

Task-Specific Technical Change and Comparative Advantage*

Lukas Althoff[†]

Hugo Reichardt[‡]

[\[Most recent version here\]](#)

June 10, 2026

Abstract

Artificial intelligence (AI) reshapes workers' comparative advantage by altering the tasks they perform and the skills those tasks require. We develop a dynamic task-based model to quantify the general-equilibrium effects of task-specific technical change. Workers have multidimensional skills, choose occupations, and accumulate skills on the job; occupations combine tasks, and productivity depends on how workers' skills match task requirements. We develop a computationally efficient procedure to estimate the model using panel data and a new database of task-level skill requirements. We apply the model to AI, allowing it to augment, automate, and simplify tasks. We find that AI narrows wage inequality and raises average wages across scenarios ranging from slow to rapid AI progress. The key equalizing force is simplification: by lowering tasks' skill requirements, AI lets lower-skill workers compete for previously inaccessible jobs. Adoption costs, highest for lower-skill workers, dampen but do not eliminate the decline in inequality.

*We thank Ran Abramitzky, Vladimir Asriyan, Marco Bellifemine, Léonard Bocquet, Kyle Herkenhoff, Erik Hurst, Ethan Ilzetzki, Ben Moll, Aureo de Paula, Ricardo Reis, Pascual Restrepo, Todd Schoellman, Jaume Ventura and numerous seminar participants for their insightful comments. Cristian Alamos, Pedro Carvalho, and Bruno Daré Riotta provided excellent research assistance. This work was supported by the Stanford Institute for Economic Policy Research (SIEPR).

[†]Stanford University. lalthoff@stanford.edu

[‡]CREi and Barcelona School of Economics. hreichardt@crei.cat

1 Introduction

Artificial intelligence is transforming how people work—the tasks they perform, the occupations they choose, and the skills they need. Workers’ comparative advantage shifts as AI makes tasks faster to complete, simplifies them, or automates them entirely. How do wages change? Do their skills become more or less valuable? Can workers move to related occupations, and how quickly can they retrain? Answering these questions—jointly and in general equilibrium—requires a quantitative framework that links skill formation, occupational choice, and wage determination.

This paper develops and estimates a dynamic task-based model of the labor market that quantifies the general equilibrium effects of any task-specific technical change—observed or counterfactual. We apply this model to study the labor market effects of artificial intelligence.

We model technical change as the augmentation, automation, and simplification of tasks. Augmentation and automation are standard features of task-based models, capturing increases in human productivity and the substitution of labor by capital (e.g., [Acemoglu and Autor, 2011](#); [Acemoglu and Restrepo, 2018, 2019](#)). Beyond those standard forces, we introduce *simplification*, reflecting that technologies can reduce the skills a task requires, raising the productivity of lower-skill workers more than higher-skill workers.¹ This channel is motivated by emerging experimental evidence that lower-skill workers experience larger productivity gains from AI than higher-skill workers (e.g., [Dell’Acqua et al., 2023](#); [Noy and Zhang, 2023](#); [Brynjolfsson et al., 2025](#); [Cruces et al., 2026](#)).

Our model combines the key features necessary to understand the labor market effects of task-specific technical change. Specifically, workers have multi-dimensional skills and their productivity depends on how well these skills match a task’s requirements. Workers accumulate skills by performing tasks where requirements exceed their current skills. Each period, workers choose an occupation (a bundle of tasks) weighing wages against the future returns to learning on the job.

To quantify the model, we must identify (i) how skills determine workers’ productivity across tasks and (ii) how workers accumulate skills over their careers. The assumption that workers optimally allocate their time across tasks yields a closed-form mapping from task-level productivity to occupation-level productiv-

¹While standard augmentation is usually modeled as skill-neutral, simplification is equivalent to lower-skill biased augmentation.

ity and observed wages. We recover this mapping using detailed data on workers (skills, occupations, and wages) and tasks' skill requirements to recover task-level productivity, identifying (i). We estimate skill accumulation from workers' occupational histories and the evolution of their wages, identifying (ii).

We develop a computationally feasible procedure to estimate the model by maximum likelihood on panel data from the National Longitudinal Survey of Youth 1979 (NLSY79). The computational challenge lies in the model's dynamic and high-dimensional nature. In our quantitative application, workers accumulate five-dimensional skills and differ in their rate of skill accumulation. Two insights make estimation feasible. First, we recover equilibrium prices directly from data, avoiding the need to solve the model inside the estimation loop. Second, we estimate many parameters either by linear regression or by fast iterative routines conditional on the remaining parameters, reducing the dimension of the parameter space we must search non-linearly.

The quantified model offers a laboratory to translate measures of task exposure to technical change into general-equilibrium outcomes. The model predicts how technologies change workers' comparative advantage, how workers reallocate and retrain, and how prices adjust in general equilibrium. We develop computationally efficient algorithms to solve for the new steady state and the transition path. We do so by simulating workers' choices and skill accumulation, iterating on prices until markets clear.

We apply this machinery to AI, predicting its effects on individual workers and the labor market. To construct the model's inputs, we update and expand existing assessments of AI's capabilities to augment, automate, and simplify each task, using large language models (LLMs) to scale established survey designs (following Eloundou et al., 2024; Aghion and Bunel, 2024; Acemoglu, 2025; Asirvatham et al., 2026). Because AI's capabilities are evolving and inherently uncertain, we assess three scenarios—slow, moderate, and rapid AI progress (following Karger et al., 2026). To ensure replicability and guard against idiosyncratic errors, we use several open- and closed-weight LLMs.² We find that these measures align with independent expert assessments, experimental estimates of AI's productivity effects, and changes in skill requirements in job postings.

Our first finding is that AI reduces wage inequality and raises average wages. In the moderate AI scenario, the 90-10 wage ratio falls by 24 percent, and average wages rise by 28 percent. The key force behind the reduction in inequality is sim-

²We make those data accessible via https://github.com/lukasalthoff/ai_labor_markets.

plification. It lowers skill-based barriers by increasing the relative productivity of lower-skill workers in tasks and occupations previously reserved for higher-skill workers. Automation and augmentation, in contrast, have small distributional effects but drive most of the average wage increase.

Second, AI decreases the returns to skills. Verbal skills depreciate most, while manual and math skills retain much of their value. As simplification disincentivizes skill accumulation, average skill levels fall, which dampens AI's positive growth effects. Mapping skills to college majors, we find that engineering, mechanics, and architecture gain value, while law, political science, and philosophy lose value.

Third, we find that AI changes the occupational landscape. For example, administrative occupations (e.g., financial clerks) see a large decline, while construction and manufacturing occupations become more important. The occupations with the largest employment gains are often also those whose relative wages fall the most. This negative relationship between employment and wage effects arises from simplification, which makes jobs easier to perform, enlarges the pool of qualified workers, and suppresses average wages through selection and increased competition (Autor and Thompson, 2025). In contrast, augmentation and automation barely affect relative wages because they leave workers' comparative advantage largely unchanged.

Fourth, despite the cost of adjusting to an unexpected transition, we find that AI generates welfare gains for current workers. In the moderate AI scenario, these gains are equivalent to a rise in lifetime wages between 15 and 45 percent. The gains are largest for the lowest-skill workers who benefit disproportionately from simplification. They vary most among older workers: some have skills and occupations that happen to match the needs of the AI era, while others do not.

Finally, we have so far assumed that all workers adopt AI—yet adoption has been slower among lower-skill workers (Humlum and Vestergaard, 2025; Bick et al., 2026). Does this adoption gap undo AI's equalizing effect? We extend our model to incorporate skill-dependent adoption costs, estimated from data on current AI adoption (Bick and Blandin, 2023; Bick et al., 2026). We find that these costs dampen but do not eliminate AI's equalizing effect: in the moderate scenario, the 90–10 ratio of wages still falls by 15 percent (versus 24% without adoption costs). Average wage gains, however, shrink from 28 to 0.3 percent. Lowering adoption costs is thus a key policy lever for realizing AI's potential to raise wages, especially among lower-skill workers.

Related literature. To understand the labor market effects of task-specific technical change, we integrate three literatures into one empirically tractable framework: task-based production, multi-dimensional skill accumulation, and dynamic occupational choice. First, we provide methods to estimate workers’ comparative advantage across tasks—key for workers’ ability to reallocate in response to technical change (Autor et al., 2003; Acemoglu and Autor, 2011; Autor and Handel, 2013; Acemoglu and Restrepo, 2018, 2022; Hurst et al., 2024). The absence of such methods has made it difficult to quantify the effects of counterfactual (future) technical change (Restrepo, 2024; Woessmann, 2025). Second, while task-based models typically treat workers’ skills as fixed, we incorporate multi-dimensional skill accumulation—also key for workers’ adaptation. Third, we embed task-based production in a general equilibrium model of dynamic occupational choice (Keane and Wolpin, 1997; Heckman et al., 1998; Lee and Wolpin, 2006; Dix-Carneiro, 2014; Traiberman, 2019; Humlum, 2021; Smeets et al., 2025).

We introduce simplification as a channel of task-specific technical change alongside automation and augmentation. Simplification reduces a task’s skill requirements, increasing the relative productivity of lower-skill workers.³ Autor and Thompson (2025) show that historically, the effect of automation on occupations’ wages has depended on whether “expert” or “non-expert” tasks are automated—a form of indirect simplification.⁴ Downey (2021) and Danieli (2025) also consider mechanisms related to simplification.⁵ We quantify all three channels within one general equilibrium framework and find that simplification drives AI’s distributional effects. Our model quantifies a dynamic effect of simplification: because workers build skills by performing demanding tasks, lowering requirements slows skill accumulation.

This paper shows how workers’ multi-dimensional skills shape their comparative advantage across tasks. Recent work emphasizes the multi-dimensionality of skills (Lindenlaub, 2017; Guvenen et al., 2020; Lise and Postel-Vinay, 2020; Bailey et al., 2022; Ide and Talamàs, 2026). Task-based models are inherently multi-dimensional because workers are characterized by their productivity across many tasks. Quantifying these models requires estimating each worker’s productivity, including in tasks they do not currently perform. We overcome this identification problem by providing a microfoundation of task-level productivity as a func-

³Simplification is thus a form of low-skill-biased technical change (Katz and Murphy, 1992).

⁴By indirect simplification we mean the occupation-level decrease in skills (“expertise”) required by automating tasks with relatively high skill requirements.

⁵Simplification is also similar in spirit to “non-autonomous AI,” which allows lower-skill workers to solve hard problems with AI, raising their relative wages (Ide and Talamàs, 2025).

tion of workers' measured skills and tasks' skill requirements. We also provide a new database of these task-level skill requirements, expanding beyond the occupational aggregates available in O*NET (as in [Lise and Postel-Vinay, 2020](#); [Baley et al., 2022](#)).

Finally, by applying our general framework to AI, we add to the rapidly growing evidence on AI's labor market effects. [Freund and Mann \(2025\)](#) introduce a framework to study how automation affects wages through changes in the importance of tasks within occupations. Our approach differs in three ways: we model simplification as a distinct channel through which AI reshapes comparative advantage, we allow workers' skills to evolve over their careers (making the model dynamic), and we estimate task-level comparative advantage directly from workers' measured skills and tasks' skill requirements.⁶ [Hosseini and Lichtinger \(2026\)](#) show that even a simple model—static and with one-dimensional skills—can capture part of the simplification mechanism. [Hampole et al. \(2025\)](#) provide a structural framework to quantify AI's effect on occupational demand through observed adoption across firms. By contrast, we predict how AI affects individual workers, who may respond by switching occupations.

2 Model

We develop a model that describes how workers choose occupations, perform tasks, and accumulate skills over their careers. Overlapping generations of workers live for A periods. Productivity and wages depend on the match between workers' skills and the skill requirements of the tasks relevant to the occupation of choice. Workers accumulate skills on the job, so that their occupational choice depends on the current wage as well as the learning benefits that the job offers. Prices are determined in general equilibrium. Technical change can take the form of augmentation, automation, and simplification of tasks.

2.1 The Firm's and Worker's Problems

Occupations and tasks. Each occupation produces a distinct good by combining a unique set of tasks. These tasks are combined with a constant elasticity of

⁶In contrast, [Freund and Mann \(2025\)](#) infer the distribution of task-level productivity indirectly from data on wages and occupational choices.

substitution ρ , so that the production function of occupation j is

$$Y_j = \left(\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau}^{\frac{1}{\rho}} y_\tau^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \quad (1)$$

where \mathcal{T}_j is the discrete set of relevant tasks, y_τ is the output of task τ , and $\theta_{j,\tau}$ is task τ 's importance weight that satisfies $\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} = 1$.

Task-level productivity and skills. The production function for task τ in occupation j depends on whether the task is automatable, i.e., $\tau \in \mathcal{A}_j$, or not, i.e., $\tau \in \mathcal{N}_j$:

$$y_\tau(\mathbf{h}, \ell_\tau, k_\tau) = \begin{cases} \ell_\tau \gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau) & \text{if } \tau \in \mathcal{N}_j \\ \ell_\tau \gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau) + \phi_\tau k_\tau & \text{if } \tau \in \mathcal{A}_j \end{cases} \quad (2)$$

where ℓ_τ represents the share of the worker's time allocated to task τ , γ_τ and ϕ_τ capture the task-specific productivity of humans and capital respectively, $\mathbf{h} = (h_s)_{s \in \mathcal{S}}$ is a vector capturing workers' multi-dimensional skills, $\mathbf{r}_\tau \equiv (r_{\tau,s})_{s \in \mathcal{S}}$ is the skill requirement of task τ , and k_τ is capital devoted to task τ . If the task is automatable, capital and labor are perfect substitutes.

