Job Transformation, Specialization, and the Labor Market Effects of AI

Lukas Freund
Boston College

Lukas Mann
ASU

August 29, 2025

SITE Conference

Who will gain and who will lose from AI-induced task automation?

Who will gain and who will lose from AI-induced task automation?

- Much of the discourse on AI's impact on labor markets: job elimination

[Frey and Osborne, 2017; Susskind, 2020]

Who will gain and who will lose from AI-induced task automation?

- Much of the discourse on AI's impact on labor markets: job elimination

  [Frey and Osborne, 2017; Susskind, 2020]

- Our focus: **job transformation** $\sim$ shift in the task content of jobs

Job transformation: the case of weavers in the 19th century

Period	Preparatory tasks		Tasks while machine running							Tasks while power loom stopped							
	Prepare warp	Dress warp	Let off warp	Pick shuttle	Beat reed	Take up cloth	Adjust warp tension	Replace empty bobbin	Monitoring	Fix smashes	Adjust temples	Back up loom	Replace empty shuttle	Fix broken weft	Fix broken warp end	Remove cloth, cleaning	Replace warp
Handloom	●	●	●	●	●	●	●		●		●		●	●	●	●	●
Early power loom (∼1820)							●	●	●	●	●	●	●	●	●	●	●
1833							●	●	●	●		●	●	●	●	●	●
1883							○	●	●	●			●	●	●	●	○

Notes. ● = Task performed; ○ = Reduced frequency; Empty = Task not performed.
Based on Bessen (2012), who draws on the records of the Lawrence Company, MA.

Who will gain and who will lose from AI-induced task automation?

- Much of the discourse on AI's impact on labor markets: job elimination

  [Frey and Osborne, 2017; Susskind, 2020]

- Our focus: **job transformation** $\sim$ shift in the task content of jobs
  - weavers *[Bessen, 2012]* & machinists *[Bartel et al., 2007]*
  - systematic historical evidence *[Autor et al., 2003; Spitz-Oener, 2006; Atalay et al., 2006]*
  - seems to hold true for AI, too *[Bonney et al., 2024; Gathmann et al., 2024]*

Who will gain and who will lose from AI-induced task automation?

- Much of the discourse on AI's impact on labor markets: job elimination

  [Frey and Osborne, 2017; Susskind, 2020]

- Our focus: **job transformation** $\sim$ shift in the task content of jobs
  - weavers *[Bessen, 2012]* & machinists *[Bartel et al., 2007]*
  - systematic historical evidence *[Autor et al., 2003; Spitz-Oener, 2006; Atalay et al., 2006]*
  - seems to hold true for AI, too *[Bonney et al., 2024; Gathmann et al., 2024]*

- State-of-art **models** abstract from job transformation...

- ...as **measurement** is hard

  1. workers' portfolios of **task-specific skills**
  2. which **tasks** will be **automated**

Who will gain and who will lose from AI-induced task automation?

- Much of the discourse on AI's impact on labor markets: job elimination

  [Frey and Osborne, 2017; Susskind, 2020]

- Our focus: **job transformation** $\sim$ shift in the task content of jobs
  - weavers *[Bessen, 2012]* & machinists *[Bartel et al., 2007]*
  - systematic historical evidence *[Autor et al., 2003; Spitz-Oener, 2006; Atalay et al., 2006]*
  - seems to hold true for AI, too *[Bonney et al., 2024; Gathmann et al., 2024]*

- State-of-art **models** abstract from job transformation...

- ...as **measurement** is hard
  1. workers' portfolios of **task-specific skills**
  2. which **tasks** will be **automated**

This paper:
unify theory & measurement
to quantify how
AI-induced job transformation
will affect worker earnings

What we do

▸ Literature

1. **Theory**: propose task-based model with bundling + occupational choice

2. **Measurement:** estimate task-specific skills

3. **Application:** quantify LLM-induced job transformation effects

What we do

▸ Literature

1. **Theory**: propose task-based model with bundling + occupational choice
   - occupations bundle tasks, performed by workers or machines
   - workers have heterogeneous portfolios of task-specific skills, choose occ. & earn wage

   $\rightarrow$ characterize implications of task bundling

2. **Measurement:** estimate task-specific skills

3. **Application:** quantify LLM-induced job transformation effects

What we do

▸ Literature

1. **Theory**: propose task-based model with bundling + occupational choice
   - occupations bundle tasks, performed by workers or machines
   - workers have heterogeneous portfolios of task-specific skills, choose occ. & earn wage

   $\rightarrow$ characterize implications of task bundling

2. **Measurement:** estimate task-specific skills
   - LLMs: occupational task weights for 30+ tasks (clustering of $\sim$ 20,000 O*NET tasks)
   - NLSY: worker panel of occ. choices & wages

   $\rightarrow$ estimate skill distribution using model

3. **Application:** quantify LLM-induced job transformation effects

What we do

Literature

1. **Theory**: propose task-based model with bundling + occupational choice
   - occupations bundle tasks, performed by workers or machines
   - workers have heterogeneous portfolios of task-specific skills, choose occ. & earn wage

     $\rightarrow$ characterize implications of task bundling

2. **Measurement:** estimate task-specific skills
   - LLMs: occupational task weights for 30+ tasks (clustering of $\sim$ 20,000 O*NET tasks)
   - NLSY: worker panel of occ. choices & wages

     $\rightarrow$ estimate skill distribution using model

$\rightarrow$ map tasks
to exposure measures

3. **Application:** quantify LLM-induced job transformation effects
   - LLMs automate information-processing tasks *[Eloundou et al., 2023]*

LLM-driven automation of information- processing: big picture argument

1. **Occupation-level automation exposure $\not\Rightarrow$ adverse worker-level impacts**
   - large reallocation flows following AI automation $\rightarrow$ shifting worker composition
   - ambiguous relationship b/w exposure & average wage change at occupational level
   - winners and losers *within* occupation

2. Even absent job *elimination*, LLM automation of information-processing tasks creates **large and heterogeneous wage effects through job *transformation***

Theory

Environment: task-based production meets Roy

- Discrete time (t), repeated static model
- **Production technology:**
  - production is Cobb-Douglas over discrete task set $\mathcal{T}$
  - **occupation** $o \in \mathcal{O}$ **bundles tasks** with weights $\{\alpha_{o,\tau}\}_{\tau \in \mathcal{T}}$

economist, teacher, ...

analyzing data, moving objects, ...

Environment: task-based production meets Roy

- Discrete time (t), repeated static model
- **Production technology:**
  - production is Cobb-Douglas over discrete task set $\mathcal{T}$
  - **occupation** $o \in \mathcal{O}$ **bundles tasks** with weights $\{\alpha_{o,\tau}\}_{\tau \in \mathcal{T}}$

economist, teacher, ...

analyzing data, moving objects, ...

- **Firms:**
  - infinite supply of entrepreneurs who perfectly compete for a worker's labor
  - assign tasks ex-ante optimally to humans ($\rightarrow \mathcal{T}_l$) or machines w prod. $\{z_\tau\}_{\tau \in \mathcal{T}}$ ($\rightarrow \mathcal{T}_m$)
  - match with 1 worker, rent machines from inf. elastic capital market at exog. rate r

Environment: task-based production meets Roy

- Discrete time (t), repeated static model
- **Production technology:**
  - production is Cobb-Douglas over discrete task set $\mathcal{T}$
  - **occupation** $o \in \mathcal{O}$ **bundles tasks** with weights $\{\alpha_{o,\tau}\}_{\tau \in \mathcal{T}}$

economist, teacher, ...

analyzing data, moving objects, ...

