Artificial Intelligence and the Labor Market

M. Hampole[1] D. Papanikolaou[2,3] L. Schmidt[4] B. Seegmiller[2]

[1]Yale School of Management

[2]Kellogg School of Management

[3]NBER

[4]MIT Sloan

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

- Investment in ‘artificial intelligence’ (ML, big data analytics, etc) has been happening for well over a decade, yet measuring the labor market impact has been elusive

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

- Investment in 'artificial intelligence' (ML, big data analytics, etc) has been happening for well over a decade, yet measuring the labor market impact has been elusive

- **Challenge:** AI can increase or decrease labor demand for specific occupations

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

Investment in 'artificial intelligence' (ML, big data analytics, etc) has been happening for well over a decade, yet measuring the labor market impact has been elusive

Challenge: AI can increase or decrease labor demand for specific occupations

AI can substitute for labor in certain tasks.

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

- Investment in ‘artificial intelligence’ (ML, big data analytics, etc) has been happening for well over a decade, yet measuring the labor market impact has been elusive
- **Challenge:** AI can increase or decrease labor demand for specific occupations
  - AI can substitute for labor in certain tasks.
  - But workers may become better at doing non-exposed tasks, increasing productivity and demand for the job.

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

Investment in ‘artificial intelligence’ (ML, big data analytics, etc) has been happening for well over a decade, yet measuring the labor market impact has been elusive

Challenge: AI can increase or decrease labor demand for specific occupations

AI can substitute for labor in certain tasks.

But workers may become better at doing non-exposed tasks, increasing productivity and demand for the job.

As firm productivity increases, overall labor demand may increase (spillovers).

Introduction

Recent advances in AI technology have magnified concerns about its impact on labor demand

- Investment in 'artificial intelligence' (ML, big data analytics, etc) has been happening for well over a decade, yet measuring the labor market impact has been elusive

- **Challenge:** AI can increase or decrease labor demand for specific occupations

  - AI can substitute for labor in certain tasks.
  - But workers may become better at doing non-exposed tasks, increasing productivity and demand for the job.
  - As firm productivity increases, overall labor demand may increase (spillovers).

This Paper: Can we tease out some of these forces in the data?

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

What We Find:

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

What We Find:

- Tasks exposed to AI technology experience lower labor demand.

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

What We Find:

- Tasks exposed to AI technology experience lower labor demand.
- Within firms:

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

What We Find:

- Tasks exposed to AI technology experience lower labor demand.
- Within firms: AI induced a net substitution away from highly-paid occupations.

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

What We Find:

- Tasks exposed to AI technology experience lower labor demand.
- Within firms: AI induced a net substitution away from highly-paid occupations.
- In aggregate:
  1. AI adoption leads firm to grow faster.

This Paper

Measurement: Combine NLP tools w/ data on online resumes and O*NET task descriptions to measure both AI adoption and exposure at a granular level (firm–occ–year).

Theory: Maps the distribution of AI exposure across tasks into sufficient statistics for occupation labor demand.

What We Find:

- Tasks exposed to AI technology experience lower labor demand.
- Within firms: AI induced a net substitution away from highly-paid occupations.
- In aggregate:
  1. AI adoption leads firm to grow faster.
  2. Employment share of highly-paid occupations increases as higher-paid occupations concentrated in AI-adopting firms.

Related Work

- AI adoption, firm growth, and labor demand: (Acemoglu et al., 2022; Acemoglu et al., 2023; Acemoglu, 2024; Eloundou et al., 2023; Eisfeldt, Schubert, Taska, and Zhang, 2023; Babina, Fedyk, He, and Hodson, 2023,2024; Humlum & Vestergaard, 2024)
- Direct measures of labor-saving technologies and labor outcomes: (Acemoglu and Restrepo, 2021; Aghion et al, 2021; Graetz and Michaels, 2018; Felten, Raj, & Seamans, 2018; Humlum, 2019; Webb, 2020; Aghion, et al., 2020; Dauth, et al., 2021; Koch, et al., 2021; Bonfiglioli et al., 2020; de Souza and Li, 2023; Kogan et al., 2023; Autor et al., 2024; Mann and Püttmann, 2023; Dechezleprêtre et al. , 2021; Jiang et al, 2025)

Key contributions:

- ▶ Use new corpus to build detailed, highly granular measures of AI adoption + worker AI exposure (firm $\times$ occ $\times$ time-varying)
- ▶ Theoretically & empirically: emphasize gains from reallocation across tasks + firms

Data and Measurement

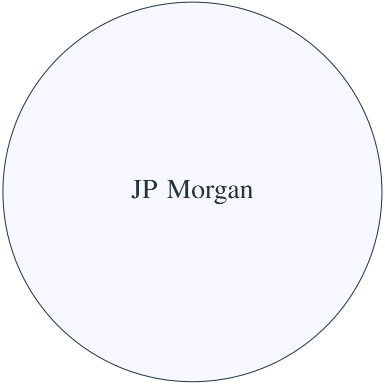
JP Morgan

J: AI developer in JP Morgan:

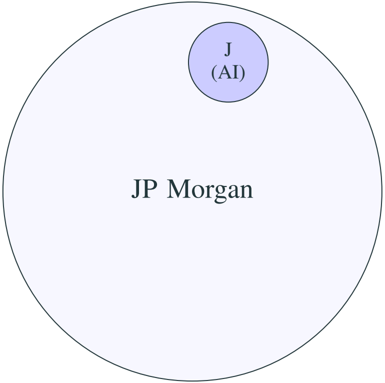

"Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

Step 1: Identify AI Developers using AI terms

J: AI developer in JP Morgan:

resumes $\rightarrow$ measure adoption of specific AI applications

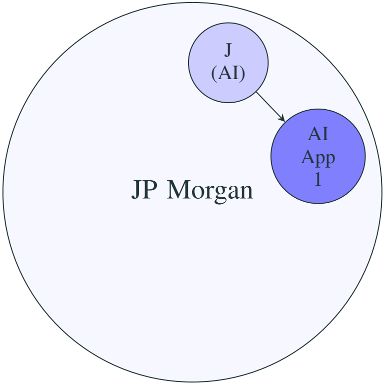

"Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

Step 1: Identify AI Developers using AI terms

Step 2: Extract and Clean up AI applications

Use LLMs to extract and clean phrases containing descriptions of AI specific uses ("AI applications")

A, B, C, D: Other workers in JP Morgan: potentially exposed to AI

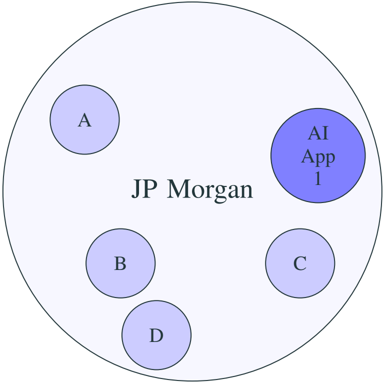

"Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

Step 1: Identify AI Developers using AI terms

Step 2: Extract and Clean up AI applications
Use LLMs to extract and clean phrases containing descriptions of AI specific uses ("AI applications")

A, B, C, D: Other workers in JP Morgan: potentially exposed to AI

- from distance between O*NET task descriptions and AI app 1

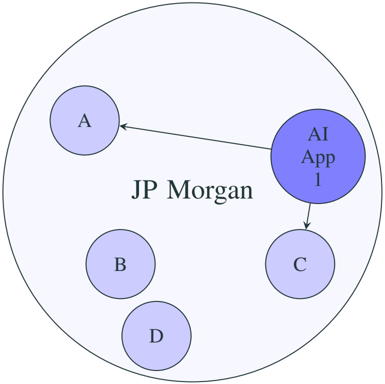

"Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

Step 1: Identify AI Developers using AI terms

Step 2: Extract and Clean up AI applications

Step 3: Task exposure

Use document embeddings (vector representations of text meaning) to get similarity of AI applications with O*NET occupational task descriptions

A, B, C, D: Other workers in JP Morgan: potentially exposed to AI

- from distance between O*NET task descriptions and AI app 1

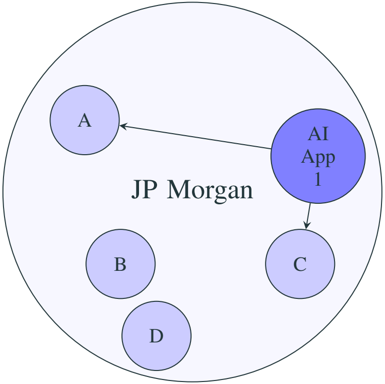

"Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

Step 1: Identify AI Developers using AI terms

Step 2: Extract and Clean up AI applications

Step 3: Task exposure
Use document embeddings (vector representations of text meaning) to get similarity of AI applications with O*NET occupational task descriptions

> **Most similar O*NET task to first application:**
> "Prepare reports that include the degree of risk involved in extending credit or lending money." (Credit Analysts)

K, M: AI developers in Walmart:

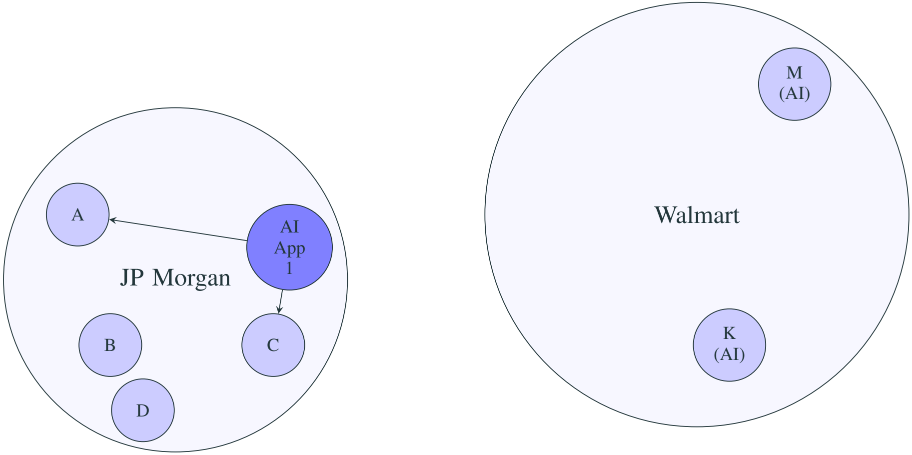

K, M: AI developers in Walmart: resumes → measure adoption of specific AI applications