The function $f(\cdot)$ determines how workers with different skills \mathbf{h} are differentially productive in tasks depending on their skill requirements, \mathbf{r}_τ . Intuitively, this function captures how workers' productivity depends on the "match" between their skills and the skills required to complete the task (Lise and Postel-Vinay, 2020; Baley et al., 2022). For our quantification, we will assume a functional form for $f(\cdot)$ (in section 2.3, equation 14) and show how its parameters can be identified and estimated.

Technical change. We consider three different ways in which technical change affects the task-level production function:

- Augmentation* Enhancing human productivity, increasing γ_τ ;
- Automation* Substituting labor with capital, expanding \mathcal{A}_j ;
- Simplification* Simplifying tasks for humans, reducing \mathbf{r}_τ .

The first two are standard in the task-based literature (e.g., Acemoglu and Autor, 2011). We introduce simplification to allow for heterogeneity in technologies' impact on productivity depending on workers' skills—essentially skill-biased augmentation at the task-level (Katz and Murphy, 1992). Experimental evidence has

shown that AI's productivity effects tend to be stronger for lower-skill workers, suggesting that simplification is an important force in practice (e.g., [Dell'Acqua et al., 2023](#); [Noy and Zhang, 2023](#); [Brynjolfsson et al., 2025](#); [Cruces et al., 2026](#)). While we call this type of technical change "simplification," the methodology equally allows for increases in skill requirements (or "complication").

The firm's problem. Each good j is produced by a representative firm that takes the equilibrium wage $\{w_j(\mathbf{h})\}_h$ and the costs of capital R as given. The firm chooses how many workers of each skill set \mathbf{h} to hire, how to allocate their time across tasks, and how much capital to use for each task. Formally, the firm solves the following profit maximization problem:

$$\begin{aligned} \max_{\{n_j(\mathbf{h}), \ell_{j,\tau}(\mathbf{h}), k_{j,\tau}(\mathbf{h})\}} & \int n_j(\mathbf{h}) \left(p_j Y_j(\mathbf{h}) - w_j(\mathbf{h}) - R \sum_{\tau \in \mathcal{A}_j} k_{j,\tau}(\mathbf{h}) \right) d\mathbf{h} \\ \text{s.t. } Y_j(\mathbf{h}) &= \left(\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau}^{\frac{1}{\rho}} \tilde{y}_\tau(\mathbf{h})^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \\ & \sum_{\tau \in \mathcal{T}_j} \ell_{j,\tau}(\mathbf{h}) = 1 \forall \mathbf{h} \\ & \tilde{y}_\tau(\mathbf{h}) \equiv y_\tau(\mathbf{h}, \ell_{j,\tau}(\mathbf{h}), k_{j,\tau}(\mathbf{h})) \text{ given by (2),} \end{aligned} \quad (3)$$

where $n_j(\mathbf{h})$ is the amount of labor employed with skill \mathbf{h} , $\ell_{j,\tau}(\mathbf{h})$ is the share of time allocated to task τ , and $k_{j,\tau}(\mathbf{h})$ denotes the capital per worker allocated to workers with skill \mathbf{h} working on task τ . Because the marginal product of capital depends on a worker's skills, the firm generally does not allocate an equal amount of capital to each worker.

Since the labor market is perfectly competitive the worker's wage must equal their marginal product. That is,

$$w_j(\mathbf{h}) = p_j Y_j(\mathbf{h}) - R \sum_{\tau \in \mathcal{A}_j} k_{j,\tau}(\mathbf{h}). \quad (4)$$

A worker's wage thus depends on the allocation of their time and capital to the various tasks in the occupation. If the worker only cares about the allocation insofar as it affects the wage, the worker and firm would always agree that the allocation should maximize the worker's value added. However, when the allocation also affects future payoffs (such as through skill accumulation as will be the case in this model) the worker would in principle be willing to accept a lower wage

in exchange for an allocation that yields higher future payoffs (through increased learning). To rule out such an exchange, we make the following assumption:

Assumption 1 (Control and non-contractibility of task assignment). The firm controls the workers' time allocation across tasks which is not contractible.

Assumption 1 implies that the firm can freely determine how workers allocate their time across tasks. Because time allocation is non-contractible, any promised task assignment cannot be enforced by workers. As a result, the firm retains full flexibility to adjust workers' task allocations as it sees fit. This assumption implies that the firm chooses the allocation that maximizes profits (output net of capital costs) for any given wage $w_j(\mathbf{h})$. In equilibrium, wages must thus equal the worker's marginal product given this allocation.

A second assumption is that automatable tasks are cheaper to perform with capital than with labor, ensuring the worker's time is only allocated to non-automatable tasks:

Assumption 2 (Full automation of automatable tasks). The unit cost of producing a task with capital is lower than the cost of producing it with labor for all occupations j , skills \mathbf{h} , and tasks $\tau \in \mathcal{A}_j$,

$$\frac{w_j(\mathbf{h})}{\gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau)} > \frac{R}{\phi_\tau}.$$

Optimal time allocation. Assumptions 1 and 2 together imply that a worker's time is allocated to non-automatable tasks to maximize the production of these tasks. That is, the firm solves the following time allocation problem:

$$\{\ell_{j,\tau}(\mathbf{h})\}_{\tau \in \mathcal{N}_j} = \arg \max_{\{\ell_\tau\}_{\tau \in \mathcal{N}_j}} \left(\sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau}^{\frac{1}{\rho}} (\ell_\tau \gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau))^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \quad \text{s.t.} \quad \sum_{\tau \in \mathcal{N}_j} \ell_\tau = 1.$$

The solution to this problem for a given task $\tau \in \mathcal{N}_j$ is

$$\ell_{j,\tau}(\mathbf{h}) = \frac{\theta_{j,\tau} \gamma_\tau^{\rho-1} f(\mathbf{h}, \mathbf{r}_\tau)^{\rho-1}}{\sum_{k \in \mathcal{N}_j} \theta_{j,k} \gamma_k^{\rho-1} f(\mathbf{h}, \mathbf{r}_k)^{\rho-1}}, \quad (5)$$

which shows that more time is spent on tasks with greater weight $\theta_{j,\tau}$. If tasks are substitutes ($\rho > 1$), a worker's time is allocated to the tasks they are most productive in, i.e., tasks for which $\gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau)$ is greater. If instead tasks are complements ($\rho < 1$), workers spend more time on less productive tasks.

Optimal capital allocation. In choosing how much capital to allocate to each worker-task pair, the firm balances the marginal benefit of increased output against the cost of capital. The first order condition implies that for all tasks $\tau \in \mathcal{A}_j$

$$k_{j,\tau}(\mathbf{h}) = \theta_{j,\tau} \phi_\tau^{\rho-1} Y_j(\mathbf{h}) \left(\frac{p_j}{R} \right)^\rho, \quad (6)$$

where $Y_j(\mathbf{h})$ is the profit-maximizing level of output when a worker with skill vector \mathbf{h} works in occupation j . Clearly, the lower the cost of capital relative to the price of the output, the more capital the firm uses. Also, firms allocate more task-automating capital to workers that are more productive in the non-automated tasks, i.e., workers for which $Y_j(\mathbf{h})$ is larger.

Occupational productivity. Given the optimal allocation of workers' time and capital to tasks, the output per worker of type \mathbf{h} in occupation j is

$$Y_j(\mathbf{h}) = \Gamma_j^{\frac{\rho}{1-\rho}} \left(\sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} f(\mathbf{h}, \mathbf{r}_\tau)^{\rho-1} \right)^{\frac{1}{\rho-1}} \quad (7)$$

where

$$\Gamma_j = 1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \left(\frac{R/\phi_\tau}{p_j} \right)^{1-\rho}$$

is the labor share. Equation (7) shows that a worker's output in occupation j is a weighted function of their productivity across all non-automated tasks $\tau \in \mathcal{N}_j$.

Wages. Combining equations (4), (6), and (7) yields the wage when a worker of skill \mathbf{h} chooses occupation j :

$$w_j(\mathbf{h}) = p_j Y_j(\mathbf{h}) \Gamma_j = p_j \Gamma_j^{\frac{1}{1-\rho}} \left(\sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} f(\mathbf{h}, \mathbf{r}_\tau)^{\rho-1} \right)^{\frac{1}{\rho-1}}. \quad (8)$$

Equation (8) shows that if none of the tasks are automatable, i.e. $\mathcal{A}_j = \emptyset$ and $\Gamma_j = 1$, a worker's income equals total revenue $w_j(\mathbf{h}) = p_j Y_j(\mathbf{h})$.

Skill accumulation. Before entering the labor market at age $a = 1$, each worker draws an initial skill vector \mathbf{h}_1 after which they accumulate further skills on the job. We assume that a worker's human capital accumulation depends on their

current skills, their “ability to learn” ψ , and the skill requirements of the tasks in their job j : $\mathbf{h}' = g_j(\mathbf{h}, \psi)$.

Occupational choice. Every period, each worker chooses from a discrete set of occupations to maximize utility. Workers consume their income each period and live for A periods. The expected lifetime utility of a worker aged a , with learning ability ψ , and previous occupation k is represented by the value function

$$V_a(\mathbf{h}, \psi, k) = \mathbb{E} \left[\max_j \log w_j(\mathbf{h}) + \log \varepsilon_j + \mu_j - \kappa(k, j) + \beta V_{a+1}(g_j(\mathbf{h}, \psi), \psi, j) \right] \quad (9)$$

where $\mathbb{E}[\cdot]$ is the expectation over the occupation-specific productivity shocks ε_j . These productivity shocks are relevant to wages: a worker’s observed wage equals the deterministic part ($w_j(\mathbf{h})$) times the stochastic part (ε_j).

Each occupation j has an amenity value μ_j . $\kappa(k, j)$ is a cost of switching from occupation k to j . In our quantitative application, we set this to $\kappa(k, j) = \kappa \mathbb{1}[j \neq k]$ for some constant κ . $g_j(\mathbf{h}, \cdot)$ is next period’s human capital when choosing occupation j . Because the worker’s life is finite, the value after terminal age A is zero, $V_{A+1}(\cdot) = 0$.

We assume that the productivity shocks $\log \varepsilon_j$ follow a type I generalized extreme value (Gumbel) distribution with mean 0 and scale parameter ζ .⁷ This assumption implies that the conditional probability of choosing occupation j has the closed-form solution

$$\mathbb{P}_a(j \mid \mathbf{h}, \psi, k) = \frac{\exp\left(\frac{1}{\zeta} (\log w_j(\mathbf{h}) + \mu_j - \kappa(k, j) + \beta V_{a+1}(g_j(\mathbf{h}, \psi), \psi, j))\right)}{\sum_{l=1}^J \exp\left(\frac{1}{\zeta} (\log w_l(\mathbf{h}) + \mu_l - \kappa(k, l) + \beta V_{a+1}(g_l(\mathbf{h}, \psi), \psi, l))\right)} \quad (10)$$

so that the value function in (9) can be simplified to

$$V_a(\mathbf{h}, \psi, k) = \zeta \log \sum_{j=1}^J \exp\left(\frac{1}{\zeta} (\log w_j(\mathbf{h}) + \mu_j - \kappa(k, j) + \beta V_{a+1}(g_j(\mathbf{h}, \psi), \psi, j))\right) \quad (11)$$

Since $V_{A+1}(\cdot) = 0$, equation (11) solves the value function, and thus the occupational choice problem, by backward iteration from age A to 1 for a given sequence of prices. Note that the notation above omits any dependence on time. In principle, prices vary over time, so that the wage schedule $w_{j,t}(\mathbf{h})$, and thus the value

⁷The CDF is $\Pr(\log \varepsilon < x) = \exp\left(-\exp\left(-\frac{x+\zeta\bar{\gamma}}{\zeta}\right)\right)$ where $\bar{\gamma} \approx 0.577$ is Euler’s constant. This implies that ε_j follows a Fréchet distribution.

functions, are time-dependent.

2.2 Equilibrium

The price of each occupational good p_j is determined in equilibrium through demand and supply. The supply is characterized by the solution to the worker's problem. The workers, in turn, consume and generate demand for the occupational goods. We assume that demand is characterized by a homothetic and invertible demand function $D(\{p_j\}_{j=1}^J)$ that maps prices p_j into relative demand for each occupational good. In our application, we use (standard or nested) CES demand. Having specified demand, we can now define the competitive equilibrium.