- **Firms:**
  - infinite supply of entrepreneurs who perfectly compete for a worker's labor
  - assign tasks ex-ante optimally to humans ($\rightarrow \mathcal{T}_l$) or machines w prod. $\{z_\tau\}_{\tau \in \mathcal{T}}$ ($\rightarrow \mathcal{T}_m$)
  - match with 1 worker, rent machines from inf. elastic capital market at exog. rate r
- **Workers:**
  - log utility over consumption
  - heterogeneous, fixed **task-specific skills** $s_i = \{s_{i,\tau}\}_{\tau \in \mathcal{T}_l}$ where $\underset{\text{$|\mathcal{T}_l| \times 1$ vector}}{s_i} \sim \mathcal{N}(\bar{s}, \Sigma_s)$
  - period t: draw shocks, choose occupation o, match with entrepreneur, produce & earn

Firm's optimal production problem

- **Output** of firm in occ o with worker i given idiosyncratic shock $\varepsilon_{i,t} \sim \mathcal{N}(0, \varrho)$:

$$y_{i,o,t}(\cdot) = \underbrace{\prod_{\tau \in \mathcal{T}_l} \left(\exp\left(s_{i,\tau} + \varepsilon_{i,t}\right) \cdot \ell_{i,\tau,t}\right)^{\alpha_{o,\tau}}}_{\text{worker-produced}} \underbrace{\prod_{\tau \in \mathcal{T}_m} \left(\exp\left(z_\tau\right) \cdot m_{i,\tau,t}\right)^{\alpha_{o,\tau}}}_{\text{machine-produced}}$$

- **Profits:**

$$\pi_{i,o,t} = \max_{\{m_{i,\tau}\}_{\tau \in \mathcal{T}_m}, \{\ell_{i,\tau}\}_{\tau \in \mathcal{T}_l}} y_{i,o,t}\left(\{\ell_{i,\tau,t}\}_{\tau \in \mathcal{T}_l}, \{m_{i,\tau,t}\}_{\tau \in \mathcal{T}_m}\right) - \exp\left(w_{i,o,t}\right) - r \sum_{\tau \in \mathcal{T}_m} m_{i,\tau,t}$$

$$\text{s.t.} \sum_{\tau \in \mathcal{T}_l} \ell_{i,\tau,t} = 1$$

Firm's optimal production problem

- **Output** of firm in occ o with worker i given idiosyncratic shock $\varepsilon_{i,t} \sim \mathcal{N}(0, \varrho)$:

$$y_{i,o,t}(\cdot) = \underbrace{\prod_{\tau \in \mathcal{T}_l} \left(\exp\left(s_{i,\tau} + \varepsilon_{i,t}\right) \cdot \ell_{i,\tau,t}\right)^{\alpha_{o,\tau}}}_{\text{worker-produced}} \underbrace{\prod_{\tau \in \mathcal{T}_m} \left(\exp\left(z_{\tau}\right) \cdot m_{i,\tau,t}\right)^{\alpha_{o,\tau}}}_{\text{machine-produced}}$$

- **Profits:**

$$\pi_{i,o,t} = \max_{\{m_{i,\tau}\}_{\tau \in \mathcal{T}_m}, \{\ell_{i,\tau}\}_{\tau \in \mathcal{T}_l}} y_{i,o,t}\left(\{\ell_{i,\tau,t}\}_{\tau \in \mathcal{T}_l}, \{m_{i,\tau,t}\}_{\tau \in \mathcal{T}_m}\right) - \exp\left(w_{i,o,t}\right) - r \sum_{\tau \in \mathcal{T}_m} m_{i,\tau,t}$$

$$\text{s.t.} \sum_{\tau \in \mathcal{T}_l} \ell_{i,\tau,t} = 1$$

- **Optimality:**

FOC capital

$$\ell_{i,\tau,t} = \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}}$$

Firm's optimal production problem

- **Output** of firm in occ o with worker i given idiosyncratic shock $\varepsilon_{i,t} \sim \mathcal{N}(0, \varrho)$:

$$y_{i,o,t}(\cdot) = \underbrace{\prod_{\tau \in \mathcal{T}_l} \left(\exp\left(s_{i,\tau} + \varepsilon_{i,t}\right) \cdot \ell_{i,\tau,t}\right)^{\alpha_{o,\tau}}}_{\text{worker-produced}} \underbrace{\prod_{\tau \in \mathcal{T}_m} \left(\exp\left(z_\tau\right) \cdot m_{i,\tau,t}\right)^{\alpha_{o,\tau}}}_{\text{machine-produced}}$$

- **Profits:**

$$\pi_{i,o,t} = \max_{\{m_{i,\tau}\}_{\tau \in \mathcal{T}_m}, \{\ell_{i,\tau}\}_{\tau \in \mathcal{T}_l}} y_{i,o,t}\left(\{\ell_{i,\tau,t}\}_{\tau \in \mathcal{T}_l}, \{m_{i,\tau,t}\}_{\tau \in \mathcal{T}_m}\right) - \exp\left(w_{i,o,t}\right) - r \sum_{\tau \in \mathcal{T}_m} m_{i,\tau,t}$$

$$\text{s.t.} \sum_{\tau \in \mathcal{T}_l} \ell_{i,\tau,t} = 1$$

- **Optimality:**

FOC capital

$$\ell_{i,\tau,t} = \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} \longrightarrow \textit{matrix A: } |\mathcal{O}| \times |\mathcal{T}_l|$$

Occupational task-weight matrix

Remark: Task-weight matrix.

The matrix A summarizes the relative weights attached to each task $\tau \in \mathcal{T}_l$ across occupations $o \in \mathcal{O}$:

$$A = \begin{pmatrix} \frac{\alpha_{1,1}}{LS_1} & \frac{\alpha_{1,2}}{LS_1} & \cdots & \frac{\alpha_{1,n_{\text{skill}}}}{LS_1} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\alpha_{n_{\text{occ}},1}}{LS_{n_{\text{occ}}}} & \frac{\alpha_{n_{\text{occ}},2}}{LS_{n_{\text{occ}}}} & \cdots & \frac{\alpha_{n_{\text{occ}},n_{\text{skill}}}}{LS_{n_{\text{occ}}}} \end{pmatrix} \in \mathbb{R}^{|\mathcal{O}| \times |\mathcal{T}_l|}$$

where $\text{LS}_o = \sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}$ denotes the labor share in occupation o.
The row vector $A_o := A_{o,\cdot}$ contains the task weights for occupation o.

Wage equation

▸ Intercept term

$$w_{i,o,t} = \overbrace{\mu_o}^{\text{occ.-specific intercept}} + \overbrace{\sum_{\mathcal{T}_l} \frac{\alpha_{o,\tau}}{LS_o} \cdot s_{i,\tau}}^{\text{weighted skills}} + \overbrace{\varepsilon_{i,t}}^{\text{idiosyncratic productivity shock}}$$

$$= \mu_o + \underbrace{\frac{1}{n_{\text{skill}}} \sum_{\mathcal{T}_l} s_{i,\tau}}_{\text{scalar absolute advantage}} + \text{Cov}\left(n_{\text{skill}} \cdot \frac{\alpha_{o,\cdot}}{LS_o}, \underbrace{s_{i,\cdot} - \frac{1}{n_{\text{skill}}} \sum_{\mathcal{T}_l} s_{i,\tau}}_{\text{specialization vector}} \right) + \varepsilon_{i,t}$$

Occupational choice

- Each period, worker i chooses occ. subject to preference shock $u_{i,o,t} \sim \text{Gumbel}(0, \nu)$:

$$\hat{o}_{i,t} = \text{argmax}_o w_{i,o,t} + u_{i,o,t}$$

- **Occupational choice probabilities:**

$$P(\hat{o} = o | w_{i,\cdot,t}) = \frac{\exp(w_{i,o,t}/\nu)}{\sum_{o'} \exp(w_{i,o',t}/\nu)}$$

- No exogenous switching costs

Automation in the model

- **Automation** of task τ^*: a one-time, permanent rise in machine productivity $z_{\tau^\star}$ that is large enough to make it optimal to reassign τ^* from humans to machines

$$\mathcal{T}_l' = \mathcal{T}_l \backslash \tau^\star \qquad\qquad \mathcal{T}_m' = \mathcal{T}_m \cup \tau^\star$$

- **Job transformation:** weight on $\tau^* \downarrow$ & $\uparrow$ weight on all other entries proportional to their occupation-specific weight

$$\begin{aligned} A_o' - A_o &= \begin{pmatrix} \frac{\alpha_{o,1}}{LS_o'} \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} & \frac{\alpha_{o,2}}{LS_o'} \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} & \dots & -\frac{\alpha_{o,\tau^\star}}{LS_o} & \dots \end{pmatrix} \\ &= \frac{\alpha_{o,\tau^\star}}{LS_o} \times \begin{pmatrix} \frac{\alpha_{o,1}}{LS_o'} & \frac{\alpha_{o,2}}{LS_o'} & \dots & -1 & \dots \end{pmatrix} \end{aligned}$$