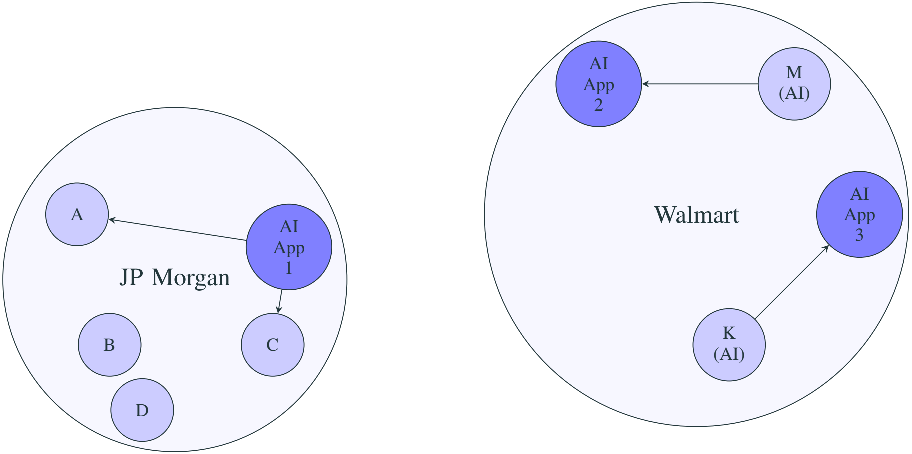

O, P, Q, R, S, T: Other workers in Walmart (potentially exposed to AI)

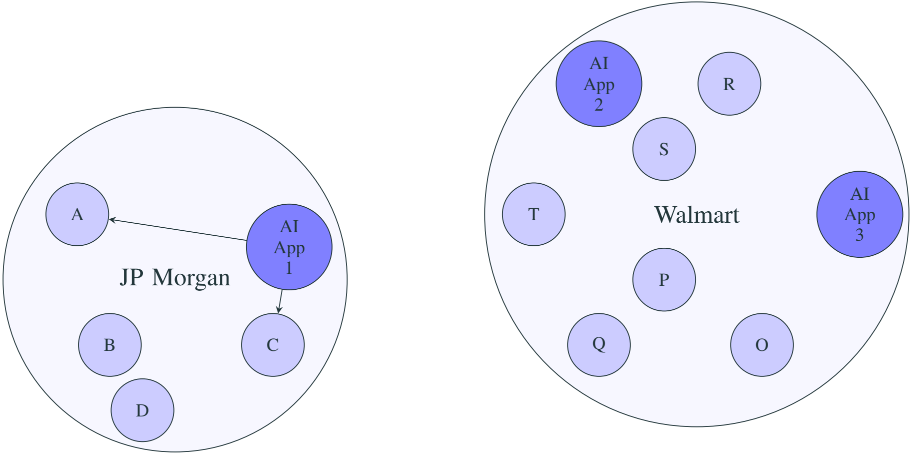

O, P, Q, R, S, T: Other workers in Walmart (potentially exposed to AI)

- based on distance between ONET task descriptions and AI app 2 and 3

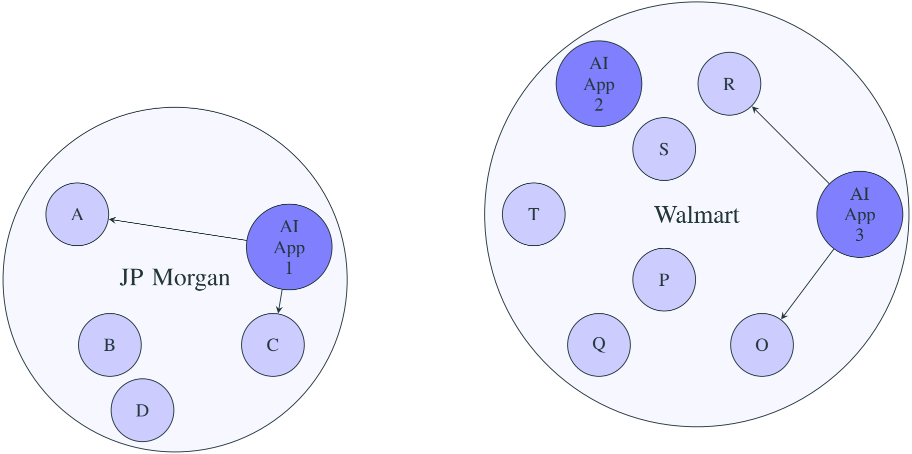

AI exposure: Granular measure that varies across occupations, firms, and time.

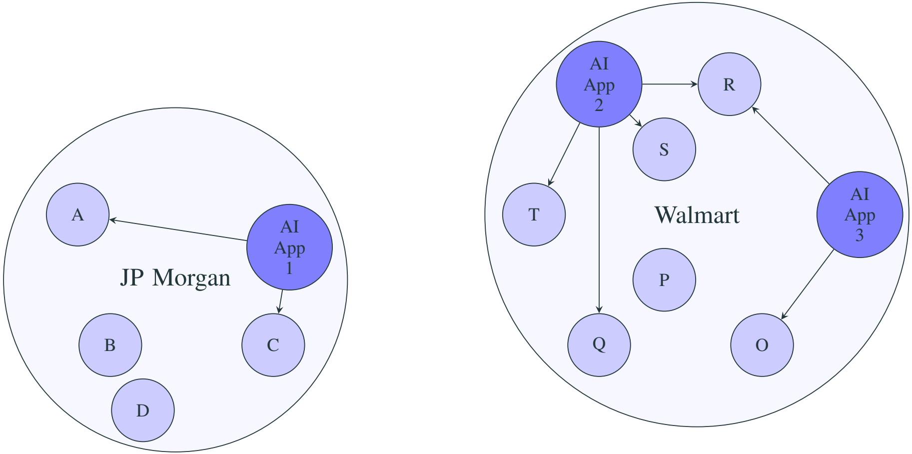

Example: Overview of AI applications at JP Morgan Chase

Application summary description	Examples of highly exposed tasks	Associated occupations
Fraud Detection, AML & Risk Mitigation	Collect and analyze data to detect deficient controls, duplicated effort, extravagance, fraud, or non-compliance with laws, regulations, and management policies.	Accountants and Auditors
	Research or evaluate new technologies for use in fraud detection systems.	Other Financial Specialists
Predictive Modeling & Financial Forecasting	Consult financial literature to ensure use of the latest models or statistical techniques.	Other Financial Specialists
	Research or develop analytical tools to address issues such as portfolio construction or optimization, performance measurement, attribution, profit and loss measurement, or pricing models.	Other Financial Specialists
Customer Engagement & Personalization	Monitor customer preferences to determine focus of sales efforts.	Sales Managers
	Identify interested and qualified customers to provide them with additional information.	Models, Demonstrators, and Product Promoters

Other clusters: Data Engineering & Analytics Infrastructure; Automation & Workflow Optimization

Example: Overview of AI applications at Walmart

Application summary description	Examples of highly exposed tasks	Associated occupations
Forecasting, Pricing, and Supply Chain Optimization	Analyze market and delivery systems to assess present and future material availability.	Purchasing Managers
	Monitor and analyze sales records, trends, or economic conditions to anticipate consumer buying patterns, company sales, and needed inventory.	Wholesale and Retail Buyers, Except Farm Products
Process Automation and Operational Efficiency	Plan and modify product configurations to meet customer needs.	Sales Engineers
	Monitor and adjust production processes or equipment for quality and productivity.	Other Engineering Technologists And Technicians, Except Drafters
Fraud, Security, and Anomaly Detection	Analyze retail data to identify current or emerging trends in theft or fraud.	Other Managers
	Monitor machines that automatically measure, sort, or inspect products.	Inspectors, Testers, Sorters, Samplers, and Weighers

Other clusters: Personalization, Recommendations, and Enhanced Search;

Data Pipelines, Integration, and Big Data Infrastructure Back

In Paper: Overview of AI applications across sectors

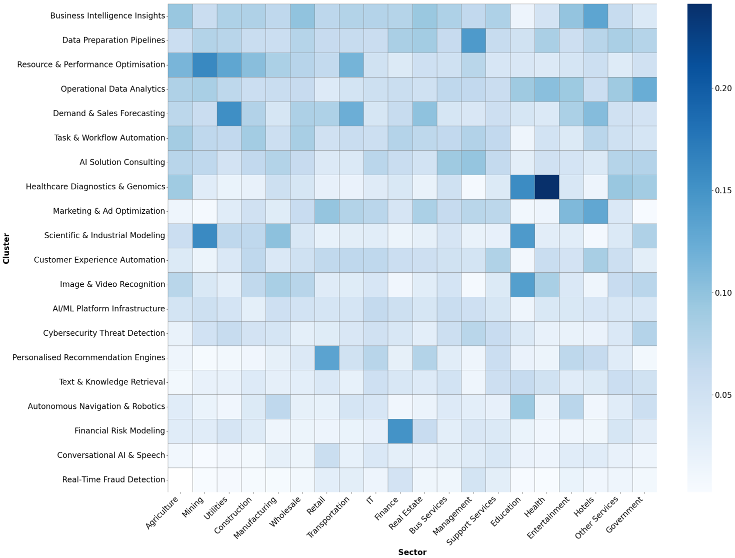

Measurement Validations

- Resume-implied AI utilization rates by sector × firm size correlate highly w/ average firm-reported AI utilization rates in Census BTOS surveys (ρ ≈ 0.9): Details Sector × Size