Definition (Competitive equilibrium). Given an initial joint distribution of age, skills, ability, and occupations, $G_{a,t}(\mathbf{h}, \psi, k)$, a distribution of human capital at birth $G_{1,t}(\mathbf{h}, \psi)$, and the supply of capital $\{\mathcal{K}_t\}_{t=1}^\infty$, a *competitive equilibrium* is defined as a sequence of prices $\{p_{1,t}, \dots, p_{J,t}, R_t\}_{t=1}^\infty$ such that

- Workers' occupational choices maximize the present value of lifetime utility given the sequence of prices. That is, their occupational choice probabilities are as in equation (10);
- The distribution over states follows from occupational choices:

$$G_{a+1,t+1}(\mathbf{h}', \psi, j) = \sum_{k=1}^J \int_{g_j(\mathbf{h}, \psi) \leq \mathbf{h}'} \mathbb{P}_{a,t}(j | \mathbf{h}, \psi, k) dG_{a,t}(\mathbf{h}, \psi, k); \quad (12)$$

- Demand for goods equals supply: $D(\{p_{j,t}\}_{j=1}^J) \propto \mathcal{Y}_{j,t}$, where

$$\mathcal{Y}_{j,t} \equiv \sum_{a=1}^A \sum_{k=1}^J \int Y_j(\mathbf{h}) \mathbb{P}_{a,t}(j | \mathbf{h}, \psi, k) \mathbb{E}[\varepsilon_j | j, \mathbf{h}, \psi, k] dG_{a,t}(\mathbf{h}, \psi, k);^8 \quad (13)$$

- Demand for capital equals supply:

$$\mathcal{K}_t = \sum_{a=1}^A \sum_{k=1}^J \sum_{j=1}^J \sum_{\tau \in \mathcal{A}_j} \int k_{j,\tau}(\mathbf{h}) \mathbb{P}_{a,t}(j | \mathbf{h}, \psi, k) dG_{a,t}(\mathbf{h}, \psi, k).$$

⁸ $\mathbb{E}[\varepsilon_j | j, \mathbf{h}, \psi, k]$ is the expectation of the productivity shock conditional on choosing occupation j when your states were \mathbf{h}, ψ, k . The Gumbel distribution of $\log \varepsilon_j$ implies that this expectation has a closed-form solution: $\mathbb{E}[\varepsilon_j | j, \mathbf{h}, \psi, k] = \exp(-\zeta\gamma)\Gamma(1 - \zeta)\mathbb{P}_{a,t}(j | \mathbf{h}, \psi, k)^{-\zeta}$.

Solution algorithms. We provide algorithms to solve for a stationary competitive equilibrium and the transition path after an unexpected technological shock.

These solution algorithms iterate over vectors of equilibrium prices for each occupational good at each point in time. For a given guess for such prices, we solve the worker’s problem, compute the implied supply of each occupational good, and update the price accordingly. The algorithm exploits a convenient feature of equation (11) to avoid solving the value function for each occupation separately despite occupations being a state variable. Appendix C describes the algorithms formally and in more detail.

2.3 Parametrization

In our quantitative application, we make assumptions on the functions that govern task-level productivity and skill accumulation.

Production. We specify the task-level production function as

$$f(\mathbf{h}, \mathbf{r}_\tau) = \prod_{s \in \mathcal{S}} h_s^{\omega_s} \exp\left(-\eta \min\{h_s - r_{\tau,s}, 0\}^2\right). \quad (14)$$

This production function is similar to that proposed by [Lise and Postel-Vinay \(2020\)](#). The first term in equation (14) reflects a force that makes workers with higher skills more productive in *any* task, independent of its skill requirements. The second (exponential) term captures the degree to which the worker’s productivity is diminished when performing tasks for which they are “underqualified.” Figure A.1a shows the functional form graphically.

This functional form assumption implies that the wage function in (4) equals

$$w_j(\mathbf{h}) = p_j \Gamma_j^{\frac{1}{1-\rho}} \prod_{s \in \mathcal{S}} h_s^{\omega_s} \left(\sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} \exp\left(-\eta \sum_{s \in \mathcal{S}} \min\{h_s - r_{\tau,s}, 0\}^2\right)^{\rho-1} \right)^{\frac{1}{\rho-1}}. \quad (15)$$

Skill accumulation. We assume that skill accumulates according to the following functional form:

$$g_{j,s}(\mathbf{h}, \psi) = (1 - \delta)h_s + \sum_{\tau \in \mathcal{N}_j} \ell_{j,\tau}(\mathbf{h}) \max\{r_{\tau,s} - h_s, 0\} e^{-\lambda(\psi) \max\{r_{\tau,s} - h_s, 0\}}, \quad (16)$$

where $\ell_{j,\tau}(\mathbf{h})$ indicates workers’ time spent on task τ , defined in equation (5). Equation (16) has several intuitive implications for skill accumulation. First, workers’ learning is most affected by the tasks they spent most time on, i.e., for which $\ell_{j,\tau}(\mathbf{h})$ is greatest.⁹ Second, workers learn by performing tasks that have skill requirements above their current skill levels. However, workers learn most from tasks that are not “too hard.” As tasks become harder relative to the workers’ skill, the rate at which skills catch up decreases; $\lambda(\psi) \geq 0$ governs the rate of this slowdown. Figure A.1b illustrates how learning varies with the distance between the worker’s skills and the task’s requirements.¹⁰ Similar to Heckman et al. (1998), we allow for workers to differ in their ability to learn ψ . Lastly, some skill depreciation occurs independently of which tasks are performed, governed by δ .

3 Data

We use three data sources to estimate the model’s parameters. First, we rely on O*NET to measure each occupation’s tasks and skill requirements. Second, we provide and validate a new database that extends O*NET’s occupation-level data on skill requirements to the task level using LLMs. Third, we use panel data on wages, occupational choices, and multi-dimensional skills from the NLSY79.

For the application to AI, we also require data on AI’s capabilities to augment, automate, and simplify each task. We follow the literature in using LLMs to estimate AI’s task-level capabilities (e.g., Eloundou et al., 2024; Acemoglu, 2025). Because these capabilities are uncertain and evolving, we do not rely on a single measure but instead construct three capability scenarios—slow, moderate, and rapid (Karger et al., 2026)—and elicit each from several open- and closed-weight LLMs. We compare the resulting measures with human expert ratings and experimental evidence, with additional validation from job postings data.

3.1 Estimation Data

3.1.1 Occupations and Tasks (O*NET)

O*NET is the leading database on occupations, tasks, and skills in the US economy (e.g., Autor et al., 2003; Acemoglu and Autor, 2011; Lise and Postel-Vinay, 2020).

⁹The workers’ time allocation across tasks thus affects skill accumulation. Assumption 1 implies that firms do not take this into account.

¹⁰Since $f(x) = x \exp(-\lambda x)$ is strictly increasing for $x < \frac{1}{\lambda}$ and strictly decreasing after, learning from task τ is maximized when $r_{\tau,s} - h_s = 1/\lambda$, yielding a learning gain of $1/(e\lambda)$.

O*NET contains detailed descriptions of 19,530 tasks linked to 974 occupations. We rely on these data to define both the occupations j and tasks \mathcal{T}_j in our model.¹¹ We set the weights of each task τ in an occupation, $\theta_{j,\tau}$, to the importance measure of that task as reported in O*NET.

We also use O*NET’s definition of worker skills across 35 dimensions (e.g., “reading comprehension” or “social perceptiveness”). O*NET rates skills on a scale from 1 to 7 and provides anchors for each level (e.g., level 2 in reading comprehension corresponding to “read[ing] step-by-step instructions for completing a form” and 4 to “understand[ing] an email from management describing new personnel policies”).¹²

We reduce O*NET’s dimensions to five skill categories: manual, math, social, technical, and verbal.¹³ Table B.1 shows this mapping.

3.1.2 Task-level Skill Requirements

Workers’ comparative advantage in a task is governed by the match between their skills and the task’s skill requirements, r_τ . While O*NET provides data on occupation-level skill requirements, it lacks task-specific data. To address this gap, we use several LLMs to estimate task-level skill requirements. To ensure consistency with O*NET and a valid survey design, we replicate O*NET’s occupation-level questionnaire at the task level, by using its questionnaire format, skill dimensions, and skill anchors. This process covers 19,530 task descriptions across 35 skills using 683,550 queries per model. See Appendix E.1 for further details on our prompt design.

We validate our data by comparing aggregations of our newly generated task-level data with O*NET’s occupation-level measures. For each occupation, we calculate importance-weighted average task-level skill requirements ($\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} r_{\tau,s}$) and compare these with corresponding O*NET values. The five aggregated skills have high agreement rates, with correlations ranging from 0.79 to 0.93 (see Figure A.2).¹⁴ The high agreement with O*NET suggests that LLMs can accurately identify differences in skill requirements across occupations.

¹¹Instead of the most granular definition of occupations, our analysis relies on 93 aggregate (3-digit) occupation groups.

¹²The original O*NET questionnaire and skill level descriptions are available [here](#).

¹³We equally weight each component in the aggregation of skills to our five final dimensions. Relative to [Addison et al. \(2020\)](#); [Baley et al. \(2022\)](#); [DeLoach et al. \(2022\)](#), we include manual as a separate skill because its interaction with AI is of particular interest.

¹⁴Agreement rates are also high across most of the 35 original O*NET skill dimensions (see Figure A.3).

We also collect new human expert ratings of skill requirements for 200 skill-task combinations in the occupation *Economist*. Skill requirements are highly correlated between LLMs and the average of two human experts ($\rho = 0.70$; see Figure A.4a), with correlations between LLMs and individual experts ranging from 0.59 to 0.65. This correlation provides evidence that LLMs can also accurately distinguish skill requirements of different tasks within an occupation.

3.1.3 Skills, Occupational Choice, and Wages (NLSY79)

We use data from the NLSY79 to estimate the task-level production function and the skill accumulation function. The data contain information on wages, occupations, and multi-dimensional skill assessment scores.

We follow the literature in measuring skills in the NLSY79. Following [Addison et al. \(2020\)](#); [Baley et al. \(2022\)](#), we measure skills with the Armed Services Vocational Aptitude Battery (ASVAB): (i) manual skills via the average of standardized scores on *auto and shop information* and *mechanical comprehension*; (ii) math skills via *mathematics knowledge* and *arithmetic reasoning* scores; (iii) technical skills via *general science* and *electronics information*; (iv) verbal skills via *paragraph comprehension* and *word knowledge*. For social skills, we use a composite measure of self-reported sociability as a young adult, sociability at age 6, the *Rotter Locus of Control Scale*, and the *Rosenberg Self-Esteem Scale* (see also [Deming, 2017](#); [Addison et al., 2020](#); [Guvenen et al., 2020](#)).

These data only provide an ordinal measure of skills. We observe $\tilde{h} = F(h)$ where $F(\cdot)$ is the initial skill distribution. We do not directly observe the cardinal measure h that is on the same scale as the skill requirements and estimate the marginal distribution of skills together with all other parameters (see section 4).

We follow the NLSY79 cohort’s labor market history from age 25 to the survey in 2022. We retain information on all jobs held, including their start and end dates, the occupational code, the hourly wage, and the number of hours worked per week. Similar to [Lise and Postel-Vinay \(2020\)](#), we only consider workers for whom the maximum gap between observed jobs is no larger than 18 months. We collapse this data to a worker panel of yearly frequency.

3.2 Data on AI’s Capabilities

In the model of Section 2, technical change takes three forms: augmentation, automation, and simplification. We construct scenarios for AI’s task-specific capabil-

ities in each channel using LLMs. We define slow, moderate, and rapid progress scenarios following [Karger et al. \(2026\)](#) and elicit each from several open- and closed-weight LLMs (Alibaba’s Qwen 2.5, OpenAI’s gpt-oss, and OpenAI’s GPT-4o). The open-weight elicitations, run at zero temperature, are exactly replicable. We report results for all three scenarios, using the moderate scenario as our reference (summarized in Table B.2).

Augmentation. In measuring AI’s potential to augment human productivity, we follow [Eloundou et al. \(2024\)](#) who asked human raters and OpenAI’s GPT-4 whether they believed that LLMs can reduce the worker’s time required to complete a task by at least half. We replicate their exercise with several more recent LLMs except that we ask for a continuous estimate of the percentage of time saved (rather than a binary measure), and consider generative AI more broadly (rather than LLMs). On average, we estimate that generative AI in the moderate scenario saves 25.2 percent of worker time (see Table B.2). Appendix E.1.1 describes our prompt design.

The LLMs’ predictions align with experimental estimates of average productivity effects of generative AI in various tasks and occupations (see Table B.3). We also find that our estimates are strongly correlated and comparable with both the human- and LLM-rated data from [Eloundou et al. \(2024\)](#)—see Figure A.4c.

Average worker productivity in a task can be affected through both augmentation (making all workers equally more productive) and simplification (making lower-skill workers differentially more productive). To recover augmentation from aggregate productivity effects, we first compute simplification’s average productivity effects across all workers.¹⁵ We then compute augmentation’s productivity effect as the difference between the aggregate and the simplification-specific productivity effects of AI.

Automation. We follow [Eloundou et al. \(2024\)](#) by eliciting automatability by task from LLMs. For each of the tasks in O*NET, we ask whether AI can complete it autonomously (i.e., whether task τ is in the automatable set \mathcal{A}_j). [Eloundou et al. \(2024\)](#) classify tasks as having either “no”, “low”, “moderate”, “high” or “full” exposure to automation.¹⁶ We classify a task as “automatable” if it has high or full automation exposure, which is restricted to cases where LLMs indicate that generative AI can complete at least 90 percent of the components of the tasks. In

¹⁵We use the distribution of skills by occupation implied by the model’s pre-AI equilibrium.

¹⁶The specific prompt is documented in ([Eloundou et al., 2024](#), Supplementary Materials).

the moderate AI scenario, 22.4 percent of all tasks are classified as automatable by generative AI (see Table B.2). The prompt is documented in Appendix E.1.2.

We find high agreement between our measures and those obtained by [Eloundou et al. \(2024\)](#). The share of tasks that are automatable is almost identical across the measures. Importantly for our exercise, we find strong agreement across tasks (see Figure A.4d).