Wage effects of automation

Change in expected log (potential) wage for i in occupation o:

$$\mathbb{E}\left[w_{i,o,t+1} - w_{i,o,t}\right] = \Delta\mu_o + \underbrace{(A'_o - A_o)s_i}_{\text{job transformation effects}}$$

where

$$\Delta\mu_o = \underbrace{\frac{\alpha_{o,\tau^\star}}{LS_o - \alpha_{o,\tau^\star}}(z_{\tau^\star} - \log r + \mu_o)}_{\text{productivity \& displacement effect}}$$

Wage effects of automation

Change in expected log (potential) wage for i in occupation o:

$$\mathbb{E}\left[w_{i,o,t+1} - w_{i,o,t}\right] = \Delta\mu_o + \underbrace{(A'_o - A_o)s_i}_{\text{job transformation effects}}$$

where

$$\Delta\mu_o = \underbrace{\frac{\alpha_{o,\tau^\star}}{LS_o - \alpha_{o,\tau^\star}}(z_{\tau^\star} - \log r + \mu_o)}_{\text{productivity \& displacement effect}}$$

$$\mathbb{E}\left[w_{i,o,t+1} - w_{i,o,t}\right] = \Delta\mu_o + \underbrace{\frac{\alpha_{o,\tau^\star}}{LS_o}}_{\text{occupational exposure}}\left(\sum_{\mathcal{T}_l \setminus \tau^\star} \underbrace{\frac{\alpha_{o,\tau}}{LS_o - \alpha_{o,\tau^\star}} s_{i,\tau} - s_{i,\tau^\star}}_{\text{relative specialization}}\right)$$

The role of task bundling

Remark: Task bundling

An occupation features **task-bundling** if

$$|\{\tau \in \mathcal{T}_l : \alpha_{o,\tau} > 0\}| > 1.$$

The economy features a **no-bundling property** if no occupation features task-bundling:

$$|\{\tau \in \mathcal{T}_l : \alpha_{o,\tau} > 0\}| = 1 \quad \forall o \in \mathcal{O}.$$

$\Rightarrow$ In a no-bundling economy, wage changes are solely driven by $\Delta\mu_o$

$\Rightarrow$ With task bundling, wages also change due to **job transformation**

Decomposition of occupation-level wage changes

▸ Approximation

Remark: Decomposition

$$
\begin{aligned}
&\mathbb{E}[w'_o|\hat{o}' = o] - \mathbb{E}[w_o|\hat{o} = o] \\
&= \overbrace{\mathbb{E}[w'_o|\hat{o} = o] - \mathbb{E}[w_o|\hat{o} = o]}^{\Delta w_o \text{ of incumbents}} + \overbrace{\mathbb{E}[w'_o|\hat{o}' = o] - \mathbb{E}[w'_o|\hat{o} = o]}^{\text{re-sorting}} \\
&= \underbrace{\Delta\mu_o}_{\text{productivity and displacement}} + \underbrace{(A'_o - A_o)\cdot\bar{s}}_{\text{task shift}} + \underbrace{(A'_o - A_o)(\bar{s}_{|o} - \bar{s})}_{\text{selection}} \\
&\quad + \underbrace{\mathbb{E}[w'_o|\hat{o}' = o] - \mathbb{E}[w'_o|\hat{o} = o]}_{\text{re-sorting}}
\end{aligned}
$$

Measurement

Theory meets data: overview

Why not O*NET GWAs/importance weights

- **Goals:** parametrize the model at same 'resolution' as task-exposure measures
- **Step 1:** map model tasks & occupations to data, construct A
  - O*NET: $\sim$ 19,000 task statements ($\sim$ most exposure measures) $\rightarrow$ *cluster* them
  - occupations: 90+ SOC-2000 minor groups ($\sim$ 3d)
- **Step 2:** estimate unobserved skill distribution $(\bar{s}, \Sigma_s)$ using MLE
  - given A + NLSY '79 + model structure

Step 1: constructing the task-weight matrix *A*

▸ Validation ▸ Examples

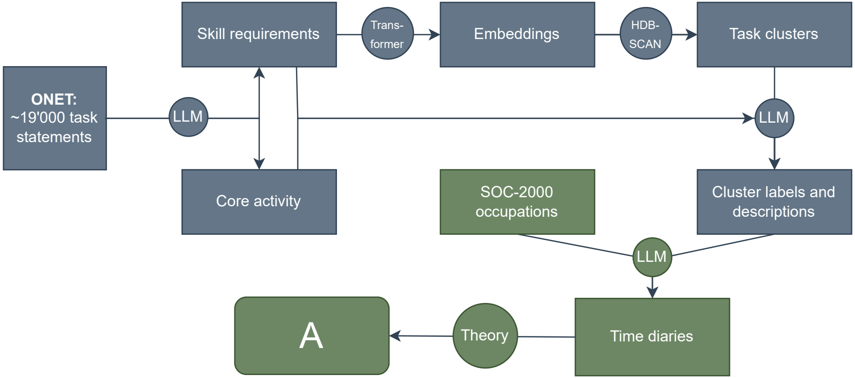

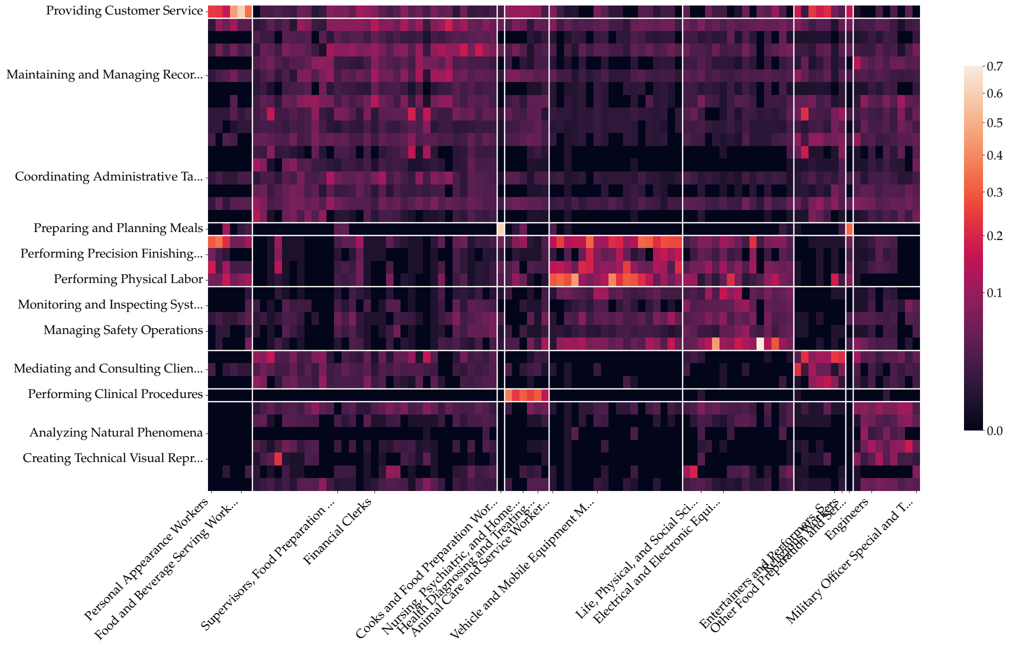
Providing Customer Service
Maintaining and Managing Recor...
Coordinating Administrative Ta...
Preparing and Planning Meals
Performing Precision Finishing...
Performing Physical Labor
Monitoring and Inspecting Syst...
Managing Safety Operations
Mediating and Consulting Clien...
Performing Clinical Procedures
Analyzing Natural Phenomena
Creating Technical Visual Repr...
Personal Appearance Workers
Food and Beverage Serving Work...
Supervisors, Food Preparation ...
Financial Clerks
Cooks and Food Preparation Wor...
Nursing, Psychiatric, and Home...
Health Diagnosing and Treating...
Animal Care and Service Worker...
Vehicle and Mobile Equipment M...
Life, Physical, and Social Sci...
Electrical and Electronic Equi...
Entertainers and Performers, S...
Other Food Preparation and Ser...
Religious Workers
Engineers
Military Officer Special and T...
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0

Step 2: estimation of task-specific skills

- **Measurement challenge #1:** skill distribution is unobserved
- **Solution:** use the structure of the model to estimate $(\bar{s}, \Sigma_s)$
  - variation: realized wages & occupational choices
  - intuition: economist vs software engineer
- **Data:** NLSY '79 + A matrix
  - worker-level panel of occupational choices and wages
- **Formalization:** max. likelihood
- **Implementation:** MC integration + auto-diff. + stochastic gradient descent ▸ Details
- **Validation:** Monte Carlo exercise