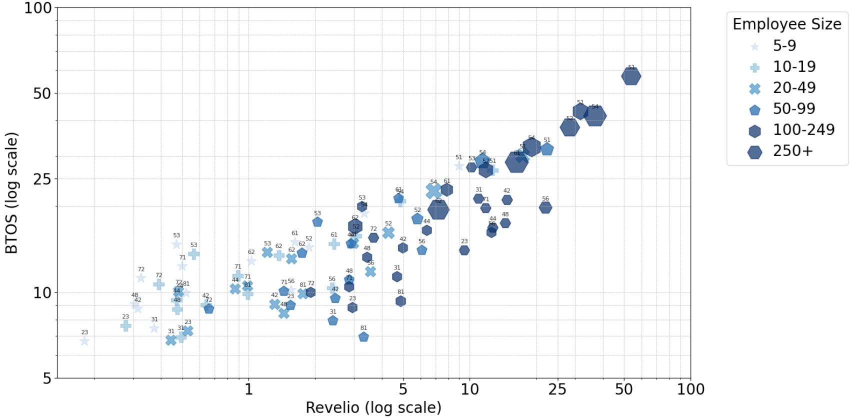

Measurement Validations

- Resume-implied AI utilization rates by sector $\times$ firm size correlate highly w/ average firm-reported AI utilization rates in Census BTOS surveys ($\rho \approx 0.9$): Details Sector $\times$ Size
- Firm-level resume-implied additions of new AI workers co-occur with firm job postings seeking new AI hires Details

Measurement Validations

- Resume-implied AI utilization rates by sector $\times$ firm size correlate highly w/ average firm-reported AI utilization rates in Census BTOS surveys ($\rho \approx 0.9$): Details Sector $\times$ Size
- Firm-level resume-implied additions of new AI workers co-occur with firm job postings seeking new AI hires Details
- Firm AI resume use strongly related with AI patenting Details
- Firms that adopt AI are larger, more productive, and pay more Details

    Consistent with survey evidence (Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo, and Zolas, 2023)

AI applications and tasks

1. Compute the similarity scores between all 20k occupation task $\times$ 1.3m AI use pairs

AI applications and tasks

1. Compute the similarity scores between all 20k occupation task $\times$ 1.3m AI use pairs
2. Task is exposed to a particular AI use: the task–application similarity score above the unconditional 95th percentile.

AI applications and tasks

1. Compute the similarity scores between all 20k occupation task $\times$ 1.3m AI use pairs
2. Task is exposed to a particular AI use: the task–application similarity score above the unconditional 95th percentile.
3. Aggregate across all applications in a given firm.

$$\text{Exposure Probability}_{j,f,t} = \frac{1}{N_{f,t}} \sum_{i=1}^{N_{f,t}} \text{Above 95th percentile indicator}_{j,i}$$

AI applications and tasks

1. Compute the similarity scores between all 20k occupation task $\times$ 1.3m AI use pairs
2. Task is exposed to a particular AI use: the task–application similarity score above the unconditional 95th percentile.
3. Aggregate across all applications in a given firm.

$$\text{Exposure Probability}_{j,f,t} = \frac{1}{N_{f,t}} \sum_{i=1}^{N_{f,t}} \text{Above 95th percentile indicator}_{j,i}$$

4. Account for intensity of AI adoption (number of AI workers)

$$\text{Task-Level AI Exposure}_{j,f,t} = \text{Exposure Probability}_{j,f,t} \times \log(1 + N_{f,t})$$

Fact 1: Average Task AI Exposure Probability is (Mostly) Increasing in Salary Rank

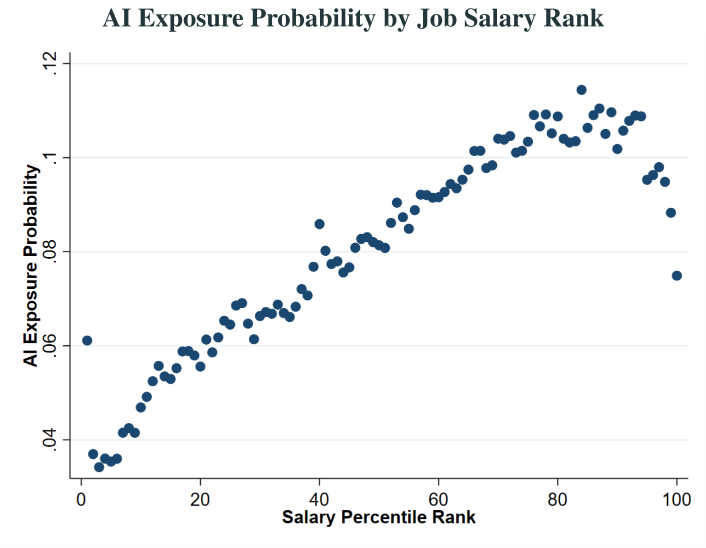

Fact 1: Average Task AI Exposure Probability is (Mostly) Increasing in Salary Rank

AI Exposure Probability by Job Salary Rank

Most Exposed Occupations

- Market Research Analysts and Marketing Specialists
- Management Analysts
- Logisticians
- Computer Hardware Engineers
- Financial Specialists
- Computer and Information Systems Managers
- Sales Engineers
- Financial Risk Specialists

Fact 2: Effort shifts away from AI exposed tasks

Map skills listed in job postings into O*NET tasks using sentence embeddings

Dep. Variable: 100 $\times$ 5-year DHS growth in share of job posting skills related to task (see paper for examples)	OLS		
	(1)	(2)	(3)
Task-level AI Exposure	-4.78***	-4.75***	-4.80***
	(-13.46)	(-13.97)	(-14.14)
Observations (task–occ–firm–year)	13.2m	13.2m	13.2m
Controls			
ONET Task Importance	X	X	X
Mean Occ Task Exposure	X	X	
Firm $\times$ Year FE	X	X	
Occ $\times$ Year FE		X	
Firm $\times$ Occ $\times$ Year FE			X

Model

Summary of Theoretical Framework

Output is nested CES of firm, occupation and task inputs with different elasticities

Tasks combine labor effort and a capital (AI)

Summary of Theoretical Framework

Output is nested CES of firm, occupation and task inputs with different elasticities

Tasks combine labor effort and a capital (AI)

Workers allocate effort across tasks j, with decreasing returns

Summary of Theoretical Framework

Output is nested CES of firm, occupation and task inputs with different elasticities

Tasks combine labor effort and a capital (AI)

Workers allocate effort across tasks j, with decreasing returns

AI technology is a direct substitute for the associated task j

AI improvements: decline in cost of capital

Additional Model Details

Impact of a given technology on occupation labor demand

Log-linearizing around a symmetric equilibrium, we find that

$$\underbrace{\frac{1}{\zeta}\Delta_\varepsilon \log N(o,f)}_{\Delta_\varepsilon \log W(o,f)} \approx \underbrace{\Big(\eta_o + \eta_c\,(J-1)\Big)}_{<0} m(\varepsilon)$$

$m(\varepsilon)$ denotes the mean of AI-induced technology productivity improvement across job tasks,

$$m(\varepsilon) \equiv \frac{1}{J}\sum_{j\in J} \varepsilon(j)$$

Impact of a given technology on occupation labor demand

Log-linearizing around a symmetric equilibrium, we find that

$$\underbrace{\frac{1}{\zeta}\Delta_\varepsilon \log N(o,f)}_{\Delta_\varepsilon \log W(o,f)} \approx \underbrace{\Big(\eta_o + \eta_c\,(J-1)\Big)}_{<0} m(\varepsilon)$$

$m(\varepsilon)$ denotes the mean of AI-induced technology productivity improvement across job tasks,

$$m(\varepsilon) \equiv \frac{1}{J}\sum_{j\in J}\varepsilon(j)$$

Implication: Within-firm employment growth:

- decreases in average AI task exposure

Impact of a given technology on occupation labor demand

Log-linearizing around a symmetric equilibrium, we find that

$$\underbrace{\frac{1}{\zeta}\Delta_\varepsilon \log N(o,f)}_{\Delta_\varepsilon \log W(o,f)} \approx \underbrace{\Big(\eta_o + \eta_c\,(J-1)\Big)}_{<0} m(\varepsilon) + \underbrace{\frac{1}{2\beta}\,(\eta_o+\eta_c)^2}_{>0}\, C(\varepsilon) + \text{Spillovers}$$

$m(\varepsilon)$ and $C(\varepsilon)$ denote the mean and concentration of AI-induced technology productivity improvement across job tasks,

$$m(\varepsilon) \equiv \frac{1}{J}\sum_{j\in J}\varepsilon(j) \qquad \text{and} \qquad C(\varepsilon) \equiv \frac{1}{J}\sum_{j\in J}\Big(\varepsilon(j) - m(\varepsilon)\Big)^2$$

Implication: Within-firm employment growth:

- decreases in average AI task exposure

Impact of a given technology on occupation labor demand

Log-linearizing around a symmetric equilibrium, we find that

$$\underbrace{\frac{1}{\zeta}\Delta_\varepsilon \log N(o,f)}_{\Delta_\varepsilon \log W(o,f)} \approx \underbrace{\Big(\eta_o + \eta_c\,(J-1)\Big)}_{<0} m(\varepsilon) + \underbrace{\frac{1}{2\beta}\left(\eta_o + \eta_c\right)^2}_{>0} C(\varepsilon) + \text{Spillovers}$$

$m(\varepsilon)$ and $C(\varepsilon)$ denote the mean and concentration of AI-induced technology productivity improvement across job tasks,

$$m(\varepsilon) \equiv \frac{1}{J}\sum_{j\in J}\varepsilon(j) \qquad \text{and} \qquad C(\varepsilon) \equiv \frac{1}{J}\sum_{j\in J}\Big(\varepsilon(j) - m(\varepsilon)\Big)^2$$

Implication: Within-firm employment growth:

- decreases in average AI task exposure
- increases in concentration, bc reallocating time most beneficial when AI improvements concentrated in a subset of tasks

Firm-Level Productivity Spillovers Across All Occupations

AI-related cost improvements generates a productivity spillover effect across occupations

$$\underbrace{\frac{1}{\zeta}\Delta_\varepsilon \log N(o,f)}_{\Delta_\varepsilon \log W(o,f)} \approx \text{Direct Effects} + \eta_z \Delta_\varepsilon \log Z_f$$

where $\eta_z > 0$ and Z_f is firm productivity.