Simplification. Lastly, we elicit the degree to which AI changes tasks' skill requirements, r_T . In addition to our new data on pre-AI task-level skill requirements, we prompt several LLMs to evaluate the task's skill requirements before and after workers gain access to generative AI (similar to [Hosseini and Lichtinger, 2026](#)). The prompt can be found in Appendix E.1.3. We estimate that across all tasks and skill dimensions, a one-step reduction (out of seven) is the most common change in the moderate AI scenario (see Table B.2).

According to these predictions, there is large heterogeneity in what AI can simplify across skills: the strongest simplification occurs in time management, writing, judgment and decision making, and critical thinking, versus the least simplification in the manual skills of equipment maintenance, repairing, and installation.

We find that occupations with larger predicted simplification experience a decline in skill requirements in real-time job postings. We use a random subsample of US job postings from Revelio Labs' Cosmos database, which covers the full text of over 5 billion job postings from company websites and job boards since 2007. Following [Autor and Thompson \(2025\)](#), we compute expertise indices by occupation over time based on the relative frequency of words specific to a professional domain. Figure 1 shows that occupations predicted to be more simplified by AI exhibit no differential trend in required expertise prior to 2022, but a sharp and accelerating decline thereafter—reaching 0.4 standard deviations by early 2026. This pattern is corroborated by two additional measures of occupational skill requirements in Figure A.5a: both the frequency of skill-related language and the rate of technical keywords in job postings fall post-2022 in occupations our model predicts to be more simplified.

Lastly, that AI tends to simplify tasks aligns with experimental evidence that has consistently shown larger AI-led productivity effects among lower-skill workers across several occupations (e.g., [Dell'Acqua et al., 2023](#); [Noy and Zhang, 2023](#); [Brynjolfsson et al., 2025](#); [Cruces et al., 2026](#)).

FIGURE 1: SIMPLIFICATION VIA JOB POSTINGS



Notes: This figure shows that occupations predicted to be more simplified by AI exhibit a decline in required expertise in real-time job postings. Continuous treatment exposure is measured using our occupation-level predicted AI-led simplification (slow AI scenario to avoid factoring in post-2026 changes in AI capabilities). Following [Autor and Thompson \(2025\)](#), the outcome is an expertise index based on the relative frequency of words that are highly specific to a professional domain. We use a random subsample of US job postings in Revelio Labs’ Cosmos database. Each coefficient gives the differential effect of a 1 standard deviation higher predicted simplification in t relative to 2022-H1. We include occupation and time fixed effects; standard errors are clustered at the occupation level.

4 Estimation

We jointly estimate the parameters governing productivity, skill accumulation, occupational choices, and the initial skill distribution using maximum likelihood on the NLSY79 estimation sample. We provide a computationally efficient methodology to do so using direct inference. The algorithm recovers the equilibrium prices directly, avoiding the need to solve for the equilibrium within the estimation loop. Table B.4 shows an overview of all model parameters and their estimated values.

4.1 Estimation strategy

The goal of the estimation strategy is to find the parameters that maximize the likelihood of the observed wages and occupational choices. A full maximum-likelihood approach over the complete parameter space is computationally prohibitive given the dynamic and high-dimensional nature of the model. We reduce the computational burden in two main ways.

First, we maximize the likelihood using a sequential approach. We partition the large parameter space into an outer vector θ_1 and a set of objects that can be consistently estimated using a fast procedure for any given θ_1 : the vector of task-level productivity parameters $\theta_y = \{\eta, \{\omega_s\}_{s \in S}\}$, occupational amenities $\mu = \{\mu_j\}_{j=1}^J$, and equilibrium prices $p = \{p_j\}_{j=1}^J$. We then maximize the likelihood

$$L(\theta_1, \theta_y(\theta_1), p(\theta_1), \mu(\theta_1))$$

with respect to θ_1 only, obtaining consistent estimates of $\theta_y(\theta_1)$, $p(\theta_1)$, and $\mu(\theta_1)$ via closed-form expressions and a fast iterative algorithm. This reduces the dimensions of the nonlinear search by 95 percent.

Second, we only maximize the likelihood of the occupational choices for workers at the terminal age A . Since the occupational choice problem for these workers is static, this avoids repeated solution of the dynamic value function.

Inner algorithm. The inner algorithm proceeds in three steps: infer workers' skills, estimate the task-level productivity parameters using linear regression, and recover occupational amenities from occupational choices.

The first step in the inner algorithm is to compute the workers' skills given θ_1 . The NLSY79 provides multi-dimensional skill scores. However, we observe those skills i) only as percentile scores, not as cardinal measures, and ii) only at labor market entry, not later. We first map percentile scores into cardinal skills h using the marginal distribution, approximated with a Beta distribution with parameters included in the outer vector θ_1 .¹⁷ That is, we recover the cardinal skills as $h = F^{-1}(\tilde{h})$, where $F(\cdot)$ is the estimated Beta distribution of initial skills and \tilde{h} are the workers' percentile skill scores. We then successively apply the skill accumulation function $g_j(\cdot, \psi)$ in equation (16) to infer workers' skills at later ages. That is, given worker i 's occupational history $j_i^{a-1} \equiv \{j_{i,1}, \dots, j_{i,a-1}\}$, worker i 's skill level at age a is

$$h_{i,a}(\lambda(\psi_i), \delta) \equiv \left(g_{j_{i,a-1}}(\cdot, \psi_i) \circ g_{j_{i,a-2}}(\cdot, \psi_i) \circ \dots \circ g_{j_{i,1}}(\cdot, \psi_i) \right) (h_{i,1}).$$

where $(f \circ g)(x) \equiv f(g(x))$.¹⁸ Following Heckman et al. (1998), we proxy ψ_i by the Armed Forces Qualification Test (AFQT) score. The parameters $\{\lambda(\psi)\}_{\psi=1}^4$ and δ are in θ_1 .

We then estimate the task-level productivity parameters and equilibrium prices

¹⁷The Beta distribution is a flexible distribution characterized by two parameters B_a and B_b with support on $[0,1]$. We assume that this distribution is common across skill dimensions.

¹⁸To save notation, it is left implicit above that $h_{i,a}$ depends on $\lambda(\psi_i)$ and δ through $g_j(\cdot, \psi_i)$.

using a linear regression. The derived occupational wage function in equation (15) governs how skills translate into earnings in each occupation depending on the occupation's price and tasks, and the parameters of the production function $\{\omega_s\}_{s \in S}$ and η . A log-linearization of this function around zero mismatch yields

$$\log w_j(\mathbf{h}_i) \approx \log p_j + \sum_{s \in S} \omega_s \log(h_{is}) - \eta \sum_{s \in S} \sum_{\tau \in \mathcal{T}_j} \tilde{\theta}_{j,\tau} \min\{h_{is} - r_{\tau,s}, 0\}^2 \quad (17)$$

where $\tilde{\theta}_{j,\tau} \equiv \theta_{j,\tau} \gamma_\tau^{\rho-1}$ (see Appendix D.2 for the proof).¹⁹ Equation (17) implies that we can estimate the equilibrium prices p_j and the parameters of the production function using a simple OLS regression of wages on occupational fixed effects, skills, and skill mismatch. Since the equilibrium prices can be recovered from the occupational fixed effects in the regression, no equilibrium solution is required within the estimation loop.

We correct for selection on the occupation-specific productivity shocks using the control-function approach of [Dubin and McFadden \(1984\)](#). When observed wages equal the deterministic component $w_j(\mathbf{h})$ times the stochastic shock ε_j , selection arises because workers observe ε_j when choosing their occupation. The Gumbel distribution of $\log \varepsilon_j$ implies that its expectation conditional on choosing occupation j is $\mathbb{E}[\log \varepsilon_j \mid j, \mathbf{h}, \psi, k] = -\zeta \log \mathbb{P}_a(j \mid \mathbf{h}, \psi, k)$ where $\mathbb{P}_a(j \mid \cdot)$ refers to the probability of choosing j (see e.g., [Dubin and McFadden, 1984](#)).²⁰ We thus control for this selection term in the regression.²¹

In the last step of the inner algorithm, we estimate the occupational amenities $\boldsymbol{\mu}$ using a fast iterative procedure to match the occupational choices of terminal-age workers. At age $a = A$, equation (10) simplifies to

$$\mathbb{P}_A(j \mid \mathbf{h}, \psi, k) = \frac{\exp\left(\frac{1}{\zeta} (\log w_j(\mathbf{h}) + \mu_j - \kappa(k, j))\right)}{\sum_{l=1}^J \exp\left(\frac{1}{\zeta} (\log w_l(\mathbf{h}) + \mu_l - \kappa(k, l))\right)}. \quad (18)$$

where ζ (scale of productivity shocks), κ (switching costs) are included in outer vector $\boldsymbol{\theta}_1$ and thus taken as given in this step. The likelihood is maximized with respect to $\boldsymbol{\mu}$ when the observed share of terminal-age workers in occupation j , denoted s_j , equals the model-implied share $\tilde{s}_j(\boldsymbol{\mu})$. We solve for this $\boldsymbol{\mu}$ using the

¹⁹For estimation we set $\sum_{\tau \in \mathcal{T}_j} \tilde{\theta}_{j,\tau} = 1$ for all j and use O*NET's task-importance weights to proxy $\tilde{\theta}_{j,\tau}$. Since we estimate the model on pre-AI data, no task is automated, i.e., $\mathcal{A}_j = \emptyset$.

²⁰While [Dubin and McFadden \(1984\)](#) derives this selection equation for a static choice problem, it extends directly to dynamic choice settings like ours.

²¹We estimate these probabilities using occupation-specific logit regressions that condition on workers' previous occupation, 10-year age bins, and each dimension of initial skill.

contraction mapping proposed by [Berry et al. \(1995\)](#):

$$\mu_j^{(r+1)} = \mu_j^{(r)} + \Psi (\ln(s_j) - \ln(\tilde{s}_j(\boldsymbol{\mu}))) \quad \text{for some } \Psi \in (0, 1] \quad (19)$$

where Ψ is a damping parameter.²²

Outer algorithm. In the outer algorithm, we optimize over the outer vector $\boldsymbol{\theta}_1$ —the parameters governing skill accumulation, the initial skill distribution, and occupational choice. We choose these parameters to maximize the joint likelihood of the wage function and workers’ occupational choices at terminal age A :

$$\hat{\boldsymbol{\theta}}_1 = \arg \max_{\boldsymbol{\theta}_1} \sum_{i=1}^N \sum_{a=1}^A \sum_{j=1}^J 1[j_{i,a} = j] \log \pi (\log w_{i,a,j} - \log \tilde{w}_{i,a,j}(\boldsymbol{\theta}_1)) + \sum_{i=1}^N \sum_{j=1}^J 1[j_{i,A} = j] \log \mathbb{P}_A(j | \mathbf{h}_{i,A}, \psi_i, k_i; \boldsymbol{\theta}_1) \quad (20)$$

where $\tilde{w}_{i,a,j}(\boldsymbol{\theta}_1)$ is the expected wage based on the inner-step given $\boldsymbol{\theta}_1$ and $\pi(\cdot)$. The first term captures the likelihood of the observed wages and the second term reflects that of the observed occupational choices. In practice, wages contain noise that workers do not act on—such as measurement error or idiosyncratic pay—which the structural model’s choice-relevant shocks $\log \varepsilon_j$ alone cannot absorb without distorting the occupational choice probabilities. We therefore augment $\log \varepsilon_j$ with an independent Gaussian term $\log \nu_j$ and approximate the total wage shock $\log \varepsilon_j + \log \nu_j$ by a normal distribution $\pi(\cdot)$.²³

Full-population amenities. As a final step after the outer algorithm has converged, we re-estimate the occupational amenities $\boldsymbol{\mu}$ using the full working-age population rather than the old workers alone. This ensures that the model matches employment shares for all workers. The procedure follows the same BLP iteration, but now solves the worker’s full dynamic value function at each iteration.²⁴ The estimates of the occupational amenities resulting from the inner algorithm and this dynamic BLP procedure are almost identical (correlation: 0.99), validating our computational shortcut.

²²In practice, we use the SQUAREM algorithm to accelerate convergence ([Varadhan and Roland, 2008](#); [Reynaerts et al., 2012](#); [Conlon and Gortmaker, 2020](#)).

²³The estimated variance of $\log \nu_j$ is substantially larger than that of $\log \varepsilon_j$, so the convolution is dominated by the Gaussian component.

²⁴In dynamic problems, this procedure is not generally a contraction mapping and we can thus not prove that the fixed point is unique (see also [Gowrisankaran and Rysman, 2012](#)). However, the procedure yields the same results for any starting value we tried.

4.2 Parameter Estimates

Table 1 reports the estimates of the task-level production function parameters in equation (14). The first five columns show the general skill elasticities ω_s , which govern the degree to which higher skills raise productivity across all tasks regardless of those tasks' specific requirements. We find that the returns to math and social skills are highest, consistent with Deming (2017). The mismatch parameter η implies that if a worker's skill falls one unit below a task's requirement in a single dimension (on O*NET's 1–7 scale), their task-specific productivity is 4.3 percent lower than that in tasks for which they meet the skill requirements in every dimension. This cost of underqualification is substantial and generates meaningful comparative advantage across tasks with differing skill requirements.

In Table B.5, we show that the estimates are not strongly affected by the log-linearization of the wage function. Estimating the exact wage function by non-linear least squares yields skill elasticities and mismatch costs that are nearly identical to the baseline estimates. Controlling for selection more flexibly using splines also does not materially change the estimates.