Validation: Monte-Carlo study

(a) Means

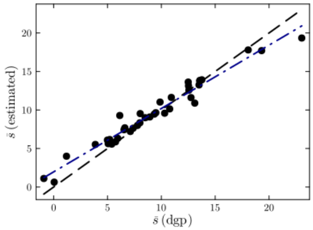

(b) Standard deviation

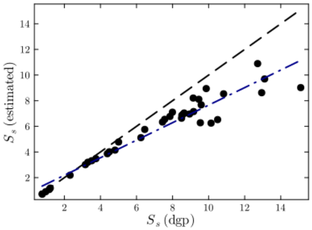

(c) Correlation

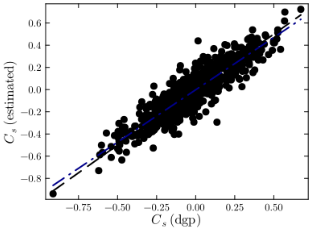

(d) Other parameters

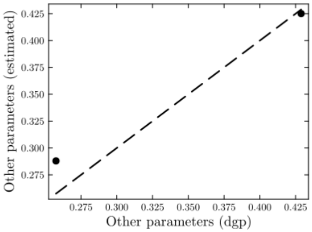

Estimated mean skills and dispersion

▸ Other parameters

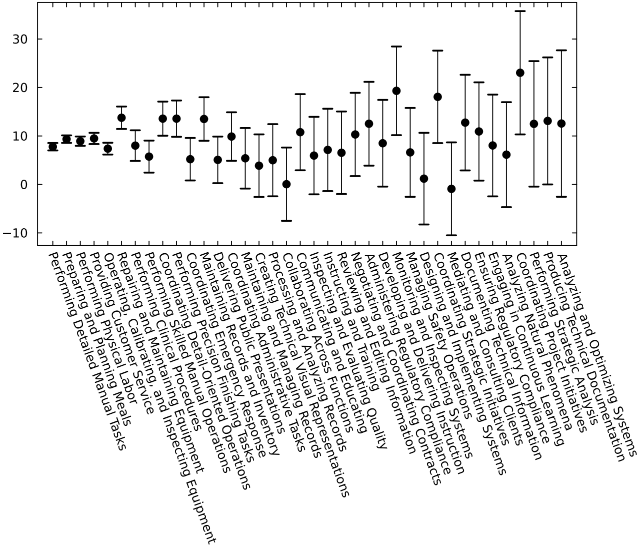

Selection based on comparative advantage

- Workers tend to select into occupations which load heavily on tasks they are relatively skilled at

relative specialization:

$s_{i,\cdot} - \frac{1}{n_{\text{skill}}} \sum_{\mathcal{T}_l} s_{i,\tau}$

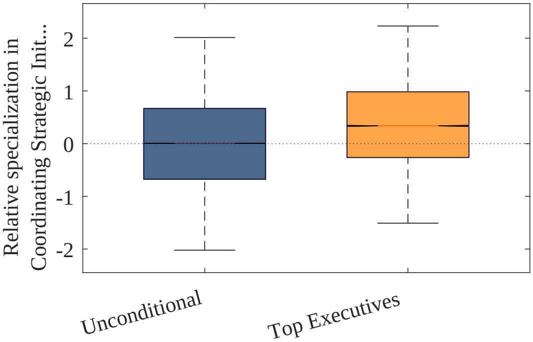

Model properties & validation

1. Wage variance decomposition
   - data: std. dev. 0.60, 28% between-occ. share
   - model: std. dev 0.70, 19% between-occ. share
2. Staying and switching probabilities Jump
3. Direction of moves driven by task requirements Jump
4. Frequency of moves shaped by specialization Jump

Model properties: occupational transition probabilities

Learning extension

- Some persistence (but not quite enough) — directionally tracks switching patterns

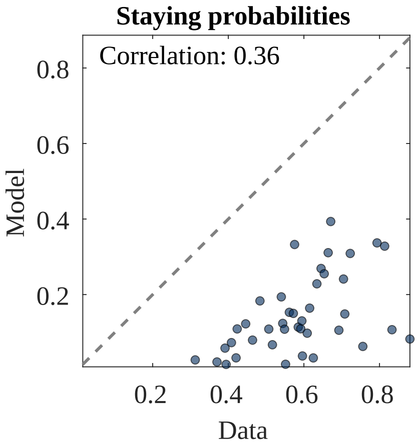

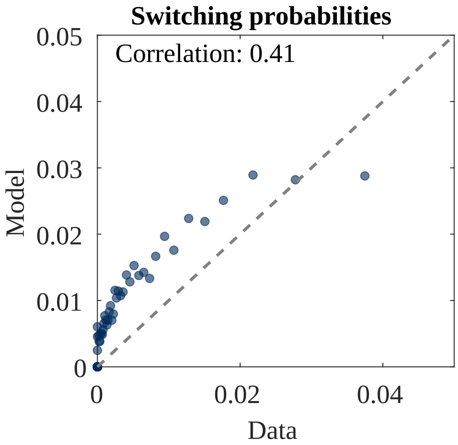

Model properties: task requirements and switching

- Workers are more likely to move to occupations with similar task requirements

[cf. Gathmann-Schoenberg, 2010]

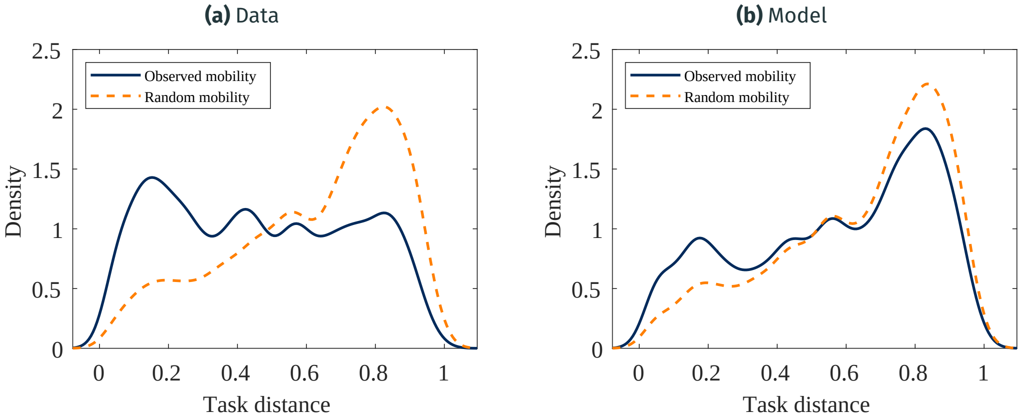

Model properties: specialization shapes switching frequency

- Evidence: skill specialization tends to generate persistence in occupational choice
  [Kambourov and Manovskii, 2008; Geel et al., 2011]

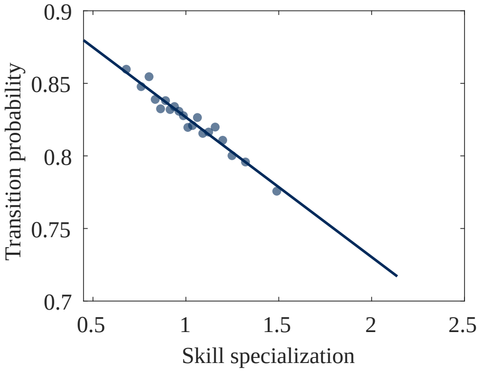

LLM-driven automation

Identifying task-specific automation shocks

- Scenario: full automation, with $z_{T^\star}$ at automation threshold –*just* productive enough...