Implication: Labor demand is increasing in the extent of AI technology use at the firm, holding occupation-specific effects constant

$\Rightarrow$ AI raises firm growth + employment

Implications

Mapping model to Data

Compute mean and concentration of task AI exposure at firm $\times$ occupation level

Assumption: technology improvement also function of the extent of AI use at the firm, measured by the number of AI applications.

Endogeneity Concerns:

1. Within firm: AI adoption targeted to specific occupations
2. Across firms: Large and productive firms tend to implement AI

Mean vs Concentration: Walmart Example

‘Shift-share’ Instrument

- Instrument the mean exposure and concentration in AI exposure of occupation o in firm f in year t with mean and concentration across all firms.
- Instrument the intensity of AI adoption of firm f with predicted adoption based on arguably exogenous shift in supply of AI workers.
  - **Predicted AI employees**: use 2005-2009 average share of employees graduated from university u, interacted with AI workers coming from university u.

Details Hiring practices are persistent

Across firms: AI adoption leads firms to grow and become more productive

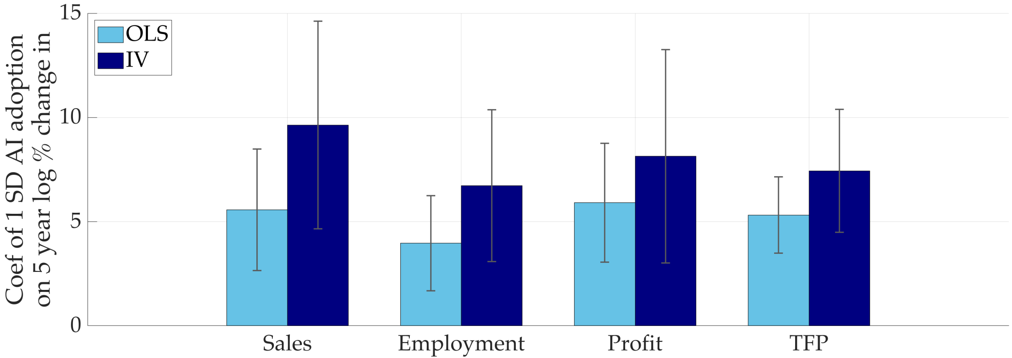

Impact of AI Exposure on Firm Outcomes: 5-year growth rates (in log p.p.)

$$\log(Y_{f,t+5}) - \log(Y_{f,t}) = \gamma \log(1 + \text{AI uses})_{f,t} + \beta X_{f,t} + \varepsilon_{f,t}$$

Robustness

AI Exposure and Within-Firm Occupational Employment Growth

Dep. Var: log 5 year growth of firm–occupation employment

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
AI Exposure Average	-8.34***	-7.73***	-5.50***	-14.1***	-14.5***	-10.4***
	(-14.35)	(-12.87)	(-10.47)	(-10.69)	(-16.94)	(-11.37)
AI Exposure Concentration	1.50***	1.87***	1.42***	7.25***	7.46***	7.33***
	(4.37)	(4.94)	(4.87)	(5.56)	(8.34)	(5.74)
$\log(1 + \text{AI uses})$	11.3***			17.2***		
	(19.28)			(14.92)		
Observations (firm-occ)	2.04m	2.04m	2.04m	2.02m	2.02m	2.02m
Controls	X	X	X	X	X	X
Industry × Year FE	X			X		
Firm × Year FE		X	X		X	X
Occ × Year FE			X			X

Robustness

AI Exposure and Within-Firm Occupational Employment Growth

Dep. Var: log 5 year growth of firm–occupation employment

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
AI Exposure Average	-8.34***	-7.73***	-5.50***	-14.1***	-14.5***	-10.4***
	(-14.35)	(-12.87)	(-10.47)	(-10.69)	(-16.94)	(-11.37)
AI Exposure Concentration	1.50***	1.87***	1.42***	7.25***	7.46***	7.33***
	(4.37)	(4.94)	(4.87)	(5.56)	(8.34)	(5.74)
$\log(1 + \text{AI uses})$	11.3***			17.2***		
	(19.28)			(14.92)		
Observations (firm-occ)	2.04m	2.04m	2.04m	2.02m	2.02m	2.02m
Controls	X	X	X	X	X	X
Industry $\times$ Year FE	X			X		
Firm $\times$ Year FE		X	X		X	X
Occ $\times$ Year FE			X			X

Robustness

AI Exposure and Within-Firm Occupational Employment Growth

Dep. Var: log 5 year growth of firm–occupation employment

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
AI Exposure Average	-8.34***	-7.73***	-5.50***	-14.1***	-14.5***	-10.4***
	(-14.35)	(-12.87)	(-10.47)	(-10.69)	(-16.94)	(-11.37)
AI Exposure Concentration	1.50***	1.87***	1.42***	7.25***	7.46***	7.33***
	(4.37)	(4.94)	(4.87)	(5.56)	(8.34)	(5.74)
$\log(1 + \text{AI uses})$	11.3***			17.2***		
	(19.28)			(14.92)		
Observations (firm-occ)	2.04m	2.04m	2.04m	2.02m	2.02m	2.02m
Controls	X	X	X	X	X	X
Industry $\times$ Year FE	X			X		
Firm $\times$ Year FE		X	X		X	X
Occ $\times$ Year FE			X			X

Robustness

AI Exposure and Within-Firm Occupational Employment Growth

Dep. Var: log 5 year growth of firm–occupation employment

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
AI Exposure Average	-8.34***	-7.73***	-5.50***	-14.1***	-14.5***	-10.4***
	(-14.35)	(-12.87)	(-10.47)	(-10.69)	(-16.94)	(-11.37)
AI Exposure Concentration	1.50***	1.87***	1.42***	7.25***	7.46***	7.33***
	(4.37)	(4.94)	(4.87)	(5.56)	(8.34)	(5.74)
$\log(1+\text{AI uses})$	11.3***			17.2***		
	(19.28)			(14.92)		
Observations (firm-occ)	2.04m	2.04m	2.04m	2.02m	2.02m	2.02m
Controls	X	X	X	X	X	X
Industry $\times$ Year FE	X			X		
Firm $\times$ Year FE		X	X		X	X
Occ $\times$ Year FE			X			X

Robustness

Aggregate effects? Impact of AI on employment across the pay distribution

Use regression coefficients and empirical distribution of exposure measures to predict average net impact of AI on labor demand across the salary distribution:

Growth in Aggregate Employment

Aggregate effects? Impact of AI on employment across the pay distribution

Use regression coefficients and empirical distribution of exposure measures to predict average net impact of AI on labor demand across the salary distribution:

Mean task-level exposure: sharply ↓ in income

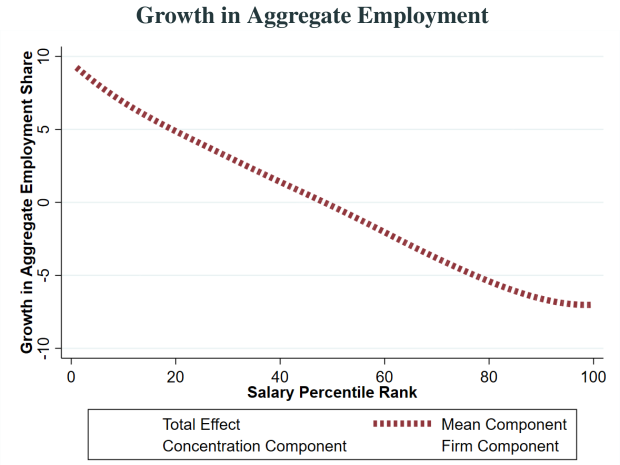

Aggregate effects? Impact of AI on employment across the pay distribution

Use regression coefficients and empirical distribution of exposure measures to predict average net impact of AI on labor demand across the salary distribution:

Mean task-level exposure: sharply ↓ in income

Concentration of task-level exposure: ↑ in income

Growth in Aggregate Employment

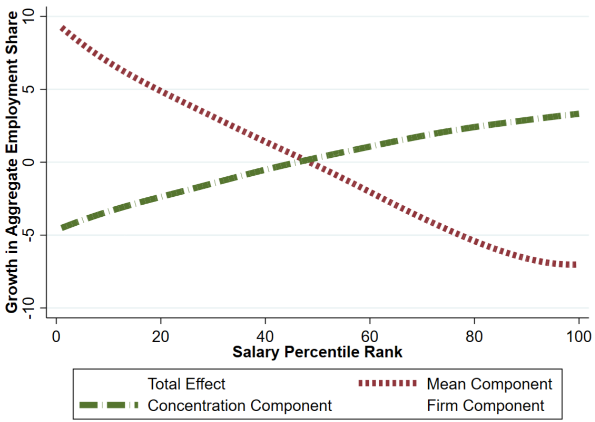

Aggregate effects? Impact of AI on employment across the pay distribution

Use regression coefficients and empirical distribution of exposure measures to predict average net impact of AI on labor demand across the salary distribution:

Mean task-level exposure: sharply ↓ in income

Concentration of task-level exposure: ↑ in income

Firm productivity effect: ↑ in income

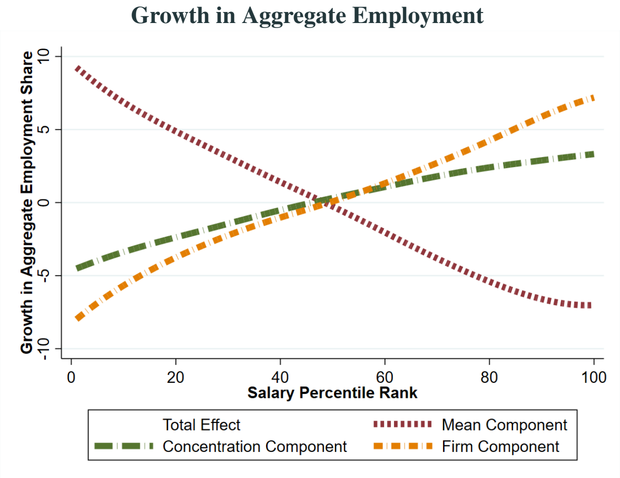

Aggregate effects? Impact of AI on employment across the pay distribution

Use regression coefficients and empirical distribution of exposure measures to predict average net impact of AI on labor demand across the salary distribution:

Mean task-level exposure: sharply ↓ in income

Concentration of task-level exposure: ↑ in income

Firm productivity effect: ↑ in income

Total effect: modestly ↑ in income

Growth in Aggregate Employment

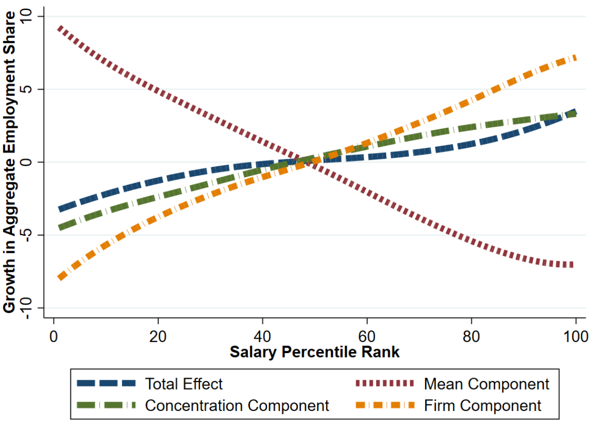

Impact of AI on employment across the pay distribution within the firm

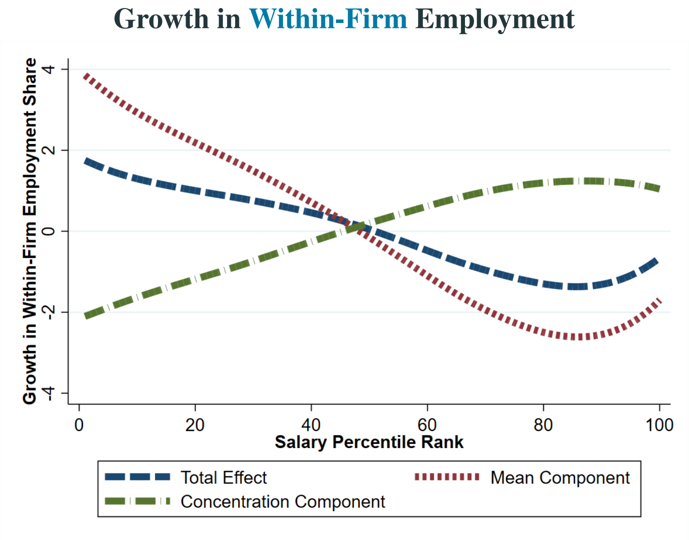

- Aggregate emp share of higher-paid jobs increase (jobs concentrated in firms who adopt AI)
- But their **within-firm employment shares have declined** (mean effect is stronger than variance effect)

Breakdown by occupation group

Conclusion

- Using NLP techniques and model as a guide, isolate different channels through which AI impacts labor demand

- Main findings:
  1. Large substitution effects reduce labor demand
  2. (1) dampened by productivity gains from reallocating time (concentration effect)
  3. Higher-paid workers employed in AI adopting firms, which grow faster

- (2) and (3) largely offset (1), so small net impact: AI has moderately increased labor demand for higher-paid workers relative to lower-paid workers.

Appendix

Measurement: Overview and Data Sources

1. Compustat (focus on publicly traded companies)
   - Examine firm growth, control for firm observables.
2. Resumes from Revelio Labs (2014–2023 period so AI $\neq$ Gen AI)
   - Resumes of AI developers to extract applications they develop for their firm.
3. ONET task descriptions
   - Distance between AI applications and task descriptions $\rightarrow$ AI task exposure.
4. Job posting text from Revelio with tagged skills from LightCast
   - Measure labor demand for specific tasks.
5. Model
   - Distribution of occupation task exposure $\rightarrow$ occupation labor demand.

Comparing Revelio Employment to Compustat

Binscatters of Revelio log employment and 5-year employment growth against Compustat equivalents:

Log Employment

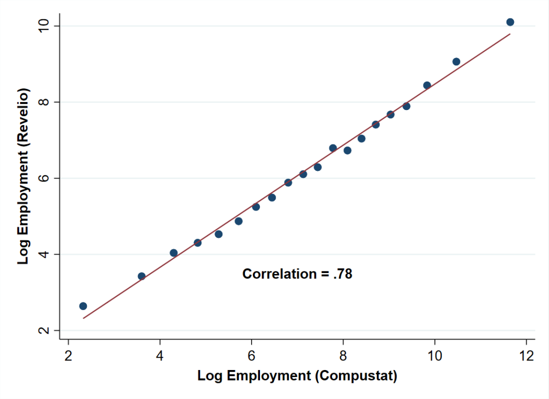

Employment Growth

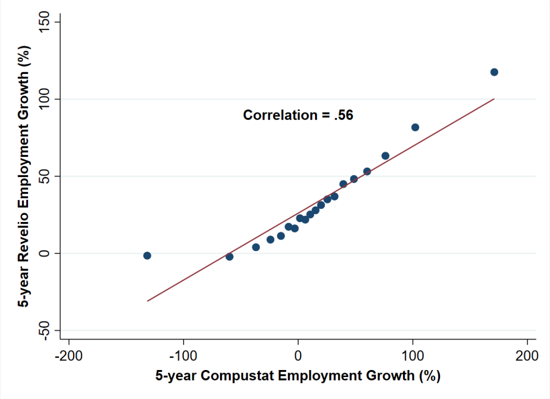

Back

Firms that adopt AI are larger, more productive, and pay more

	(1) Log Sales per worker	(2) Log Sales	(3) Log Profit	(4) Log TFP	(5) Log Average Salary
$\log(1+\text{AI uses})$	0.117***	0.310***	0.415***	0.125***	0.109***
	(6.87)	(12.39)	(17.48)	(10.37)	(18.81)
N	33541	36227	33309	17034	38211
R-sq	0.345	0.644	0.614	0.181	0.427
Revelio Emp Control	X	X	X	X	X
Ind $\times$ Year FE	X	X	X	X	X

Consistent with survey evidence (Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo, and Zolas, 2023) Back

Resume-based AI hiring and AI-related job postings

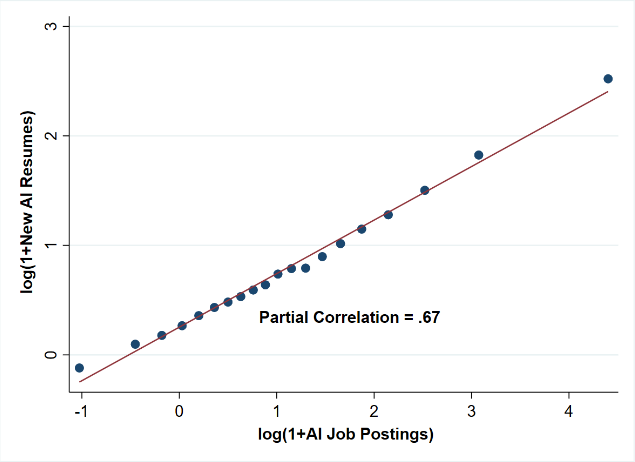

Plotted: Residualized binscatter of $\log(1+\text{AI-Related Job Postings}_{f,t})$ against $\log(1+\text{Newly Added AI Resumes}_{f,t})$

Controls: $\log(\text{Total Job Postings}_{f,t})$ and $\log(\text{Total Resume Employment}_{f,t})$

Sector × size AI utilization rates from resume data align with Census survey data

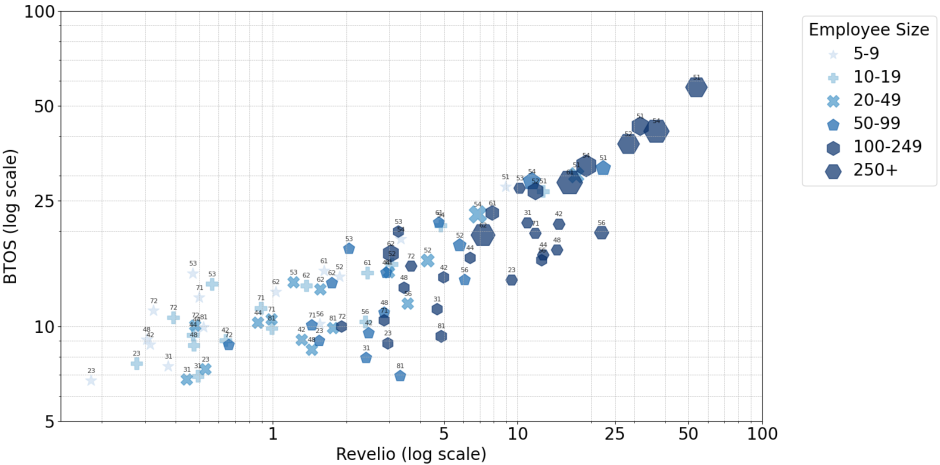

Data from firm-level Census Business Trends and Outlook Survey (BTOS), also analyzed by Bonney et al (2024) Back