TABLE 1: PRODUCTION FUNCTION: PARAMETER ESTIMATES

General skill					Mismatch
ω_{Mn}	ω_{Mt}	ω_S	ω_T	ω_V	η
0.330	0.786	0.545	0.220	0.361	0.044
(0.024)	(0.028)	(0.021)	(0.038)	(0.031)	(0.002)

Notes: This table shows estimates of the task-level productivity parameters. Subscripts Mn , Mt , S , T , V refer to manual, math, social, technical, and verbal, respectively. Estimates are obtained through OLS based on equation (17). Standard errors in parentheses (not corrected for uncertainty in other parameters).

The occupational prices are recovered from the occupational fixed effects in the production function regression. Consistent with the model's prediction that occupational prices reflect skill requirements, the estimated prices \hat{p}_j are strongly positively correlated with occupational skill requirements: skill requirements explain around 73 percent of the variance in prices across occupations (see Table B.6). To reduce noise in these price estimates, we apply empirical Bayes shrinkage toward the price predicted by skill requirements (Walters, 2024).

Table 2 presents the estimates governing initial skills and skill accumulation. The depreciation rate $\delta = 0.0003$ implies that skills are highly persistent when

workers only perform tasks for which they are overqualified. The learning cost $\lambda(\psi)$ is strictly decreasing in the AFQT quartile, implying that higher-ability workers converge more quickly toward the skill requirements of tasks for which they are underqualified. The shape parameters of the initial Beta distribution imply an average initial skill level of $B_a/(B_a + B_b) = 0.35$, corresponding to 3.11 on O*NET’s 1–7 scale, in between the “low” and “medium” skill requirement anchors. Figure A.6 plots the implied density function.

TABLE 2: SKILLS AND SKILL ACCUMULATION: PARAMETER ESTIMATES

Learning costs				Depr.	Initial dist.	
$\lambda(1)$	$\lambda(2)$	$\lambda(3)$	$\lambda(4)$	δ	B_a	B_b
3.50	2.97	2.81	2.67	0.0003	61.74	113.84

Notes: This table shows parameter estimates for the law of motion for skill accumulation in (16) and the initial skill distribution. $\lambda(\psi)$ refers to the learning cost at quartile ψ of the AFQT distribution.

Lastly, we estimate the scale parameter $\hat{\zeta} = 0.053$ and the switching cost parameter $\hat{\kappa} = 0.340$. The estimate for κ implies that the utility cost of switching occupations is equivalent to a 29 percent wage loss.

Calibrated parameters. Some parameters are calibrated externally rather than estimated. We set career length A to 40 so that each model period corresponds one year between ages 25 and 64. Following [Keane and Wolpin \(1997\)](#), the discount factor $\beta = 0.78$ (see also [Postel-Vinay and Robin \(2002\)](#) for similar estimates). We set the elasticity of substitution between tasks ρ to 0.49, as estimated by [Humlum \(2021\)](#), implying that tasks within an occupation are complements. The task sets \mathcal{T}_j , importance weights $\theta_{j,\tau}$, and the skill requirements are constructed from the data described in sections 3.1.1 and 3.1.2.

The labor share Γ_j in each occupation—the share of occupational revenue accruing to workers—is calibrated from the share of automatable tasks and the cost savings AI generates in those tasks. Appendix D.1 shows that this share satisfies

$$\Gamma_j \approx \frac{\chi^{\rho-1} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau}\right)}{\sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} + \chi^{\rho-1} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau}\right)}, \quad (21)$$

where χ is the unit cost of producing automatable tasks with AI relative to the unit cost of producing them with labor. Using experimental evidence, [Acemoglu \(2025\)](#)

estimates the cost savings of AI in automatable tasks to be 27 percent. Hence, we set $\chi = 0.73$.

We close the model by estimating the demand for each occupational good. Our baseline specification assumes that occupational goods are substituted with a constant elasticity of substitution (CES) σ , so that demand for occupation j satisfies $D_j(\mathbf{p}) \propto \alpha_j p_j^{-\sigma}$. We set the elasticity of substitution between occupations $\sigma = 1.57$ —the midpoint between 1.81 (Burstein et al., 2019) and 1.34 (Caunedo et al., 2023). The CES importance weights $\{\alpha_j\}_{j=1}^J$ are then identified from observed occupational wage shares and the estimated occupational prices because for any two occupations i and j ,

$$\frac{\alpha_i}{\alpha_j} = \left(\frac{p_i}{p_j} \right)^{\sigma-1} \times \frac{\text{Wage share of occupation } i}{\text{Wage share of occupation } j}$$

We compute occupational wage shares from the 2018 BLS Occupational Employment and Wage Statistics (OEWS) and use the occupational fixed effects from equation (17) as consistent estimates of log occupational prices. Because the supply-side parameters are estimated independently of demand, this demand specification does not affect any of the other parameter estimates.

In addition to this baseline CES structure for occupational demand, we estimate a nested CES structure. Occupations are combined to produce goods for each industry, where industries are characterized by the varying importance weights of occupations. These industries (3-digit NAICS) are nested within sectors (1-digit NAICS). For this nested CES, we set the elasticity of substitution between 1-digit sectors to 1.01, across 3-digit industries within sectors to 2.98 (as estimated by Hobijn and Nechio, 2019), and across occupation within industries to 0.90 (Goos et al., 2014). Each occupation’s weight is inferred from its wage bill share within an industry (analogous to the standard CES case). These occupation-level weights are then aggregated to obtain industry- and sector-level importance weights.

4.3 Model Fit

The model’s steady state moments fit labor market data well. The model’s moments are computed from a simulated sample of 100,000 workers living in the steady state before any technical change occurs. Figure 2 reports how well the moments from this simulated panel match the data.

First, Figure 2a shows that the model captures the unconditional distribution

of wages reasonably well. Given that some drivers of wage inequality, such as regional, racial, and gender differences, are omitted from the model, it is not surprising that inequality is somewhat underestimated. However, this underestimation is quite limited. For instance, the ratio between the 75th and the 25th percentile is 2.04 in the data, compared to 1.84 in the model. Table B.7 reports how additional inequality measures compare between the model and the data.

The model also accurately replicates patterns of occupational sorting. Figure 2b shows the correlation between average skill by occupation in the model and the NLSY79 data. We compute these correlations using young workers for which we observe the skills directly from ASVAB scores (rather than partially inferred through the skill accumulation function), making the test as stringent as possible. The correlations range between 0.6 and 0.8 across skill dimensions, implying that the skill assessment scores in the NLSY79 are predictive of occupational choices (see also [Lise and Postel-Vinay, 2020](#)) and that workers in the model select into occupations based on their skills in ways similar to that observed in the NLSY79.

Average wages by occupation and the wage gradient by age also show a reasonable fit. Figure A.7a shows that the average wage by occupation matches the data almost perfectly. This is unsurprising, since the estimation of occupational demand and amenities directly targets total wages and employment by occupation. Figure A.7b shows that the model matches the growth rate of wages from labor market entry to around age 55. However, the wage pattern in the model is not as concave as in the data, so that growth in the first years is underestimated and growth in the last 10 years overestimated.²⁵


The model also generates a realistic degree of occupational persistence. The probability of remaining in the same 3-digit occupation is 0.86 in the model and 0.90 in CPS data. This moment is directly targeted by the switching cost κ . The model additionally reproduces the untargeted fact that staying within a broader 2-digit occupational group is even more likely: 0.92 (model) and 0.94 (CPS).

Lastly, we compare the transition probabilities between occupations conditional on switching. The correlation between the (log of) the transition probabilities in the model and data is 0.56 on the 2-digit occupation level. On the 3-digit level, it is substantially lower: 0.20. The model thus accurately predicts occupational transitions across 22 broader occupational groups. Within those groups, occupa-

²⁵Besides, the model predicts markedly higher wages in the first period than those directly after. This feature is caused by the fact that occupational switching costs are only incurred after the first period. Workers are therefore more likely to choose occupations in which they are highly productive (i.e., with a high ε_j) in the first period than in any later periods.

tional transitions are harder to predict, because occupations are more similar in skill requirements within those groups.

FIGURE 2: MODEL FIT: COMPARING MODEL MOMENTS WITH DATA



Notes: Panel A shows a kernel density plot of the wage distributions of the NLSY79 data and the model's steady state. Panel B reports the correlation between the average skill by occupation in the model's first period and the NLSY79.

5 Artificial Intelligence and the Labor Market

Having estimated the model and validated its ability to match various features of the pre-AI labor market, we next use the model to assess AI's labor market effects. Specifically, we study AI's general equilibrium effects on wages, inequality, welfare, skill returns, and occupations using the predictions of the technology's capabilities to augment, automate, and simplify.

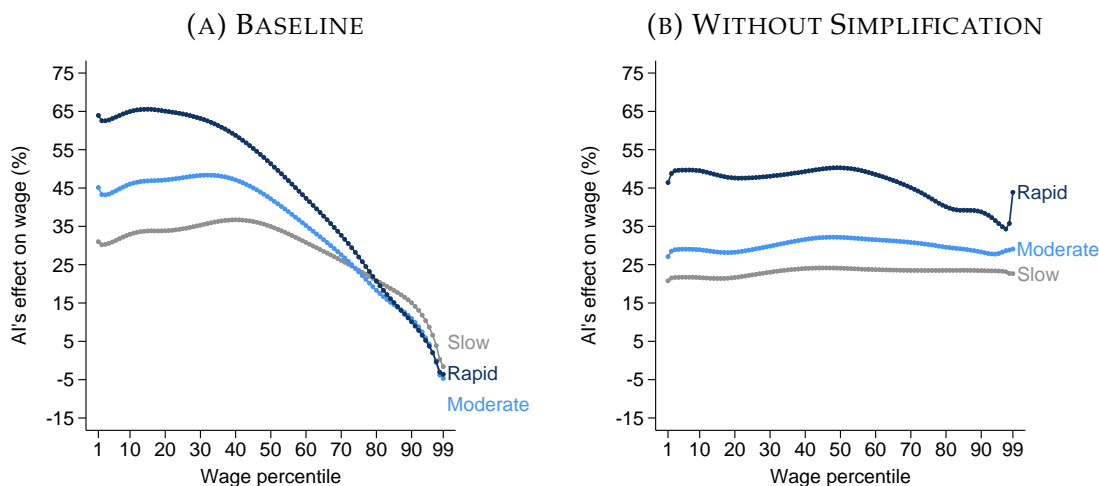
5.1 AI's Effects on Individual Workers

5.1.1 Wages and Inequality

We begin by studying AI's impact on the steady state wage distribution. We find that AI produces sizable average wage gains of 27.6 percent (see Figure 3). This effect is larger than estimates of IT's cumulative growth contribution of around 16 percent (Jorgenson et al., 2008). It is also similar to other estimates on AI's growth

effect (Baily et al., 2023; Aghion and Bunel, 2024; Filippucci et al., 2024) but larger than the estimate by Acemoglu (2025).²⁶

FIGURE 3: WAGE EFFECTS ACROSS THE DISTRIBUTION



Notes: This figure shows the distribution of wage changes induced by AI across the wage percentile distribution for three scenarios of AI progress: slow, moderate, and rapid. We follow Karger et al. (2026) in defining the scenarios and measure task-specific AI capabilities using Qwen 2.5. The left panel shows the joint effect of AI’s augmentation, automation, and simplification on each wage percentile; the right panel shows the effects if AI induced only augmentation and automation, but no simplification. The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile.

The three mechanisms of technical change shape wages through competing forces (see Figure A.9). Augmentation has an unambiguously positive effect on wages, and we find it to be quantitatively large (22% on average). Automation raises productivity but displaces labor with capital (Acemoglu and Restrepo, 2018); here the productivity effect dominates, raising average wages by 8 percent.²⁷ Simplification weakly increases productivity (see equation 14), but limits opportunities for learning; for AI, we find that the net effect is slightly negative (−3%). Workers’ skill levels decrease significantly on average, especially for verbal and technical skills (see Figure A.10). Our modeling of learning cost is consistent with experimental evidence that AI usage slows the acquisition of new skills (Shen and Tamkin, 2026).

The wage gains are concentrated at the bottom of the distribution and are nearly zero at the 99th percentile, implying a significant drop in wage inequality. In the

²⁶The productivity effects quantified by Acemoglu (2025) are based on a conservative automation scenario, whereas our evidence also incorporates augmentation.

²⁷This finding is primarily driven by the cost savings of automation, which, following Acemoglu (2025), we calibrate to 27%.

moderate scenario, the 90-10 ratio of wages falls by 24 percent. In the slow and rapid AI scenarios, it falls by 13 and 33 percent, respectively.

Simplification is the key force lowering inequality. Without simplification—that is, with only augmentation and automation operating—inequality would slightly *increase* rather than strongly *decrease*. Simplification lowers inequality in two ways. First, it reduces wage dispersion within occupations by enabling lower-skill workers to perform tasks more productively. Second, it reduces wage differences across occupations by making occupations with high skill requirements more accessible to lower-skill workers, reducing its relative price.

In contrast, automation and augmentation have small distributional effects. First, within occupations, augmentation and automation affect the relative productivity of workers with different skills only if they disproportionately affect tasks with atypical skill requirements. In such cases, augmentation and automation induce an indirect form of simplification by changing tasks' effective importance within an occupation. However, we find that these indirect effects are quantitatively small relative to direct simplification. Second, across occupations, augmentation and automation can in principle increase wage inequality by increasing the relative productivity of high-wage occupations. However, when the elasticity of substitution across occupations is close to one—as empirical estimates suggest (e.g., [Burstein et al., 2019](#); [Caunedo et al., 2023](#))—such productivity differences translate only weakly into relative wage changes.