- **Measurement challenge #2:** which specific tasks are being, or will be, automated?
    - forward-looking
    - labor share $\neq$ sufficient statistic when considering job transformation effects

- **Solution:** mapping of model to (clusters of) granular tasks that link directly to influential automation exposure measures *[Webb, 2019; Eloundou et al., 2023; Anthropic—Handa et al., 2024; ...]*

- **Focus on LLMs** using Eloundou et al. task-level measure
    - paper: industrial robots *[Webb et al., 2019]*

Aggregated task exposure measures from Eloundou et al. (2023)

▸ Webb (2020)

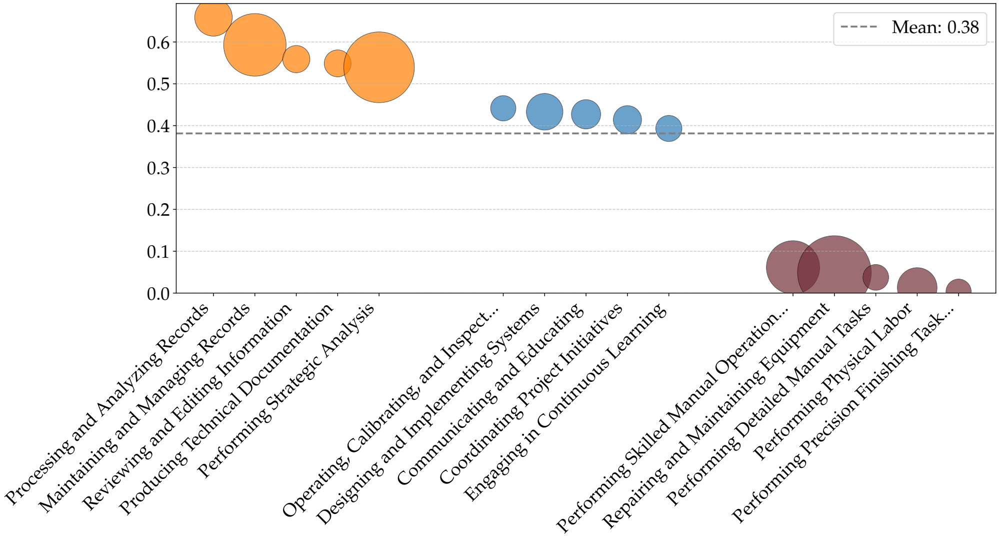

Occupation-level effects

⇒ More exposed occupations experience *larger* wage gains on average

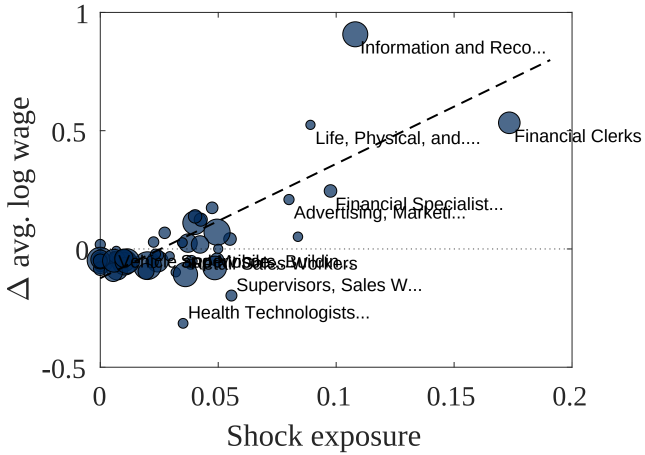

Decomposition: positive slope driven by task upgrading and resorting

$\Rightarrow$ This is b/c $\Delta\mu_o < 0$ is offset by positive task-shift & resorting effects

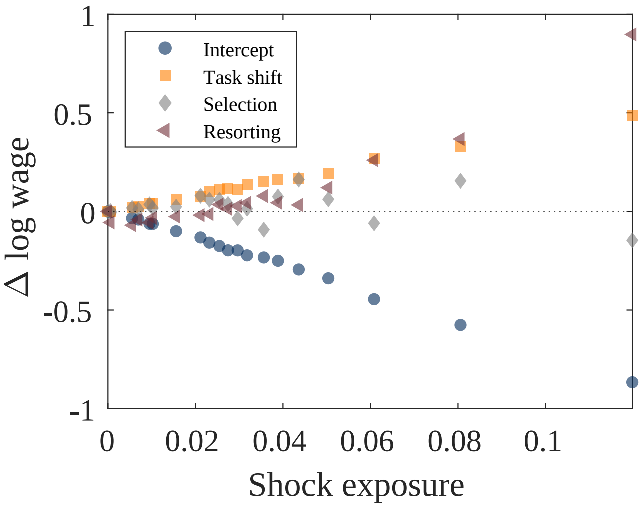

Resorting effect: comparison across tasks

$\Rightarrow$ AI-exposed tasks tend to be associated with larger skills dispersion $\rightarrow$ larger re-sorting wage effects $\rightarrow$ occupational averages provide worse guidance to worker-level outcomes

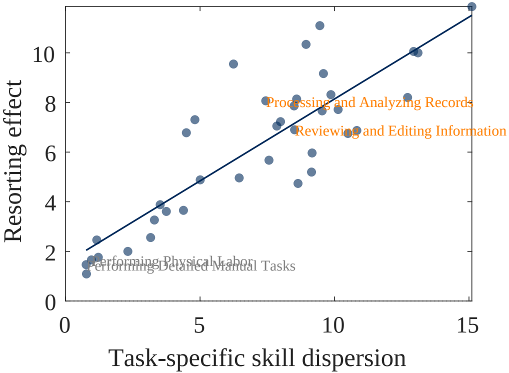

Individual-level effects for incumbents

⇒ Incumbent workers' wages in highly exposed origin occupations decline on average

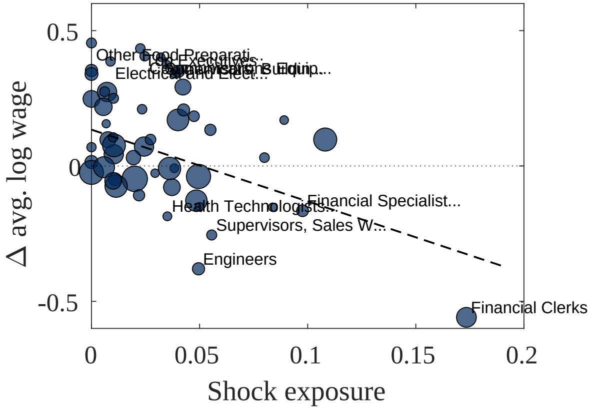

Heterogeneity among incumbents

▸ Relative specialization

⇒ Stayers win, incumbents lose (consistent with evidence on *task upgrading* for stayers *[Bartel et al., 2007; Dauth et al., 2021]* and losses for occupation switchers *[e.g. Huckfeldt, 2022]*)

(a) Incumbent stayers

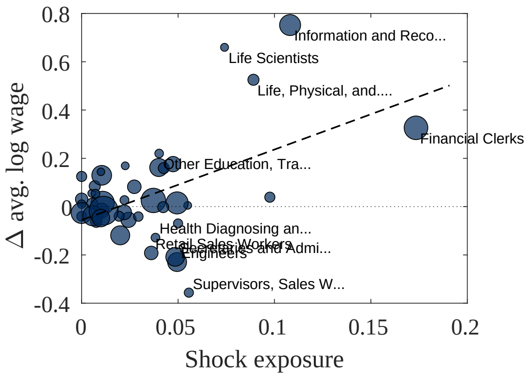

(b) Incumbent leavers

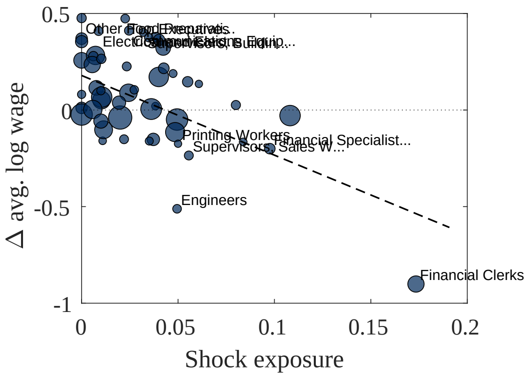

Explanation: selection

⇒ Leavers are, as a matter of selection, specialized in now-automated task

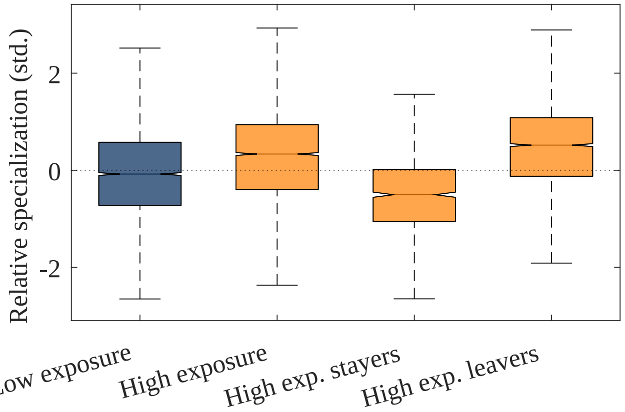

In-switchers experience large wage gains

⇒ Workers previously deterred from highly exposed occupations by skill barriers in now-automated tasks experience large

