirely non-STEM**; STEM is a **precise zero**. Business effect ≈**13%** of its 4.8% baseline.
- Mirrors **task overlap**: business training is highly AI-exposed; STEM skills more complementary.

Points are Q5-vs-Q1 DiD estimates; bars are illustrative 95% CIs consistent with reported significance ($p < 0.01/0.05$; STEM n.s.).

The sharpest test: hiring margins

Results — coefficient plot (Table 2)

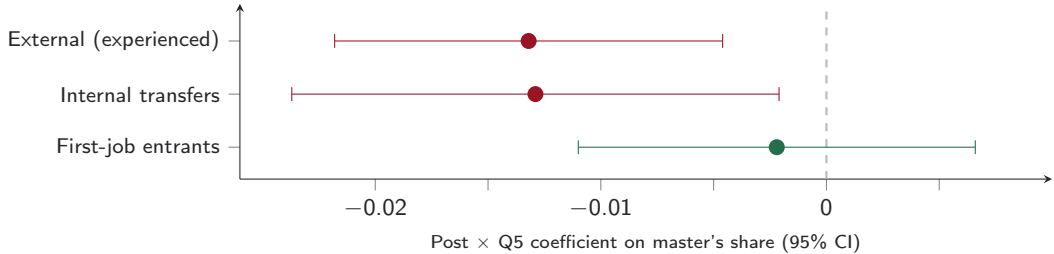

- **First-job entrants: a precise zero**—the CI straddles zero where employers have the *fewest* alternative signals of ability.
- **Experienced** (external) and **internal** hires decline significantly, once work history or performance is observed.
- Consistent with the degree retaining value as a **signal of general ability** while its **specialized-skill** value falls.

Concentrated where AI is actually adopted

Results — coefficient plot (Table 3)

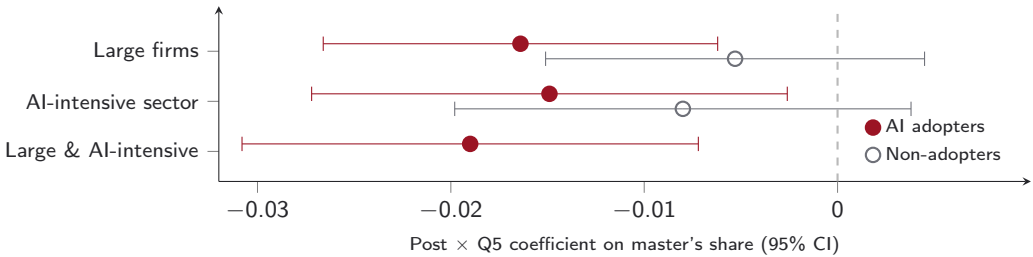

- The decline is **concentrated at AI-adopting firms**; triple-difference (size) = -0.0173^{***} ($\sim$1.7 pp larger at large firms).
- Exposure is measured **independently of adoption** (Census BTOS, Bonney et al. 2026) $\Rightarrow$ validation of the identification strategy.

It is demand, not supply

Results — ruling out supply-side stories

Supply kept growing

- Master's completions: **815k (2018) → 949k (2025)**, no break.
- Prime-age master's in labor force: **13.7M → 17.1M**.
- Business +6%, STEM +40%.

Not international students

- Non-citizen master's concentrated in *STEM* (26% vs. 7%).
- But the decline is in *non-STEM*—the opposite pattern.

Hiring below the available pool

High-AI occupations moved from hiring master's at a **premium** (+0.6 pp, 2018) to a **discount** (−1.5 pp, 2025) relative to the fixed economy-wide pool—an *employer-selection* shift.

Robustness

The result is not fragile

The central finding survives a battery of checks:

- **Missing education** (43%): no differential Q5/Q1 trend in missingness; robust to alternative treatments.
- **Composition**: role-level resampling within occupations.
- **Sample edges**: excluding graduate-required and low-credential service occupations.
- **Contemporaneous shocks**: excluding finance and technology firms.
- **Pre-trend detrending**: binary Q5-vs-Q1 essentially unchanged (-0.020, $t = -2.4$ in 2025).

Section 4

Conclusion

Conclusion

Takeaways

- Generative AI reduces **employer reliance on graduate credentials**—a margin beyond skills, tasks, productivity, wages, and employment.
- Concentrated in **AI-exposed, non-STEM** fields and in hiring contexts where credentials traditionally matter most.
- AI affects the **two components of credential value differently**: the specialized-skill certification loses value while the general-ability signal endures.

Broader implication

Technological change alters not only the value of skills but the **informational value of the credentials that certify them**. Signaling value is itself an *outcome* of the technological environment.

Managerial implications

For graduate education

- Graduate education does not become *less* valuable—but programs whose value rests on **AI-performable skills** face pressure to reinvent.
- Shift content toward **AI-complementary capabilities**: judgment, problem framing, decision-making, domain expertise, human–AI collaboration, leadership.
- From **one-time degrees** toward **continuous reskilling** and modular credentialing.
- Rising pressure on universities to demonstrate value **beyond content transmission**.

Thank you

Questions & comments welcome

Patricia Cortes · Chris Dellarocas · Qi Wang
Questrom School of Business, Boston University

Selected references I

- Altonji & Pierret (2001). Employer learning and statistical discrimination. *QJE* 116(1).
- Autor, Levy & Murnane (2003). The skill content of recent technological change. *QJE* 118(4).
- Bonney et al. (2026). The microstructure of AI diffusion. *NBER* WP 35141.
- Brynjolfsson, Li & Raymond (2023). Generative AI at work. *NBER* WP 31161.
- Deming & Kahn (2018). Skill requirements across firms and labor markets. *JOLE* 36(S1).
- Deming & Silliman (2024). Foundational versus specific skills. *NBER* WP.
- Eloundou, Manning, Mishkin & Rock (2023). GPTs are GPTs.
- Lichtinger & Hosseini Maasoum (2025). AI and hiring.
- Noy & Zhang (2023). Productivity effects of generative AI. *Science*.
- Spence (1973). Job market signaling. *QJE* 87(3).