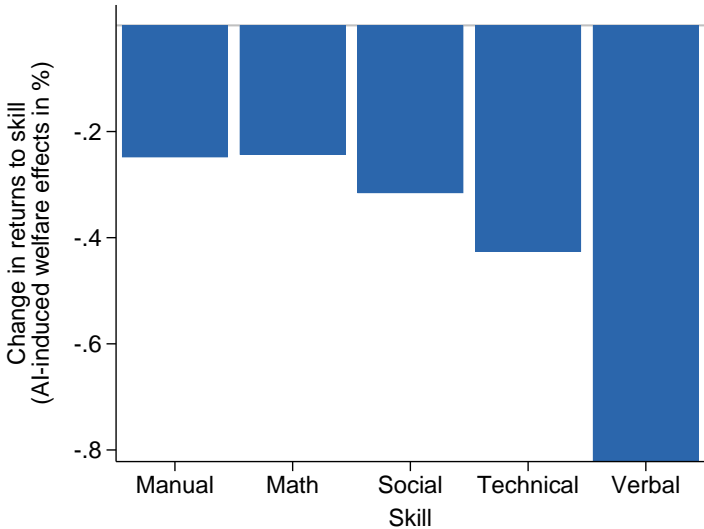
The estimates of AI's effect on average wage growth are comparable to economists' forecasts of AI's impact on growth ([Karger et al., 2026](#)). On average, economists predict that moderate AI progress would raise annual growth from 2.4 to 2.9 percent, implying 23.4 percentage points additional growth by 2050—close to the model's prediction of 27.6 percent. The general public's forecasts imply 28.5 percent additional growth, AI experts' 85.7 percent.

5.1.2 Returns to Skills

Consistent with the decrease in wage inequality, we find that the ex-ante welfare gains are largest for lower-skill workers. Specifically, we compare expected welfare at labor market entry, conditional on initial skills, in economies with and without AI. We measure welfare changes using equivalent variation, measured in terms of a permanent proportional wage increase that delivers the same welfare gain as the introduction of AI (regardless of skill and occupation). [Figure 4](#) shows the coefficients of a regression of the welfare gains on initial skill levels. Workers

with high verbal skills see the smallest increases in welfare gain: a 1-point increase in verbal skills (on the O*NET scale from 1 to 7) decreases the welfare gains from AI by 0.8 percent. Higher manual and math skills, in contrast, have the least negative association with AI-induced welfare effects. While qualitatively similar, the decreases in the returns to skills are quantitatively starker in the rapid than in the slow and moderate AI scenarios (see Figure A.18).

FIGURE 4: HOW AI CHANGES THE RETURNS TO SKILLS



Notes: This figure shows how AI’s welfare effects differ by skills (moderate AI scenario). The welfare effects are measured in equivalent permanent percentage wage increases. This figure plots the coefficient of a regression of these welfare effects on skill levels across all dimensions. For interpretability, the skills are expressed on the O*NET scale from 1 to 7. Figure A.18 shows the slow and rapid AI scenarios.

Unlike the skill-biased technical changes of recent decades, which increased the relative demand for cognitive skills and widened wage differentials (Katz and Murphy, 1992; Autor et al., 2008), AI’s distributional effects are dominated by simplification, which compresses the wage distribution. Simplification is similar in spirit to “non-autonomous AI”, which allows lower-skill workers to solve difficult problems with AI, and thus raises their wages relative to skilled workers (Ide and Talamàs, 2025).

The decline in the returns to verbal and social skills has implications beyond overall inequality. Hurst et al. (2024) show that rising returns to abstract, cognitive tasks—closely related to verbal and social skills—together with discrimination of Black workers explain why the Black-White wage gap stopped closing after 1980. Our finding that AI compresses these returns suggests that simplification could

partially reverse this.

Lastly, we merge these changes in skill returns to data on college majors' skill content to analyze how the value of each major changes (see Figure A.19). Majors high in manual and technical skill—especially, mechanics, engineering, culinary arts, cosmetology, architecture, transportation, and physics—rank among the top majors whose returns rise relative to other majors by over 1.5 percentage points. In contrast, majors intensive in verbal and social skills fare worse, especially political science, religion, women's studies, sociology, philosophy, and law.

Our key result is not sensitive to the structure of demand for occupations. Comparing standard and nested CES demand, we find that the distributional impacts are very similar (see Figure A.8). On average, wages increase slightly more with the nested versus the standard CES.

5.1.3 Welfare

While wages are a natural measure of AI's labor market impact, they capture only part of workers' overall welfare. First, wages abstract from the switching cost workers bear to reach the new equilibrium. Second, even if AI raises average wages, its distributional consequences matter: under standard concavity of utility, wage gains accruing to lower-income workers generate larger welfare improvements than equivalent gains at the top of the distribution.

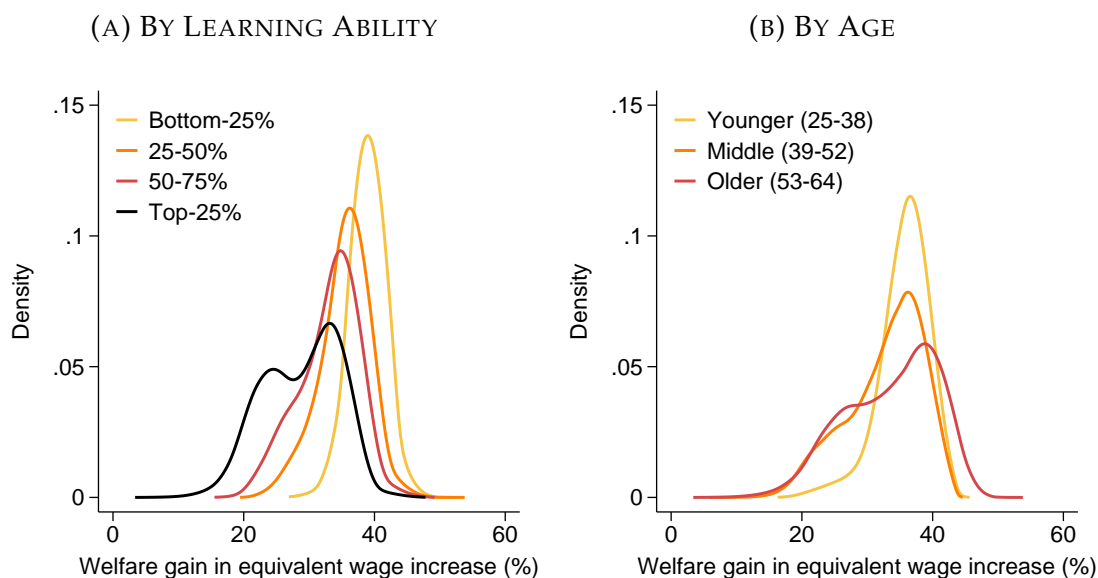
We first consider AI's welfare effects in the short run. Specifically, we consider a one-off ("MIT") shock through which AI is introduced to the economy. We measure welfare effects in terms of equivalent wage variation. That is, we ask: Absent AI, how much would a worker have to be paid more for the rest of their life to make them equally well off as they are in the transition to AI.

We find that welfare increases despite the reallocation costs that workers face during the transition to AI—for most of them substantially (see Panel A of Figure A.11). Perhaps unsurprisingly, AI affects older workers' welfare particularly heterogeneously (see Figure 5). Older workers have already invested in occupations and skills and have less time left to benefit from costly adjustments. Thus, those who have invested in the "wrong" occupations and skills gain less than others.

We also consider scenarios in which AI does not provide any augmentation, automation, or simplification. We find that a small set of workers experience negative welfare effects in the scenario where AI does not augment the worker at all (see Figure A.12 for the distribution under each scenario). However, even in that relatively pessimistic scenario the average worker experiences an 11 percent wel-

fare gain. In all other scenarios, we find positive welfare effects across the board.

FIGURE 5: SHORT-RUN WELFARE EFFECTS BY WORKER CHARACTERISTICS



Notes: This figure shows heterogeneity in AI’s welfare effects along the transition path by worker characteristics (moderate AI scenario). Panel (A) shows welfare effects by learning ability (AFQT score). Panel (B) shows welfare effects by age. The welfare effect is measured in equivalent wage variation: it represents the permanent wage increase across all occupations that yields the same welfare gain as the introduction of AI.

Last, we consider the welfare effects in the long run for workers entering the labor market in the new steady state.²⁸ We find that long-run welfare increases by 37 percent on average, ranging between 30 and 42 percent for most workers (see Panel B of Figure A.11). The welfare gains exceed the average wage increase (28 percent) because of reduced inequality: AI disproportionately raises wages at the bottom of the distribution, where the marginal utility of consumption is highest. The average welfare gain is higher in the long run than in the short run because the transition requires costly reallocation. The most important difference is that welfare gains are more unequally distributed in the short run.

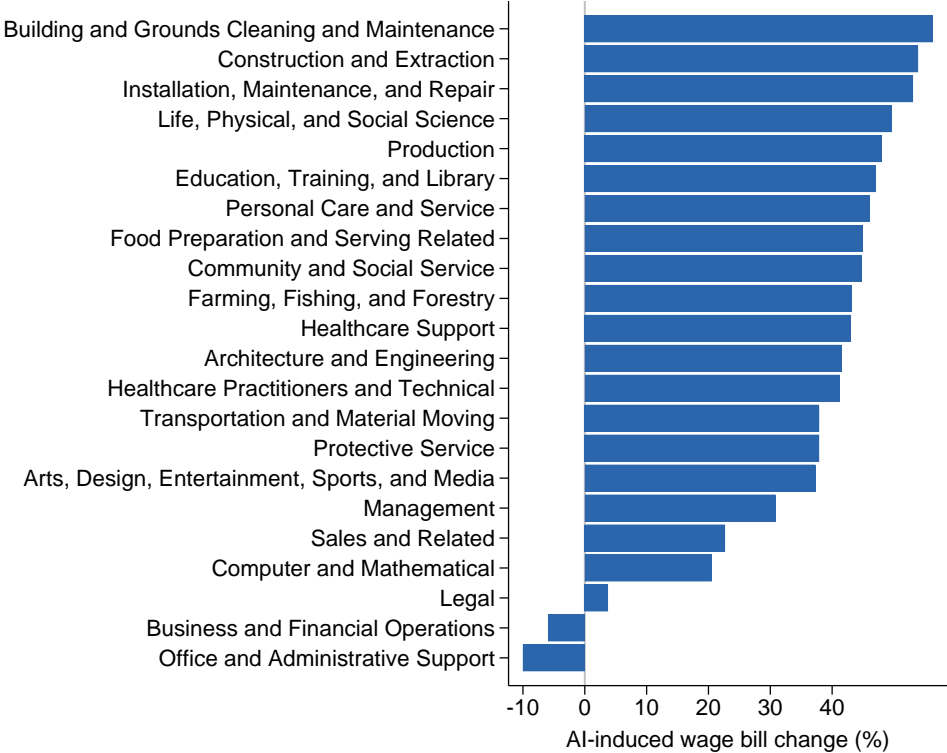
5.2 AI’s Effect on the Occupational Landscape

Workers’ response to AI significantly shifts the occupational landscape. These shifts are reflected in the changing importance of individual occupations as measured by their wage bill (employment times wages). *Building and Maintenance* as

²⁸The welfare gain is expressed in equivalent wage variation as explained above.

well as *Construction and Extraction* occupations experience the largest wage bill increase of 50 percent, while *Office and Administrative Support* as well as *Business and Financial Operations* see an absolute decline in their wage bill (see Figure 6).

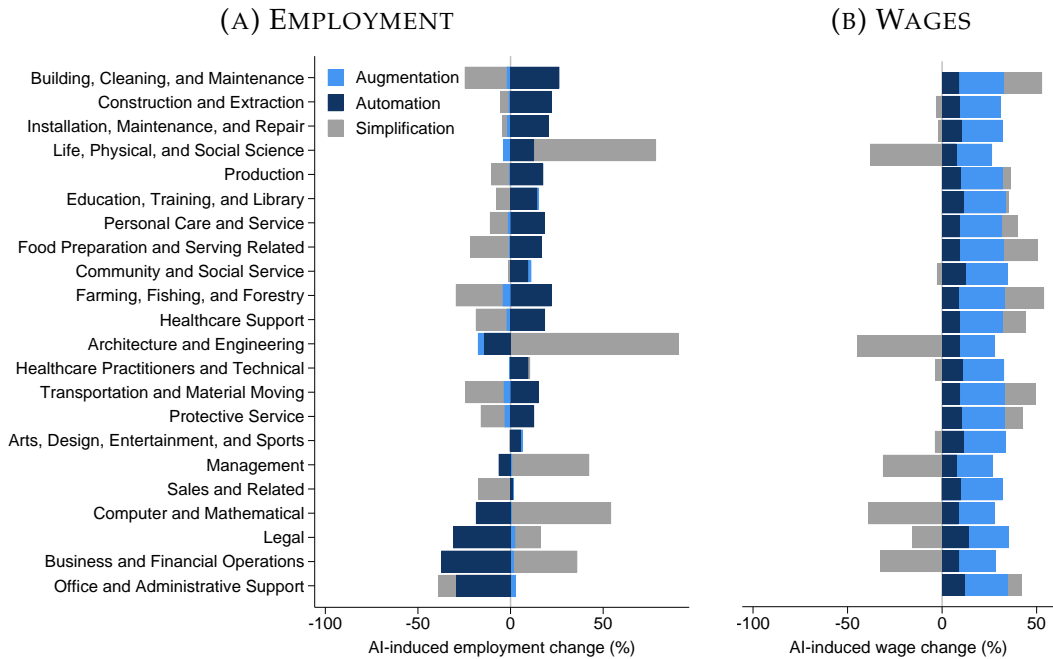
FIGURE 6: AI’S EFFECT ON OCCUPATIONS’ WAGE BILLS



Notes: This figure shows the model predictions on AI’s wage bill effects by occupational group (moderate AI scenario). Figure A.15 shows these effects on the disaggregated occupations used in the model.