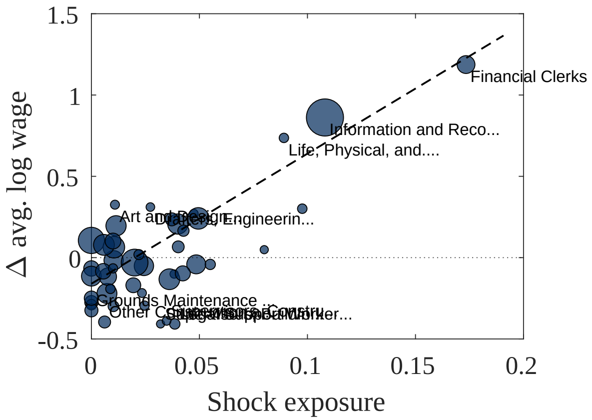

Recap of results: LLM-driven automation of information- processing tasks

- LLM-driven automation generates more occupational reallocation than in past
  - $\rightarrow$ occupation-level averages offer limited guidance for worker-level outcomes

— **Selection** on specialization generates neg. link b/w exposure & incumbent wages
  - $\rightarrow$ incumbent leavers specialized in information-processing tasks

\+ Automation **benefits** those reallocating time to tasks in which they're more skilled
  - $\rightarrow$ incumbent stayers who excel in customer-facing and coordination tasks

\+ *Or* enabled to access better occupations by **reducing skill-based entry barriers**
  - $\rightarrow$ in-switchers (think of "vibe coding")

Conclusion

Concluding remarks

- Just put out a first draft – **feedback** very welcome!
- **Core contribution:** empirically rich & tractable framework to quantify & forecast who wins and who loses from AI-induced **job transformation**
- **The big picture:**
  1. occupational exposure $\neq$ adverse individual wage effects
  2. absence of AI-induced job destruction $\neq$ absence of large labor market effects
- Planned work:
  - historical validation
  - will AI exacerbate wage inequality or might it, in fact, dampen it?

Extra Slides

What's new?

Intro

- **Measurement of job exposure to technologies** *[Brynjolfsson et al., 2018; Webb, 2019; Felten et al., 2021;* ***Eloundou et al., 2023****; Gathmann et al., 2024; Kogan et al., 2024]*

  $\Rightarrow$ map to **structural** model $\rightarrow$ individual **earnings effects** as a function of skills
- **Model-based analysis of AI** *[****Hampole et al., 2025****; Fan, 2025]*

  $\Rightarrow$ model with **bundling & skill heterogeneity** $\rightarrow$ quantify how job transformation affects heterogeneous worker's earnings
- **Task-based theory** *[Acemoglu-Autor, 2011; Acemoglu-Restrepo, 2018; Acemoglu-Restrepo, 2022; Freund, 2023;* ***Autor-Thompson, 2025****]*

  $\Rightarrow$ introduce task bundling $\rightarrow$ highlight automation effects due to Δ task content
- **Empirical literature on job transformation** *[****Autor et al., 2003****; Autor and Handel, 2013; Spitz-Oener, 2006; Atalay et al., 2020; Autor et al., 2024]*

  $\Rightarrow$ **link tasks with skills** $\rightarrow$ quantify *earnings* effects
- **Multi-dimensional skills** *[Lindenlaub, 2017; Lise-PostelVinay, 2021; Deming, 2023; Grigsby, 2023]*

  $\Rightarrow$ **estimate** distribution of high-dim. task-specific skills $\rightarrow$ **measure specialization**

AI capabilities are rapidly improving relative to humans

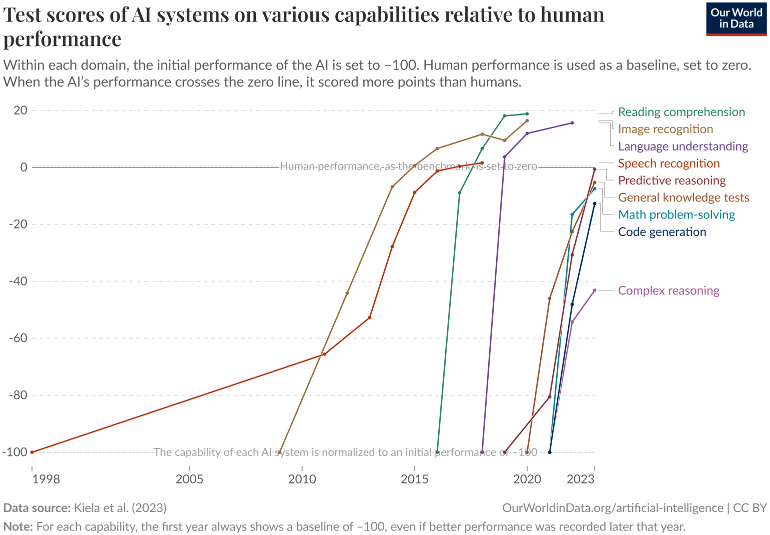

GE: plan

Back

- Missing important model feature: heterogeneous, endogenous occupation prices
  - steady-state: high-wage occ's involve scarce skills hence high o price
  - counterfactual: occupational price response as a function of demand elasticities

- Identification challenge: μ_o becomes endogenous and the following equation is satisfied by more than one pair $(\mu_o, \bar{s})$:

$$\text{mean potential wage}_o = \mu_o + A'_{o,\cdot} \circ \bar{s}$$

  where $\bar{s}$ is vector of average skills

- Options we're exploring:
  1. time variation in task shares
  2. dynamic skill accumulation
  3. identifying restriction $A \perp \mu_o$

FOCs

Back

- FOC for machines $m := \sum_{\tau \in \mathcal{T}_m} m_\tau$:

$$\left(\sum_{\tau \in \mathcal{T}_m} \alpha_{o,\tau}\right) \frac{y}{r} = m$$

and

$$m_\tau = \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_m} \alpha_{o,\tau}} m$$

- Given

$$\log y_o = \left[\sum_{\tau \in \mathcal{T}_l} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} s_{i,\tau}\right] + \varepsilon_{i,o}$$

$$+ \left[\sum_{\tau \in \mathcal{T}} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} \log(\alpha_{o,\tau})\right] - \log\left(\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}\right) + \left[\sum_{\tau \in \mathcal{T}_m} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} (z_\tau - \log r)\right],$$

Wage equation: details

Back

- Intercept

$$\mu_o = \sum_{\tau \in \mathcal{T}} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} \log\left(\alpha_{o,\tau}\right) + \left(\sum_{\tau \in \mathcal{T}_m} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} \left(z_\tau - \log r\right)\right)$$

- We assume that in the initial steady state there is only one composite machine task with productivity normalized to $\log r$, which implies that μ_o is known for all occupations.

Occupation-level decomposition: approximation

$$\mathbb{E}[w_o'|\hat{o}' = o] - \mathbb{E}[w_o|\hat{o} = o]$$

$$= \overbrace{\underbrace{\Delta\mu_o}_{\text{productivity and displacement}} + \underbrace{(A_o' - A_o)\cdot\bar{s}}_{\text{task shift}} + \underbrace{\nu^{-1}(A_o' - A_o)\Sigma\left(A_o^\intercal - \sum_{o''} h_{o''}(\bar{s}_{|o})A_{o''}^\intercal\right)}_{\text{selection}}}^{\Delta w_o \text{ of incumbents}}$$

$$+ \underbrace{\nu^{-1}A_o'\Sigma\left(\left((A_o' - A_o)^\intercal - \sum_{o''}\left(h_{o''}'(\bar{s}_{|o}')(A_{o''}')^\intercal - h_{o''}(\bar{s}_{|o})A_{o''}^\intercal\right)\right)\right)}_{\text{re-sorting}}. \tag{1}$$

where

$$\bar{s}_{|o} = \bar{s} + \nu^{-1}\Sigma \overbrace{\left(A_o^\intercal - \sum_{o''} h_{o''}(\bar{s}_{|o})A_{o''}^\intercal\right)}^{\text{relative task intensity of occupation } o} \tag{2}$$

$$h_o(s) = \frac{\exp(\nu^{-1}\mu_{o'} + \nu^{-1}A_{o'}\cdot s)}{\sum_{o''}\exp(\nu^{-1}\mu_{o''} + \nu^{-1}A_{o''}\cdot s)} \tag{3}$$