Percent of firms w/ at least 1 AI-tagged position, by major NAICS sector×size

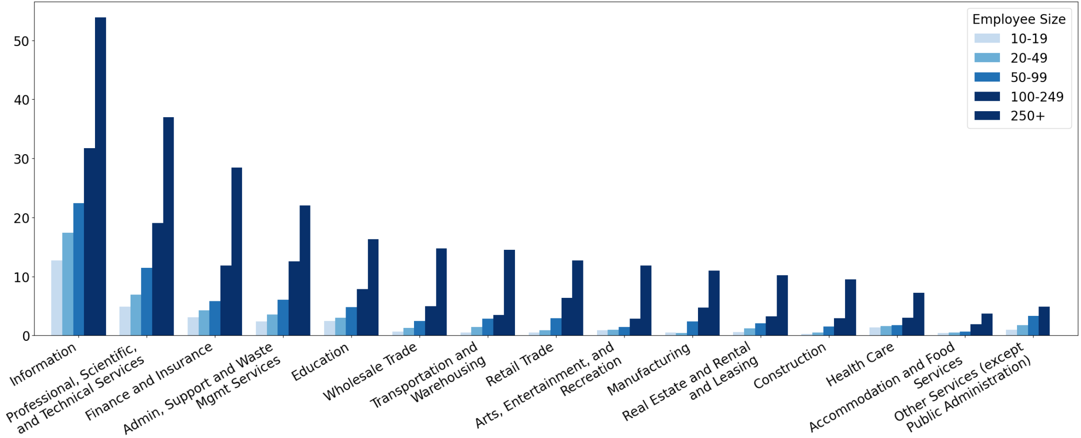

Back

Resume-based AI workers and AI-related patenting

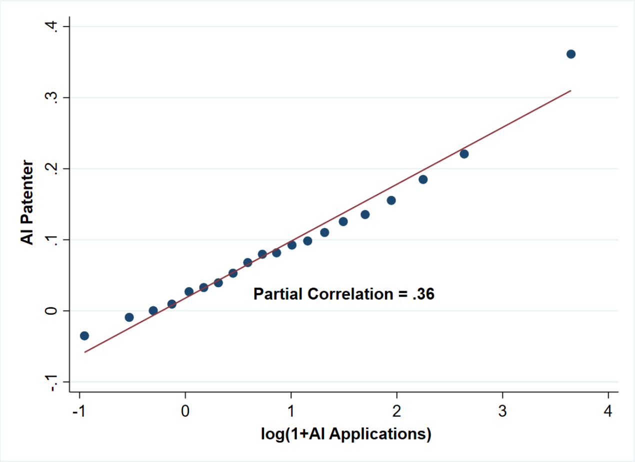

Residualized binscatter of indicator for AI patenting status against $\log(1 + \text{AI Applications}_{f,t})$:

Controls:

$\log(\text{Total Resume Employment}_{f,t})$ and non-AI patenting indicator Back

Comparison with AI employees in Babina et al (2024)

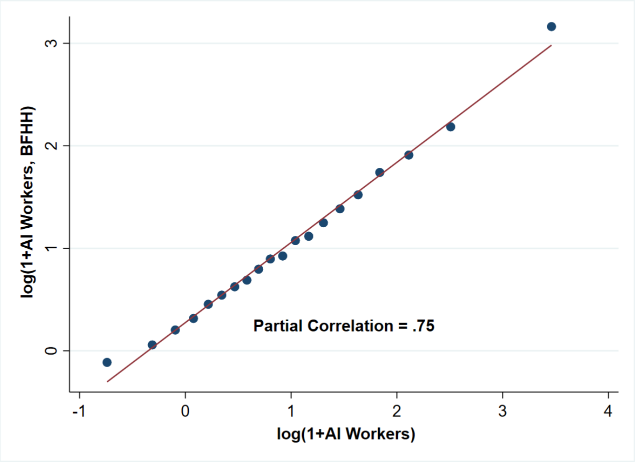

Babina et al (2024)—BFHH—count AI-related resumes with a slightly different but related method.

Plotted: Residualized binscatter of $\log(1 + \text{AI Workers (BFHH)}_{f,t})$ against $\log(1 + \text{AI Workers}_{f,t})$:

Controls: $\log(\text{Total Resume Employment}_{f,t})$ and $\log(\text{Total Resume Employment (BFHH)}_{f,t})$

Back

Theoretical Framework

Aggregate output composite of individual firm output

$$\bar{Y} = \left(\int Y(f)^{\frac{\theta-1}{\theta}} df \right)^{\frac{\theta}{\theta-1}}$$

Firm composite output of individual occupations

$$Y(f) = \left(\int Y(o,f)^{\frac{\chi-1}{\chi}} do \right)^{\frac{\chi}{\chi-1}}$$

Parameters:

- θ: elasticity of substitution across firms (demand elasticity)
- χ: elasticity of substitution across occupations

Back

Occupation Output

Occupation output composite of different tasks

$$Y(o,f) = \left(\sum_j y(j)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}}$$

where tasks are produced by labor l and capital k

$$y(j) = \left(\gamma_j \, l(j)^{\frac{\nu-1}{\nu}} + (1-\gamma_j) \, k(j)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}}$$

Parameters:

- ψ: elasticity of substitution across tasks within an occupation
- ν: elasticity of substitution between labor input and capital (AI technology)

Occupation Labor Supply

A worker i in occupation o chooses hours $h(i,j)$ across tasks j

$$l(i,j) = h(i,j)^{1-\beta} \qquad \text{subject to} \qquad \sum_j h(i,j) = 1$$

Occupation Labor Supply

A worker i in occupation o chooses hours $h(i,j)$ across tasks j

$$l(i,j) = h(i,j)^{1-\beta} \qquad \text{subject to} \qquad \sum_j h(i,j) = 1$$

Workers' labor supply to occupation o function of occupational wage index:

$$L(w_o) = \bar{\zeta} \Bigg(\underbrace{\sum_j w(j) l(i,j)}_{W(o,f)} \Bigg)^{\zeta}$$

Microfoundation: occupation-specific taste shocks, as in Lamadon-Mogstad-Setzler (2022)

AI Technology Improvements

(AI) Capital is specific to task j

Improvements in AI technology $\Rightarrow$ decline in (quality-adjusted) price q_j:

$$\varepsilon_j \equiv -\Delta \log q_j$$

Impact on labor demand of a given technology $[\varepsilon_1 \ldots \varepsilon_J$ for that occupation?

1. Capital became better so may use more capital relative to labor in task j.
2. But, if only some tasks are affected, workers can shift their effort to other tasks which can increase their productivity.
3. In addition, if the firm becomes more productive overall, it may hire more workers even from the affected occupations.

Key elasticities

Elasticity of task-specific price $w(j)$ to own-task improvements $\varepsilon(j)$

$$\eta_o \equiv$$

Key elasticities

Elasticity of task-specific price $w(j)$ to own-task improvements $\varepsilon(j)$

$$\eta_o \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j)} = -\frac{s_k}{J}\,\frac{(\nu-\chi)(1-\beta)+}{\left(s_k\nu+s_l\chi+\zeta\right)\left(1-\beta\left(1-s_k\nu-\psi s_l\right)\right)},$$

- First term, likely negative, captures two forces: substitution between labor and capital vs across occupations (workers may become more productive).

Note: elasticities: ν capital-labor; ψ across tasks, within occ; χ across occ, within firm; θ across firm, β governs strength of hours reallocation

Key elasticities

Elasticity of task-specific price $w(j)$ to own-task improvements $\varepsilon(j)$

$$\eta_o \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j)} = -\frac{s_k}{J} \frac{(\nu-\chi)(1-\beta)+\beta\left((J-1)(\nu-\psi)\zeta+\nu(\psi-\chi)+J\,(\nu-\psi)(s_k\nu+s_l\chi)\right)}{\left(s_k\nu+s_l\chi+\zeta\right)\left(1-\beta\,(1-s_k\nu-\psi s_l)\right)},$$

- First term, likely negative, captures two forces: substitution between labor and capital vs across occupations (workers may become more productive).
- Second term captures the effect from reallocation across tasks

Note: elasticities: ν capital-labor; ψ across tasks, within occ; χ across occ, within firm; θ across firm, β governs strength of hours reallocation

Key elasticities

Elasticity of task-specific price $w(j)$ to own-task improvements $\varepsilon(j)$

$$\eta_o \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j)} = -\frac{s_k}{J} \frac{(\nu-\chi)(1-\beta) + \beta\left((J-1)(\nu-\psi)\zeta + \nu(\psi-\chi) + J(\nu-\psi)(s_k\nu + s_l\chi)\right)}{\left(s_k\nu + s_l\chi + \zeta\right)\left(1-\beta(1-s_k\nu - \psi s_l)\right)},$$

- First term, likely negative, captures two forces: substitution between labor and capital vs across occupations (workers may become more productive).
- Second term captures the effect from reallocation across tasks

In general, $\eta_o < 0$ if the number of tasks is sufficiently large

Note: elasticities: ν capital-labor; ψ across tasks, within occ; χ across occ, within firm; θ across firm, β governs strength of hours reallocation

Key elasticities

Key elasticities

Elasticity of task-specific price $w(j)$ to improvements in other tasks $\varepsilon(j')$ for $j' \neq j$

$$\eta_c \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j')} = -\frac{s_k}{J} \frac{(\nu-\chi)(1-\beta)-\beta\Big(\nu(\chi-\psi)+\zeta(\nu-\psi)\Big)}{\Big(s_k\nu+s_l\chi+\zeta\Big)\Big(1-\beta(1-s_k\nu-\psi s_l)\Big)}.$$

Note: elasticities: ν capital-labor; ψ across tasks, within occ; χ across occ, within firm; θ across firm, β governs strength of hours reallocation

Key elasticities

Elasticity of task-specific price $w(j)$ to improvements in other tasks $\varepsilon(j')$ for $j' \neq j$

$$\eta_c \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j')} = -\frac{s_k}{J} \frac{(\nu-\chi)(1-\beta)-\beta\Big(\nu(\chi-\psi)+\zeta(\nu-\psi)\Big)}{\Big(s_k\nu+s_l\chi+\zeta\Big)\Big(1-\beta(1-s_k\nu-\psi s_l)\Big)}.$$