Occupations’ wage bill changes are driven by both wages and employment (see Figure A.13). Although wages rise on average, some occupations experience absolute wage declines. The effects on employment and wages often work in the opposite direction (Figure A.14). For instance, *Architecture and Engineering* as well as *Life, Physical, and Social Science* experience the largest increase in employment share and the largest decrease in average wages. In contrast, *Building Cleaning and Maintenance* experiences the largest increase in average wages and a decline in employment. Note that because of workers’ reallocation, occupation-level wage declines do not imply that individual workers experience wage losses. Since some occupations lose more than half of their employment, considering worker reallocation is key to understand how AI affects workers (Figure A.13c).

FIGURE 7: AI'S EFFECT ON OCCUPATIONAL EMPLOYMENT & WAGES



Notes: This figure shows the model's predictions on AI's employment and wage effects by occupational group (moderate AI scenario). Occupations are sorted in descending order of AI's effect on the total wage bill, so that the first listed occupation experiences the largest wage bill increase. We conduct a Shapley-Owen decomposition to separate the overall change into the contribution of each channel: augmentation, automation, and simplification.

We find that simplification shapes relative wage effects more than augmentation and automation. We decompose occupational outcomes by recomputing them under all possible channel combinations (Shapley, 1953; Owen, 1977, see Figure 7). Augmentation and automation tend to increase the productivity of all workers in an occupation equally, unless the skill requirements of affected tasks greatly differ from the occupational average. Small automation- and augmentation-driven relative wage effects imply that such *indirect* simplification has limited quantitative importance for AI.

In contrast to augmentation and automation, *direct* simplification has strong effects on relative wages. Simplification expands the pool of workers who can perform an occupation productively by lowering skill requirements. It thereby raises affected occupations' employment and compresses their average wages through increased competition. Figure A.16 shows this breakdown for detailed occupational categories.

Three occupations illustrate the mechanisms behind these predictions: radiologists, management analysts, and telemarketers.

Radiologists. In 2016, Geoffrey Hinton urged the field to “stop training radiologists,” predicting AI would replace them within five years. Radiology has since accounted for more than 75 percent of FDA-authorized clinical AI tools, and about two-thirds of US radiology departments use AI (Mousa, 2025). Yet the radiologist labor market has grown: the wage bill share rose by 6.6 percent between 2016 and 2024, driven by strong employment growth (23.2% versus 9.8% on average) and held back by below-average wage growth (30.1% versus 36.9%). The model lines up with this pattern. It predicts a 41 percent wage bill increase for Health Diagnosing and Treating Practitioners, 1.5 times the average (Figure A.15), combining above-average employment growth (11% versus 0% by construction) with near-average wage growth (27% versus 28%; Figure A.16). Simplification drives both: it raises employment and compresses radiologists’ relative wages. Automation affects only a small share of their tasks, adding to employment with little wage effect; augmentation matters little for either.

Management analysts. Management analysts (Business Operations Specialists) shrink in importance: the model predicts a 3 percent employment decline and flat wages (Figure A.16)—below the cross-occupation average on both margins. They are strongly exposed to both simplification and automation: simplification raises employment and lowers wages, while automation offsets the employment gain, so the occupation shrinks overall. This matches experimental evidence that lower-skill consultants gain more from AI (Dell’Acqua et al., 2023); in the model, such gains can only come from (direct or indirect) simplification, which raises the relative productivity of lower-skill workers.

Telemarketers. For telemarketers, AI substitutes for rather than complements labor: all 12 of their tasks are automatable by generative AI. Their broader group, *Other Sales and Related Workers*, also faces high automation exposure, sharply reducing employment and placing telemarketers among the most negatively affected occupations by both employment and wage bill (Figures A.15 and A.16). Simplification plays little role, so relative wages stay constant.

5.3 Accounting for Adoption Costs

While AI has the potential to reduce wage inequality, adoption frictions may limit this equalizing effect. Specifically, it may be most costly to adopt the technology for workers with the lowest learning ability (Humlum and Vestergaard, 2025; Bick

et al., 2026)—the workers who stand to benefit the most. In this subsection, we consider how our findings change when we take adoption costs into account.

We model adoption as a firm-level decision taken separately for each worker. Adopting AI for a worker with learning ability ψ shrinks output by a proportional factor $c(\psi) \in [0, 1)$ that captures the time, training, and reorganization required to put the technology to productive use. After paying the adoption cost, the firm benefits from augmentation, automation, and simplification. The firm adopts the technology for a worker with skills \mathbf{h} in occupation j whenever doing so makes the worker more productive despite the adoption costs, i.e.,

$$(1 - c(\psi)) Y_j^A(\mathbf{h}) \Gamma_j > Y_j^N(\mathbf{h}) \quad (22)$$

where Y_j^A and Y_j^N are output with and without AI adoption and Γ_j is the labor share that remains after automation costs as in equation (7). Since adoption affects a worker’s time allocation across tasks, it affects their skill accumulation accordingly (equation 16).

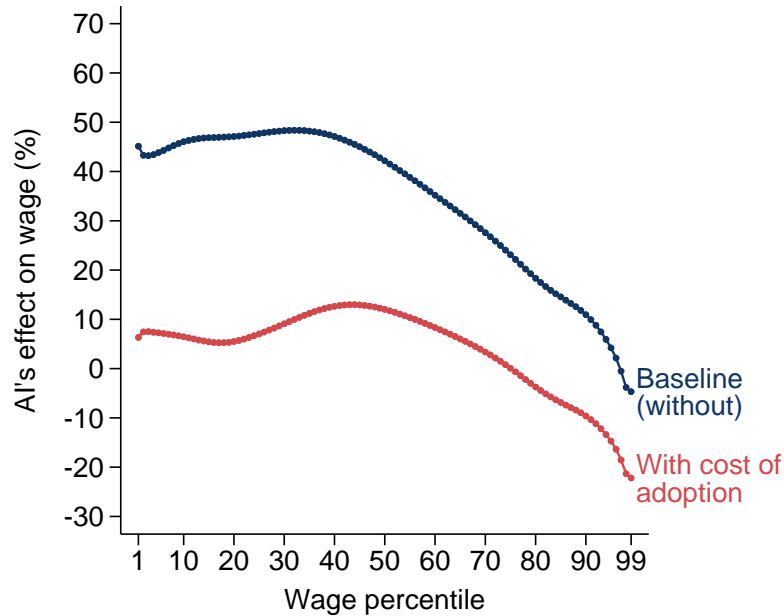
We estimate adoption costs by learning ability using micro-data on AI adoption at work provided by Bick and Blandin (2023) and Bick et al. (2026).²⁹ We only consider variation in adoption costs within occupations, so as not to confound differential adoption costs with differential productivity effects across occupations. Specifically, we regress a dummy for work-related generative AI usage on fixed effects for each occupation and each quartile of the learning ability distribution, relative to the first quartile. We proxy learning ability with education. We then find the vector of adoption costs that implies the same regression coefficients and average adoption rate in the model, holding the distribution of skills and learning ability by occupation constant at the pre-AI equilibrium.³⁰

The resulting estimates for the adoption costs are from the lowest to the highest quartile of learning ability ψ : 0.28, 0.26, 0.24, and 0.22. These estimates imply that adoption costs are lower for workers with higher learning ability and that they average around a quarter of a worker’s effective units of labor. Since generative AI is still a relatively young technology, one may expect that adoption costs decrease over time. In that sense, we view the estimated productivity effects of AI under these adoption costs as conservative.

²⁹These individual-level data on generative AI usage at work are from the Real-Time Population Survey (RPS), a nationally representative survey of US adults aged 18–64 (Bick and Blandin, 2023). We pool the June, August, and November 2024 waves, which included a generative AI module (Bick et al., 2026).

³⁰To estimate these adoption costs, we compute $Y_j^A(\mathbf{h})$ based on AI’s capabilities under the “slow” scenario, as this scenario is closest to AI’s current capabilities.

FIGURE 8: ADOPTION COSTS AND AI'S WAGE EFFECTS



Notes: This figure shows the distribution of wage changes induced by generative AI across the wage percentile distribution with versus without adoption costs (both for the moderate AI scenario). Estimates without adoption costs replicate the “moderate” estimates in Panel A of Figure 3. We follow Karger et al. (2026) in defining the scenarios and measure task-specific AI capabilities using Qwen 2.5. Figure A.17 shows estimates for the slow and rapid scenarios. The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile. Results are similar when we change the learning ability gradient of adoption cost by a factor of 0.5 or 2.

Our baseline estimates of adoption costs suggest that AI’s equalizing potential is reduced but remains intact (see Figure 8). First, even though adoption costs are lower for workers with higher ability, the model still predicts that AI reduces wage inequality substantially. The 90–10 ratio of wages still falls by 15 percent (versus 24 percent without adoption costs). Adoption costs have qualitatively similar effects under the slow and rapid AI scenarios (see Figure A.17).

Second, the adoption costs reduce the estimated wage effect from AI since workers who adopt pay a large productivity cost of doing so and some workers choose not to adopt. After accounting for adoption costs, our estimate of AI’s effect on average wage growth—0.3 percent (versus 27.6% without adoption costs)—is closer to short-run than long-run forecasts (Karger et al., 2026). The short-run forecasts are the more modest of the two, projecting additional growth of 1.1 to 3.3 percent by 2030 (versus 23.4% to 85.7% by 2050).

In sum, adoption costs attenuate AI’s equalizing effect but do not eliminate it, even though they affect workers with lower learning ability most. Lowering these costs is thus a key policy lever for realizing AI’s equalizing potential.

6 Conclusion

Technical change reorganizes production at the task level, so understanding its labor market effects requires characterizing workers' comparative advantage across occupations and tasks. This paper develops and estimates a dynamic task-based framework that recovers this comparative advantage and embeds it in a general equilibrium model of occupational choice and skill accumulation. We use this framework to study artificial intelligence as a technology that augments, automates, and simplifies tasks. The quantified model predicts that AI substantially raises wages, especially in the lower part of the wage distribution. Simplification of tasks is the key reason why our model predicts that AI will reduce inequality. However, simplification also reduces skill accumulation, dampening AI's positive effect on wage growth.

Our work has several limitations that future research may be able to overcome. First, our estimates are based on our expectations (as of 2026) of AI's future capabilities. While we bracket this uncertainty with slow, moderate, and rapid scenarios, the underlying range itself may shift as the technology develops. Second, our model does not feature unemployment. While we find that AI generates a large reallocation of workers across occupations, we ignore frictions that may generate transitional unemployment accompanying this reallocation. Third, this paper only considers the effect of technical change on wage income. Beyond wages, AI also has the potential to affect inequality through the incomes of entrepreneurs and capital owners. On the one hand, small-scale entrepreneurs specifically may benefit from the low fixed cost of using AI, making them more competitive with large-scale entrepreneurs, lowering inequality. On the other hand, producing AI requires high fixed costs, which can increase firm concentration and raise inequality (Reichardt, 2025). Furthermore, the owners of capital may benefit disproportionately if the returns to capital rise faster than wages (Moll et al., 2022).

Future research may also be able to address several additional dimensions we hold fixed. In our framework, we take the technical change brought about by AI as exogenous. One could, however, consider how simplifying technologies arise from directed innovation when particular skills are in short supply (Acemoglu, 2002; Acemoglu and Restrepo, 2018). Similarly, we hold the organization of work fixed. In practice, technological change may prompt the emergence of new tasks and a reorganization of how work is bundled into occupations. While our methodology can accommodate new tasks, measuring what such new tasks (and their skill requirements) might be is particularly challenging. Last, given that we find

empirical evidence for AI's impact on college majors, it is useful to incorporate how such educational choices respond to technical change in the model (Heckman et al., 1998).

References

- ACEMOGLU, D. (2002): "Directed technical change," *Review of Economic Studies*, 69, 781–809.
- (2025): "The simple macroeconomics of AI," *Economic Policy*, 40, 13–58.
- ACEMOGLU, D. AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of labor economics*, Elsevier, vol. 4, 1043–1171.
- ACEMOGLU, D. AND P. RESTREPO (2018): "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.
- (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of economic perspectives*, 33, 3–30.
- (2022): "Tasks, Automation, and the Rise in US Wage Inequality," *Econometrica*, 90, 1973–2016.
- ADDISON, J. T., L. CHEN, AND O. D. OZTURK (2020): "Occupational Skill Mismatch: Differences by Gender and Cohort," *ILR Review*, 73, 730–767.
- AGHION, P. AND S. BUNEL (2024): "AI and Growth: Where Do We Stand?" Working paper.
- ASIRVATHAM, H., E. MOKSKI, AND A. SHLEIFER (2026): "GPT as a Measurement Tool," Working paper.
- AUTOR, D. AND N. THOMPSON (2025): "Expertise," *Journal of the European Economic Association*, 23, 1203–1271.
- AUTOR, D. H. AND M. J. HANDEL (2013): "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages," *Journal of labor Economics*, 31, S59–S96.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2008): "Trends in U.S. Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 90, 300–323.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118, 1279–1333.
- BAILY, M., E. BRYNJOLFSSON, AND A. KORINEK (2023): "Machines of mind: The case for an AI-powered productivity boom," .