Examples of mapping from detailed tasks to clusters

Back

We cluster ~20k unstructured, detailed task statements into 38 task categories based on similarity of skill requirements

*~20k detailed O*NET task statements*

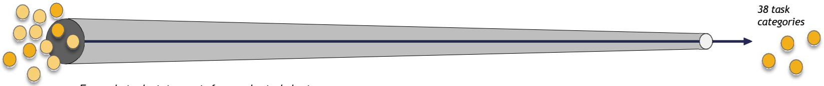

38 task categories

Example task statements from selected clusters

Example task categories

- Smooth rough spots on walls and ceilings, using sandpaper.
- Lubricate moving parts on gate-crossing mechanisms and swinging signals.
- Clean and polish vehicle windows.
- *Etc.*

Performing Detailed Manual Tasks

- Prepare reports of activities, evaluations, recommendations, or decisions.
- Prepare, examine, or analyze accounting records, financial statements, or other financial reports to assess accuracy, completeness, and conformance to reporting and procedural standards.
- Prepare and submit reports and charts to treatment team to reflect patients' reactions and evidence of progress or regression.
- *Etc.*

Processing and Analyzing Records

- Confer with officials of public health and law enforcement agencies to coordinate interdepartmental activities.
- Confer with directors and production staff to discuss issues, such as production and casting problems, budgets, policies, and news coverage.
- Plan and evaluate new projects, consulting with other engineers and corporate executives, as necessary.
- *Etc.*

Coordinating Project Initiatives

For each task, we extract skill requirements, create semantic vector embeddings for these requirements using a transformer model, and perform HDBSCAN-clustering on these embeddings to create broad task categories.

Details on the estimation strategy I

Back

- Exact likelihood:

$$\prod_i \int_s \left[\left(\int_{w_{i,\cdot,-\omega.}} \prod_t P(\hat{o}_{i,t} = \omega_{i,t} | w_{i,\cdot,\cdot}, \nu) \cdot f(w_{i,t,-\omega_t} | s, w_{i,\cdot,\omega.}, \varsigma) \right) \cdot f(s | w_{i,\cdot,\omega.}, \varsigma, \bar{s}, \Sigma_s) \right]$$
$$\cdot f(w_{i,\cdot,\omega.} | \varsigma, \bar{s}, \Sigma_s)$$

- **Strategy:** Monte Carlo integration - for all i generate n_o draws from

$$f(w_{i,\cdot,-\omega.} | w_{i,\cdot,\omega.}, \varsigma, \bar{s}, \Sigma_s) = \int_s f(w_{i,\cdot,-\omega.} | s, w_{i,\cdot,\omega.}, \varsigma) f(s | w_{i,\cdot,\omega.}, \varsigma, \bar{s}, \Sigma_s)$$

and evaluate the mean of $P(\hat{o}_{i,t} = \omega_{i,t} | w_{i,\cdot,t}, \nu)$ to obtain an estimator for $\mathcal{L}_i(\theta)$:

$$\hat{\mathcal{L}}_i(w_{i,t,\omega}, \nu, \varsigma, \bar{s}, \Sigma_s) = \left(\frac{1}{n_o} \sum_j \prod_t P(\hat{o}_{i,t} = \omega_{i,t} | w_{j,t,\cdot}, \nu) \right) \cdot f(w_{i,\cdot,\omega.} | \varsigma, \bar{s}, \Sigma_s)$$

Details on the estimation strategy II

Back

- Two numerical techniques help speed up the maximum likelihood computation
- **Auto-differentiation:** efficiently compute the gradient of this function
- **Stochastic gradient descent:**
  - basic technique: gradient descent

$$\theta_{t+1} = \theta_t - \eta \cdot \nabla \left(-\mathcal{L}(\theta_t)\right)$$

  - randomly partition individuals into n groups:

$$\{1, 2, \ldots, I\} = B_1 \cup B_2 \cup \ldots \cup B_n, \quad B_i \cap B_j = \emptyset$$

  - calculate the likelihood based on batch $B_1, \ldots, B_n$ only
  - when done, draw a new partition

Parameter estimates

Back

- For the scalar parameters, we estimate $\nu = 0.26$ and $\varrho = 0.43$.
- The estimate of ν implies that reducing prospective wages in a given occupation by 1% lowers the odds of choosing this occupation by about 3.8% since
- $\varrho = 0.43$ indicates that a one-standard-deviation occupation-specific random productivity shock can raise or lower wages by about 43% in a given year.

Why not use O*NET GWAs and importance weights

Back

- Potential alternative to our approach: use O*NET ”General Work Activities” (GWAs) and occupational importance weights

- Reasons we prefer our approach:
  1. GWAs themselves are not mutually exclusive (e.g. “Analyzing Data or Information” vs “Processing Information”) nor exhaustive (esp. regarding activities differentiating high-wage occupations, e.g. complex quantitative analyses), and some seem ambiguous (“Getting Information”)
  2. Weights available (importance/level/frequency) don’t correspond to time shares, as required to map onto the theory
  3. GWAs + LLM-generated time shares: resulting A matrix is low-rank ($\rightarrow$ poor model fit)
  4. Flexibility: our approach is consistent with different occupational classifications (e.g. SOC-2000, which can be x-walked to NLSY) and time periods

Validation of LLM-generated time shares: overview

1. LLM-generated task weights at the occupation-cluster level highly correlated with the average importance rating that O*NET assigns to detailed tasks within each cluster ✓
2. Comparison of time share measurement: LLM vs BIBB survey ✓
3. Comparison of LLM-generated time shares for GWAs to O*NET importance weights ✓
4. Internal consistency: do measurements for detailed occupations aggregate up? ✓

Validation: LLM-generated task shares vs. BIBB

(a) Occupation-level correlations

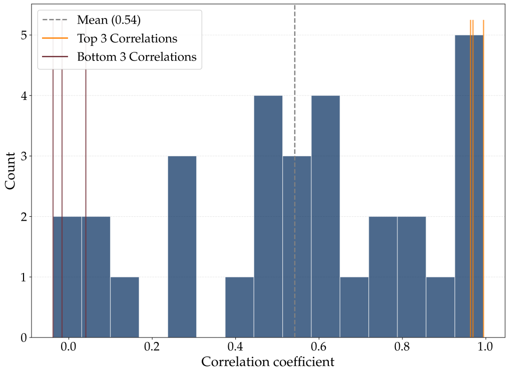

(b) Task-level correlations

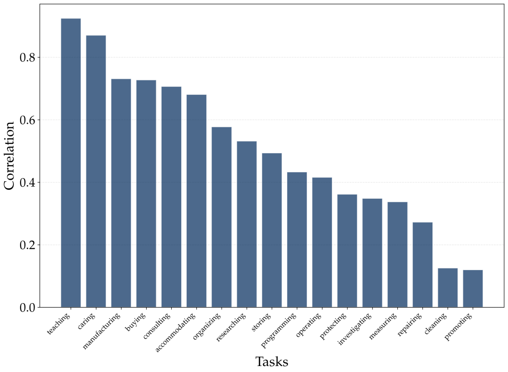

Validation: O*NET GWAs (1)

- Take O*NET GWAs (O*NET 5.0, consistent with SOC-2000), construct relative importance for each GWA by occupation, aggregate to SOC-2000-3d
- Let LLM generate *time shares* for the GWAs for each SOC-2000-3d occ
- How do LLM-time shares correlate with vector of O*NET importance weights?