Sensitivity to improvements in other occupations (spillovers):

$$\eta_z \equiv \frac{\partial \log w(j)}{\partial \log Z_f} = \frac{\theta-\chi}{s_k\nu+s_l\chi+\zeta}$$

Note: elasticities: ν capital-labor; ψ across tasks, within occ; χ across occ, within firm; θ across firm, β governs strength of hours reallocation

Allocation of hours

Improvements in AI technology that are specific to task j reduce the allocation of labor effort in directly affected tasks

$$\Delta \log h(j) \approx \underbrace{\frac{\eta_o - \eta_c}{\beta}}_{<0} (\varepsilon(j) - m(\varepsilon))$$

where $m(\varepsilon)$ is the mean technology improvement across all tasks j in occupation

$$m(\varepsilon) \equiv \frac{1}{J} \sum_{j \in J} \varepsilon(j)$$

Allocation of hours

Improvements in AI technology that are specific to task j reduce the allocation of labor effort in directly affected tasks

$$\Delta \log h(j) \approx \underbrace{\frac{\eta_o - \eta_c}{\beta}}_{<0} (\varepsilon(j) - m(\varepsilon))$$

where $m(\varepsilon)$ is the mean technology improvement across all tasks j in occupation

$$m(\varepsilon) \equiv \frac{1}{J} \sum_{j \in J} \varepsilon(j)$$

Consistent with the fact that demand for AI exposed tasks declines

Mean vs variance: example from Walmart

Distribution of AI exposure across tasks: customer service reps vs stockers and order fillers

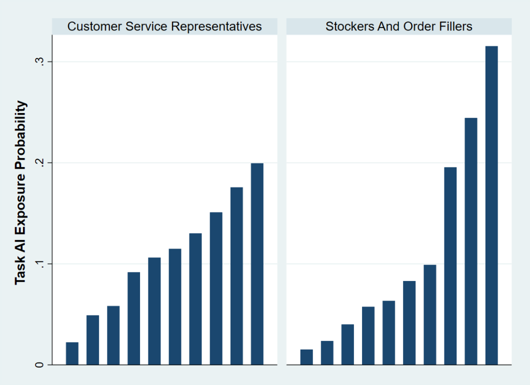

Most Exposed Tasks for Stockers and Order Fillers

Issue or distribute materials, products, parts, and supplies to customers or coworkers, based on information from incoming requisitions.

Answer customers' questions about merchandise and advise customers on merchandise selection.

Itemize and total customer merchandise selection at checkout counter, using cash register, and accept cash or charge card for purchases.

Least Exposed Tasks for Stockers and Order Fillers

Clean display cases, shelves, and aisles.

Operate equipment such as forklifts.

Complete order receipts.

These two occupations have similar mean but different variance exposure at Walmart

Back JP Morgan Example

Mean vs variance: example from JP Morgan Chase

Distribution of exposure prob. across tasks: credit analysts vs financial managers

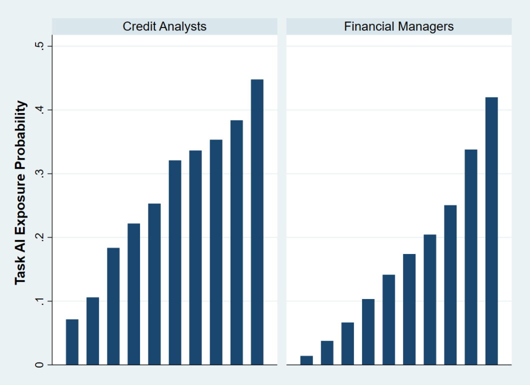

These two occupations have similar variance but different mean exposure at JPMC

Most Exposed Tasks for Financial Managers

Develop or analyze information to assess the current or future financial status of firms.

Analyze and classify risks and investments to determine their potential impacts on companies.

Analyze the financial details of past, present, and expected operations to identify development opportunities and areas where improvement is needed.

Least Exposed Tasks for Financial Managers

Direct insurance negotiations, select insurance brokers or carriers, and place insurance.

Compute, withhold, and account for all payroll deductions.

Approve, reject, or coordinate the approval or rejection of lines of credit or commercial, real estate, or personal loans.

Walmart Example: Back

Predicting AI employees

The shift-share IV is the predicted number of AI workers at the firm

$$\text{Predicted AI Employees}_{f,t} = \text{Employment}_{f,t} \times p^{AI}_{f,t}$$

and $p^{AI}_{f,t}$ is the predicted probability a given worker is an AI worker

Predicting AI employees

The shift-share IV is the predicted number of AI workers at the firm

$$\text{Predicted AI Employees}_{f,t} = \text{Employment}_{f,t} \times p_{f,t}^{AI}$$

and $p_{f,t}^{AI}$ is the predicted probability a given worker is an AI worker

$$p_{f,t}^{AI} = \sum_{u}$$

Back

Predicting AI employees

The shift-share IV is the predicted number of AI workers at the firm

$$\text{Predicted AI Employees}_{f,t} = \text{Employment}_{f,t} \times p^{AI}_{f,t}$$

and $p^{AI}_{f,t}$ is the predicted probability a given worker is an AI worker

$$p^{AI}_{f,t} = \sum_{u} \underbrace{w_{u \to f}}_{\substack{\text{share of} \\ \text{non-AI employees} \\ \text{coming from} \\ \text{university } u \text{ over} \\ \text{2005 to 2009}}}$$

Back

Predicting AI employees

The shift-share IV is the predicted number of AI workers at the firm

$$\text{Predicted AI Employees}_{f,t} = \text{Employment}_{f,t} \times p^{AI}_{f,t}$$

and $p^{AI}_{f,t}$ is the predicted probability a given worker is an AI worker

$$p^{AI}_{f,t} = \sum_u \underbrace{w_{u \to f}}_{\substack{\text{share of} \\ \text{non-AI employees} \\ \text{coming from} \\ \text{university } u \text{ over} \\ \text{2005 to 2009}}} \times \underbrace{\frac{N^{AI}_{u,t}}{N^{total}_{u,t}}}_{\substack{\text{fraction of} \\ \text{university } u \text{ grads} \\ \text{who work in AI}}}$$

IV strategy: Firms are more likely to hire AI workers if they were previously connected to universities whose graduates are more likely to do AI today. Details

Back

IV relevance tests

Lagged University Shares Predict Future Shares:

	(1) Average Share (2014-2018)
Average Share (2005-2009)	0.482***
	(26.98)
N	864502
R-sq (within)	0.118
Firm FE	X
University FE	X

Variation: university $\times$ firm

Shift-Share Predicts Firm AI Worker Share:

	(1) Actual AI Worker Share
Predicted AI Worker Share	0.582***
	(7.46)
N	16560
R-sq (within)	0.0436
Revelio Emp Control	X
Ind $\times$ Year FE	X

Variation: firm $\times$ year

Back

Back-Across Firms

IV relevance tests

Lagged University Shares Predict Future Shares:

	(1)
	Average Share (2014-2018)
Average Share (2005-2009)	0.482***
	(26.98)
N	864502
R-sq (within)	0.118
Firm FE	X
University FE	X

Variation: university $\times$ firm

Shift-Share Predicts Firm AI Worker Share:

	(1)
	Actual AI Worker Share
Predicted AI Worker Share	0.582***
	(7.46)
N	16560
R-sq (within)	0.0436
Revelio Emp Control	X
Ind $\times$ Year FE	X

Variation: firm $\times$ year

Back

Effort shifts away from AI exposed tasks: OLS vs IV

Back

Dep. Variable: 100 × 5-year DHS growth	OLS			IV		
in share of job posting skills related to task	(1)	(2)	(3)	(4)	(5)	(6)
Task-level AI Exposure	-4.71***	-4.68***	-4.73***	-4.57***	-4.14***	-4.68***
	(-13.40)	(-13.91)	(-14.08)	(-9.52)	(-10.54)	(-11.48)
Observations (task–occ–firm–year)	13.2m	13.2m	13.2m	13.2m	13.2m	13.2m
F-stat				17071.9	25249.7	27488.8
Controls						
ONET Task Importance	X	X	X	X	X	X
Mean Occ Task Exposure	X	X		X	X	
Firm × Year FE	X	X		X	X	
Occ × Year FE		X			X	
Firm × Occ × Year FE			X			X

Impact of AI on employment across occupation groups

	2-digit SOC	Mean Component	Variance Component	Firm Component	Total	% of Emp
Management	11	-2.27	1.55	0.78	0.057	19.0
Business and Financial	13	-10.1	6.18	2.04	-1.92	17.6
Architecture and Engineering	17	-5.96	2.82	0.51	-2.63	9.10
Science	19	1.60	-0.018	0.10	1.68	2.36
Community and Social Service	21	10.8	-5.76	0.30	5.32	0.33
Legal	23	10.0	-6.17	2.56	6.42	0.71
Education and Library	25	9.47	-5.03	0.072	4.51	1.00
Arts, Entertainment, Media	27	7.99	-4.82	2.09	5.26	5.38
Healthcare Practitioners	29	5.77	-2.63	-0.54	2.60	1.92
Healthcare Support	31	7.59	-3.95	0.42	4.06	0.47
Protective Service	33	9.37	-5.87	-1.46	2.05	0.43
Food Preparation and Serving	35	12.7	-7.02	-7.70	-1.99	2.75
Cleaning and Maintenance	37	14.5	-8.80	-3.37	2.28	0.46
Personal Care and Service	39	12.5	-6.81	-3.66	1.98	1.09
Sales and Related	41	1.47	-0.73	-1.60	-0.86	13.3
Office and Administrative	43	2.71	-2.45	0.61	0.87	10.6
Farming, Fishing, and Forestry	45	13.4	-7.76	-3.79	1.81	0.46
Construction and Extraction	47	6.41	-4.30	-0.44	1.67	2.07
Installation and Repair	49	4.03	-3.33	-0.99	-0.29	2.72
Production	51	5.80	-2.58	-2.40	0.82	3.94
Transportation	53	7.92	-4.47	-2.57	0.88	4.26