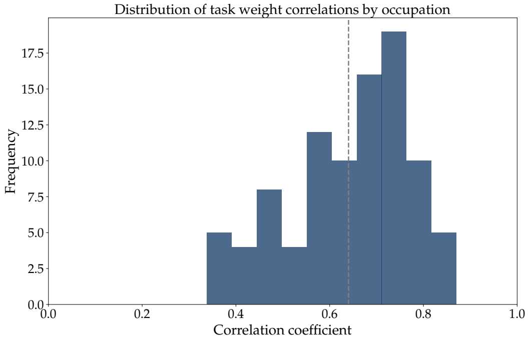

Validation: O*NET GWAs (2): correlation across occupations by task

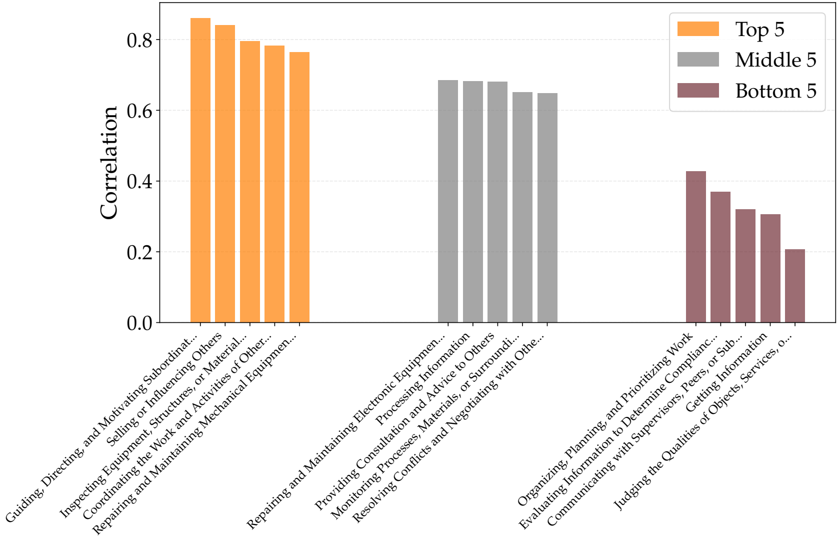

Model fit: occupational wages and employment shares

Back

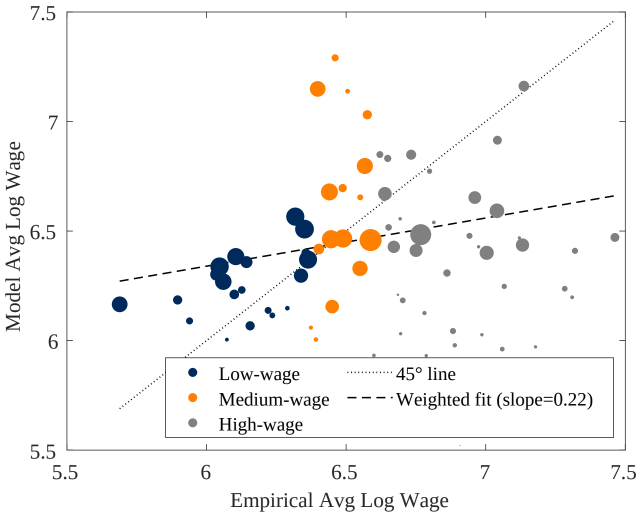

A matrix: example tasks - extracted skills - tasks

Back

Task	Activity	Skills	Cluster
Direct or coordinate an organization's financial or budget activities to fund operations, maximize investments, or increase efficiency	Direct financial operations	Financial management (expert), strategic planning (advanced), budgeting (advanced), analytical thinking (advanced)	Evaluating and Strategizing
Clean and sterilize vats and factory processing areas	Clean and sterilize processing areas	Manual dexterity (basic)	Performing Material Handling Tasks
Press switches and turn knobs to start, adjust, and regulate equipment, such as beaters, extruders, discharge pipes, and salt pumps	Operate equipment controls	Technical knowledge (intermediate), manual dexterity (basic)	Performing Precision Technical Tasks
Conduct research, data analysis, systems design, or support for software such as Geographic Information Systems (GIS) or Global Positioning Systems (GPS) mapping software	Conduct research and data analysis for GIS software	Research skills (advanced), data analysis (advanced), systems design (advanced)	Analyzing Complex Data

Webb measure: selection criteria

Back

Table A1: **Patent selection criteria.**

Technology	Definition
AI	Title/abstract include "neural network", "deep learning", "reinforcement learning", "supervised learning", "unsupervised learning", or "generative model"
Software	Title/abstract include "software", "computer", or "program" AND title/abstract exclude "chip", "semiconductor", "bus", "circuit", or "circuitry"
Robots	Title/abstract include "robot"

Notes: Patents corresponding to each technology are selected using these keyword inclusion/exclusion criteria.

Webb's (2020) exposure measures

▸ Patent criteria ▸ Elondou et al. (2023)

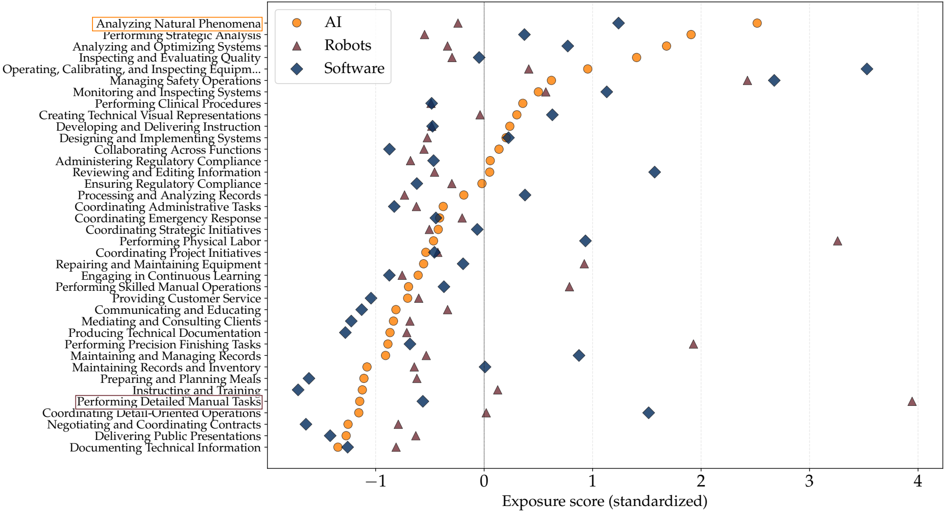

Webb's historical evidence on effects of robots

Back

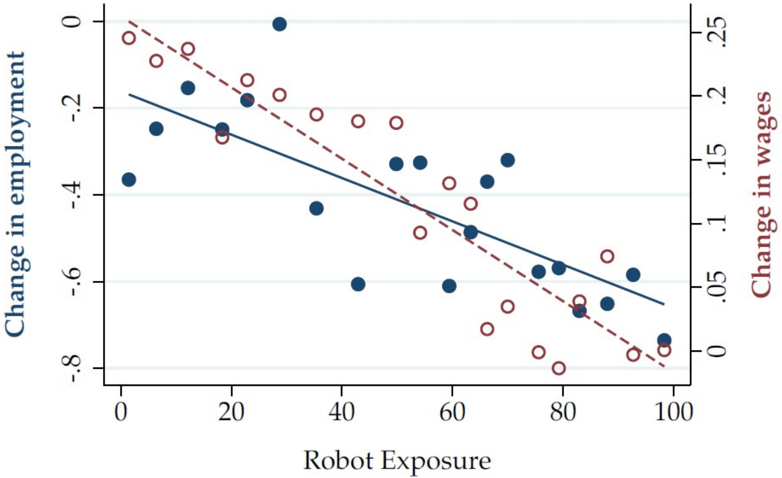

Returns to occupational experience

Main

- **Limitation** of baseline: lower occupational persistence than in data
- **Simple learning amendment:** if a worker picks o in t, if they didn't work in o in $t-1$, their productivity is 1; if they did work in o in $t-1$, their productivity is $\exp(\Delta)$ with $\Delta \geq 0$. Let the expected wages of a worker with skills s_i be

$$w^e_{i,o}(0) = \mu_o + A \cdot s_i$$

$$w^e_{i,o}(1) = \mu_o + \Delta + A \cdot s_i$$

$\Rightarrow$ Worker's (expected) value function satisfies:

$$V_o(0) = w^e_{i,o}(0) + \beta\nu \log \left[\exp\left(\frac{V_o(1)}{\nu}\right) + \sum_{o' \neq o} \exp\left(\frac{V_{o'}(0)}{\nu}\right) \right]$$

$$V_o(1) = w^e_{i,o}(1) + \beta\nu \log \left[\exp\left(\frac{V_o(1)}{\nu}\right) + \sum_{o' \neq o} \exp\left(\frac{V_{o'}(0)}{\nu}\right) \right]$$

and so $V_o(1) = V_o(0) + \Delta$

- **Paper:** higher persistence but similar counterfactual results