We find wide dispersion across detailed (6-digit) occupations Back

IV Robustness

Drop elite universities

Firm-level, Firm-Occ level

Drop top employers

Firm-level, Firm-Occ level

Exclude technology firms

Firm-level, Firm-Occ level

Add controls for trends in CS and engineering labor demand

Firm-level, Firm-Occ level

Back-Firm Outcomes Back Firm-Occ Outcomes

Firm Outcomes and IV: Dropping elite universities, firms, and tech

Exclude (1) top 50 universities by total AI grads in post-period (includes nearly all Ivy leagues+); (2) the top 50 firms (by emp of AI workers); (3) tech industries

	IV (Drop Top 50 AI Firms/Universities+Tech Industry)			
	(1)	(2)	(3)	(4)
	Sales	Emp	Profit	TFP
$\log(1+\text{AI uses})$	8.21*	7.36**	8.17*	4.55*
	(2.56)	(3.04)	(2.48)	(2.46)
N	9458	9847	8507	4256
R-sq	0.084	0.050	0.034	0.17
Controls	X	X	X	X
Ind × Year FE	X	X	X	X

Impact of AI Exposure on Firm Outcomes: 5-year growth rates (in log p.p.)

$$\log(Y_{f,t+5}) - \log(Y_{f,t}) = \gamma \log(1+\text{AI uses})_{f,t} + \beta X_{f,t} + \varepsilon_{f,t}$$

Back-Robustness

Firm-Occ Outcomes and IV: Dropping elite universities, firms, and tech

Exclude (1) top 50 universities by total AI grads in post-period (includes nearly all Ivy leagues+); (2) the top 50 firms (by emp of AI workers); (3) tech industries

	Panel A: IV (Drop Univ/Firm/Tech)			
	(1)	(2)	(3)	(4)
AI Exposure Average	-17.0***	-14.7***	-16.1***	-8.98***
	(-7.23)	(-8.01)	(-11.56)	(-6.54)
AI Exposure Concentration	10.2***	6.19***	7.07***	4.89**
	(6.01)	(3.94)	(5.87)	(2.95)
$\log(1+\text{AI uses})$	9.56***	5.93***		
	(4.84)	(3.81)		
N	1636223	1636223	1635606	1635606
R^2	0.036	0.043	-0.0032	-0.0026
F-stat (AI Exposure Average)	1058.5	1143.0	2456.7	1335.2
F-stat (AI Exposure Concentration)	821.9	913.1	2293.9	666.6
F-stat ($\log(1+\text{AI uses})$)	915.8	905.7		
Controls	X	X	X	X
Year FE	X			
Industry × Year FE		X		
Firm × Year FE			X	X
Occ × Year FE				X
Drop Firm/Univ/Tech	X	X	X	X
Shift-Share Controls				

Back-Robustness

Firm Outcomes and IV: Control for predicted growth in CS/Eng

Add shift-share controls for predicted share of employees in computer science and engineering occupations

	IV (Add shift-share controls)			
	(1)	(2)	(3)	(4)
	Sales	Emp	Profit	TFP
$\log(1+\text{AI uses})$	9.51^{***}	6.54^{***}	8.21^{**}	7.60^{***}
	(3.82)	(3.60)	(3.21)	(5.17)
N	12282	12688	11246	6035
R-sq	0.070	0.051	0.027	0.18
Controls	X	X	X	X
Shift-Share Controls	X	X	X	X
Ind $\times$ Year FE	X	X	X	X

Impact of AI Exposure on Firm Outcomes: 5-year growth rates (in log p.p.)

$$\log(Y_{f,t+5}) - \log(Y_{f,t}) = \gamma \log(1+\text{AI uses})_{f,t} + \beta X_{f,t} + \varepsilon_{f,t}$$

Back-Robustness

Firm-Occ Outcomes and IV: Control for predicted growth in CS/Eng

Add shift-share controls for predicted share of employees in computer science and engineering occupations

	Panel B: IV (Shift-Share Controls)			
	(1)	(2)	(3)	(4)
AI Exposure Average	-14.2***	-12.9***	-14.8***	-10.3***
	(-5.50)	(-5.70)	(-9.10)	(-10.16)
AI Exposure Concentration	12.7***	12.1***	13.4***	9.16***
	(5.03)	(4.95)	(7.32)	(6.77)
$\log(1+\text{AI uses})$	17.6***	16.6***		
	(12.71)	(11.78)		
N	2017628	2017628	2016990	2016990
R^2	0.040	0.034	-0.029	-0.011
F-stat (AI Exposure Average)	465.0	453.9	1136.7	1394.9
F-stat (AI Exposure Concentration)	157.4	123.3	332.6	501.0
F-stat ($\log(1+\text{AI uses})$)	2074.3	2196.1		
Controls	X	X	X	X
Year FE	X			
Industry $\times$ Year FE		X		
Firm $\times$ Year FE			X	X
Occ $\times$ Year FE				X
Drop Firm/Univ/Tech				
Shift-Share Controls	X	X	X	X

Back-Robustness

Firm Growth Rate Regressions (AI Users Only)

	OLS				IV			
	(1) Sales	(2) Emp	(3) Profit	(4) TFP	(5) Sales	(6) Emp	(7) Profit	(8) TFP
log(AI uses)	1.85	0.54	6.69**	5.48***	19.6**	14.4*	28.3***	19.7***
	(0.86)	(0.28)	(2.84)	(3.54)	(2.96)	(2.51)	(3.56)	(4.32)
N	4602	4617	4411	2879	4578	4588	4400	2879
R-sq	0.22	0.22	0.23	0.29	0.038	0.054	-0.014	0.079
F-stat					97.5	102.3	87.5	97.7
Controls	X	X	X	X	X	X	X	X
Ind × Year FE	X	X	X	X	X	X	X	X

Impact of AI Exposure on Firm Outcomes: 5-year growth rates (in log p.p.)

$$\log(Y_{f,t+5}) - \log(Y_{f,t}) = \gamma \log(\text{AI uses})_{f,t} + \beta X_{f,t} + \varepsilon_{f,t}$$

Back

NLP Resources

Our embeddings model of choice are the gte-large embeddings.

We use the https://deepinfra.com/meta-llama/Meta-Llama-3.1-70B-Instruct to tag AI applications in resumes

We access these models using an API from DeepInfra

First Llama LLM query

Your current task is to review the following descriptions of job duties being performed by employees of the same company and summarize each of the applications of AI that you see being performed. The goal is to produce an itemized list, where each item corresponds with a different use case for artificial intelligence methods being described. For each application, please describe, in a few sentences based ONLY on the resume descriptions, what functions AI tools are being applied to perform (it is important not to make predictions unless a use case is described in the text). Your answers should be focused on which tasks these AI tools are being used to perform, rather than on which tools are being used. In other words, I only want you to summarize instances in which these employees describe using AI to perform a specific function or solve a particular problem. I am looking for descriptions of the tasks and functions that *the AI tools themselves are performing*, rather than just the responsibilities or activities of the employees who are working with those tools.

To organize your efforts, I suggest you follow a four-step process. In the first step, please filter out descriptions of tasks which are unrelated to applications of artificial intelligence. If a description does not refer to how an artificial intelligence method is being used (e.g., because it describes development of hardware or other infrastructure related to AI deployment), please disregard the information. In the second step, produce your temporary itemized list from the filtered text. Now let's start the third step: Think aloud. Please audit your answers according to the original text. Sometimes, a task is clearly AI-related, but the specific application is not really specified. An example would be an employee mentioning that they are maintaining data infrastructure or deploying algorithms without saying anything about which data they are using or what the purpose of the underlying algorithms are. When reviewing your preliminary set of bullets, feel free to discard items which fall into this category of not specifying an actual application. For fourth step, please provide your final answer to improve your previous answers. Before finalizing your answer, please also reread the original body of text and identify any additional applications, if any, which were not included in the original list. Extract key applications from the following text document. Please output ONLY as a JSON list (Do not include "“" and anything else). The JSON should represent a table with three columns:

(1) The first column, labeled 'Key Application', should contain concise summaries or key insights extracted from the text.

(2) The second column, labeled 'Raw Excerpt', should include the corresponding raw excerpts from the text that support each key point.

(3) The third column, labeled 'Final Answer', should include your final answer.

TEXT TO REVIEW

Follow-on Llama LLM query (further filtering and cleaning step 1 responses)

The excerpt below describes how an artificial intelligence technology is being applied. Assume that it is already known that the excerpt refers to a use of artificial intelligence; the reader only wants to know the specific final application. Therefore, all references to any type of AI tool (e.g. natural language processing, machine learning, computer vision, generative AI, or any specific AI/ML algorithm) are redundant and should be stripped from the text. If the text only contains reference to an AI tool and without a clearly specified application, you should return 'N/A' when you filter the text.

For reference, here are a few examples of correctly applied filters:

-'AI tools are being used to measure text similarity in educational settings using NLP' should become 'Measure text similarity in educational settings'

-'Machine learning is being applied to perform tasks related to database analysis and firmware/software development for embedded environments' should become 'Perform tasks related to database analysis and firmware/software development for embedded environments'

-'AI-powered chatbots are being used to provide customers with quick solutions and answers using natural language processing capabilities.' should become 'Provide customers with quick solutions and answers.'

-'Analyzing customer reviews using NLP to understand customer needs and wants' should become 'Analyze customer reviews to understand customer needs and wants'

-'AI tool is being used to deploy computer vision model' should become 'N/A", because computer vision models themselves are an AI tool, and the exact use of computer vision is not specified.'

With this in mind, please filter the following excerpt describing an AI application. < final answer from Llama here >