- BALEY, I., A. FIGUEIREDO, AND R. ULBRICHT (2022): “Mismatch Cycles,” *Journal of Political Economy*, 130, 2943–2984.
- BECKER, J., N. RUSH, E. BARNES, AND D. REIN (2025): “Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity,” Working paper.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- BICK, A. AND A. BLANDIN (2023): “Employer reallocation during the COVID-19 pandemic: Validation and application of a do-it-yourself CPS,” *Review of Economic Dynamics*, 49, 58–76.
- BICK, A., A. BLANDIN, AND D. J. DEMING (2026): “The rapid adoption of generative AI,” *Management Science*.
- BOPPART, T., P. KRUSELL, AND K. MITMAN (2018): “Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative,” *Journal of Economic Dynamics and Control*, 89, 68–92.
- BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025): “Generative AI at Work,” *Quarterly Journal of Economics*, 140, 889–942.
- BURSTEIN, A., E. MORALES, AND J. VOGEL (2019): “Changes in Between-Group Inequality: Computers, Occupations, and International Trade,” *American Economic Journal: Macroeconomics*, 11, 348–400.
- CAMPERO, A., M. VACCARO, J. SONG, H. WEN, A. ALMAATOUQ, AND T. MALONE (2022): “A Test for Evaluating Performance in Human-Computer Systems,” Working paper.
- CAUNEDO, J., D. JAUME, AND E. KELLER (2023): “Occupational Exposure to Capital-Embodied Technical Change,” *American Economic Review*, 113, 1642–1685.
- CLARKE, S. AND B. HANRAHAN (2024): “Comparing the Effect of Different Task Types on Effective Use of GitHub Copilot,” in *Generative AI in Real-World Workplaces: The Second Microsoft Report on AI and Productivity Research*, ed. by S. Jaffe, N. P. Shah, J. Butler, A. Farach, A. Cambon, B. Hecht, M. Schwarz, and J. Teevan.
- CONLON, C. AND J. GORTMAKER (2020): “Best Practices for Differentiated Products Demand Estimation with pyblp,” *The RAND Journal of Economics*, 51, 1108–1161.
- CRUCES, G., D. FERNÁNDEZ MEIJIDE, S. GALIANI, R. H. GÁLVEZ, AND M. LOMBARDI (2026): “Does Generative AI Narrow Education-Based Productivity Gaps? Evidence from a Randomized Experiment,” Working paper.
- CUI, Z. K., M. DEMIRER, S. JAFFE, L. MUSOLFF, S. PENG, AND T. SALZ (2024): “The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers,” Working paper.

- DANIELI, O. (2025): "Skill-Replacing Technology and Bottom-Half Inequality," Working paper.
- DELL'ACQUA, F., C. AYOUBI, H. LIFSHITZ, R. SADUN, E. MOLLICK, L. MOLLICK, Y. HAN, J. GOLDMAN, H. NAIR, S. TAUB, AND K. LAKHANI (2025): "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise," Working Paper.
- DELL'ACQUA, F., E. MCFOWLAND III, E. R. MOLLICK, H. LIFSHITZ-ASSAF, K. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2023): "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Working Paper.
- DELOACH, S. B., M. KURT, AND R. SANSALE (2022): "Non-cognitive mismatch and occupational switching," *Journal of Behavioral and Experimental Economics*, 97, 101846.
- DEMING, D. J. (2017): "The Growing Importance of Social Skills in the Labor Market," *Quarterly Journal of Economics*, 132, 1593–1640.
- DILLON, E. W., S. JAFFE, N. IMMORLICA, AND C. T. STANTON (2025): "Shifting Work Patterns with Generative AI," Working paper.
- DIX-CARNEIRO, R. (2014): "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 82, 825–885.
- DOWNEY, M. (2021): "Partial Automation and the Technology-Enabled Deskilling of Routine Jobs," *Labour Economics*, 69, 101973.
- DUBIN, J. A. AND D. L. MCFADDEN (1984): "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption," *Econometrica: Journal of the Econometric Society*, 345–362.
- EDELMAN, B., J. BONO, S. PENG, R. RODRIGUEZ, AND S. HO (2024): "Randomized Controlled Trial for Microsoft Security Copilot," Working paper.
- ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2024): "GPTs are GPTs: Labor Market Impact Potential of LLMs," *Science*, 384, 1306–1308.
- FILIPPUCI, F., P. GAL, AND M. SCHIEF (2024): "Miracle or Myth? Assessing the macroeconomic productivity gains from Artificial Intelligence," Tech. rep., OECD Publishing.
- FREUND, L. B. AND L. F. MANN (2025): "Job Transformation, Specialization, and the Labor Market Effects of AI," Working paper.
- GAMBACORTA, L., H. QIU, S. SHAN, AND D. M. REES (2024): "Generative AI and Labour Productivity: A Field Experiment on Coding," Working paper.

- GOOS, M., A. MANNING, AND A. SALOMONS (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104, 2509–2526.
- GOWRISANKARAN, G. AND M. RYSMAN (2012): “Dynamics of Consumer Demand for New Durable Goods,” *Journal of Political Economy*, 120, 1173–1219.
- GUVENEN, F., B. KURUSCU, S. TANAKA, AND D. WICZER (2020): “Multidimensional Skill Mismatch,” *American Economic Journal: Macroeconomics*, 12, 210–244.
- HAMPOLE, M., D. PAPANIKOLAOU, L. D. SCHMIDT, AND B. SEEGMILLER (2025): “Artificial intelligence and the Labor Market,” Working paper.
- HECKMAN, J. J., L. LOCHNER, AND C. TABER (1998): “Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents,” *Review of Economic Dynamics*, 1, 1–58.
- HOBijn, B. AND F. NECHIO (2019): “Sticker Shocks: Using VAT Changes to Estimate Upper-Level Elasticities of Substitution,” *Journal of the European Economic Association*, 17, 799–833.
- HOSSEINI, S. M. AND G. LICHTINGER (2026): “Generative AI and Occupational Entry Barriers: the Labor-Supply Channel of Technological Change,” Working paper.
- HUMLUM, A. (2021): “Robot Adoption and Labor Market Dynamics,” *Princeton University*.
- HUMLUM, A. AND E. VESTERGAARD (2025): “The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities Among Workers,” *Proceedings of the National Academy of Sciences*, 122, e2414972121.
- HURST, E., Y. RUBINSTEIN, AND K. SHIMIZU (2024): “Task-Based Discrimination,” *American Economic Review*, 114, 1723–1768.
- IDE, E. AND E. TALAMÀS (2025): “Artificial Intelligence in the Knowledge Economy,” *Journal of Political Economy*, 133, 3762–3800.
- (2026): “The Turing Valley: How AI Capabilities Shape Labor Income,” *arXiv preprint arXiv:2408.16443*.
- JORGENSEN, D. W., M. S. HO, AND K. J. STIROH (2008): “A Retrospective Look at the US Productivity Growth Resurgence,” *Journal of Economic Perspectives*, 22, 3–24.
- KANAZAWA, K., D. KAWAGUCHI, H. SHIGEOKA, AND Y. WATANABE (2022): “AI, Skill, and Productivity: The Case of Taxi Drivers,” Working paper.
- KARGER, E., O. KUUSELA, J. ABALUCK, K. BRYAN, B. HALPERIN, T. JONES, C. MURPHY, P. TRAMMELL, M. REYNOLDS, D. MAYLAND, R. VISWANATHAN, A. MITTAL, R. CEPPAS DE CASTRO, J. ROSENBERG, AND P. E. TETLOCK (2026): “Forecasting the Economic Effects of AI,” Working paper.

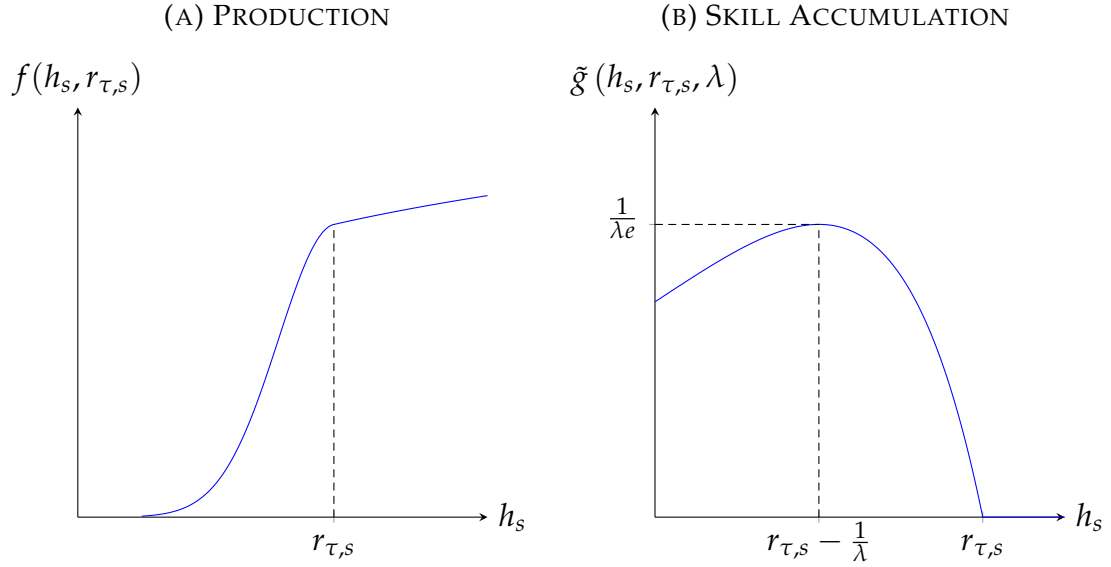
- KATZ, L. F. AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107, 35–78.
- KEANE, M. P. AND K. I. WOLPIN (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105, 473–522.
- LEE, D. AND K. I. WOLPIN (2006): “Intersectoral Labor Mobility and the Growth of the Service Sector,” *Econometrica*, 74, 1–46.
- LINDENLAUB, I. (2017): “Sorting Multidimensional Types: Theory and Application,” *Review of Economic Studies*, 84, 718–789.
- LISE, J. AND F. POSTEL-VINAY (2020): “Multidimensional Skills, Sorting, and Human Capital Accumulation,” *American Economic Review*, 110, 2328–2376.
- LOWHAGEN, N., P. SCHWENDENER, AND T. NETLAND (2025): “Can a Troubleshooting AI Assistant Improve Task Performance in Industrial Contexts?” *International Journal of Production Research*.
- MARŠÁL, A. AND P. PERKOWSKI (2025): “A Task-Based Approach to Generative AI: Evidence from a Field Experiment in Central Banking,” Working paper.
- MOLL, B., L. RACHEL, AND P. RESTREPO (2022): “Uneven Growth: Automation’s Impact on Income and Wealth Inequality,” *Econometrica*, 90, 2645–2683.
- MOUSA, D. (2025): “AI Isn’t Replacing Radiologists,” Blog post, *Understanding AI* (Substack).
- NOY, S. AND W. ZHANG (2023): “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science*, 381, 187–192.
- OWEN, G. (1977): “Values of Games with a Priori Unions,” in *Essays in Mathematical Economics and Game Theory*, ed. by R. Heim and O. Moeschlin, New York: Springer, 76–88.
- PARADIS, E., K. GREY, Q. MADISON, D. NAM, A. MACVEAN, V. MEIMAND, N. ZHANG, B. FERRARI-CHURCH, AND S. CHANDRA (2024): “How much does AI impact development speed? An enterprise-based randomized controlled trial,” *arXiv preprint arXiv:2410.12944*.
- PENG, S., E. KALLIAMVAKOU, P. CIHON, AND M. DEMIRER (2024): “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” Working paper.
- POSTEL-VINAY, F. AND J.-M. ROBIN (2002): “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70, 2295–2350.
- REICHARDT, H. (2025): “Scale-Biased Technical Change and Inequality,” Working paper.
- RESTREPO, P. (2024): “Automation: Theory, Evidence, and Outlook,” *Annual Review of Economics*, 16, 1–25.

- REYNAERTS, J., R. VARADHAN, AND J. C. NASH (2012): "Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping," Working paper.
- SCHWARCZ, D., S. MANNING, P. BARRY, D. R. CLEVELAND, J. PRESCOTT, AND B. RICH (2024): "AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice," Working Paper.
- SHAPLEY, L. (1953): "A Value for n-Person Games," in *Contributions to the Theory of Games*, ed. by H. Kuhn and A. Tucker, Princeton University Press, vol. 2, 307–317.
- SHEN, J. H. AND A. TAMKIN (2026): "How AI Impacts Skill Formation," Working paper.
- SMEETS, V., L. TIAN, AND S. TRAIBERMAN (2025): "Field Choice, Skill Specificity, and Labor Market Disruptions," Working paper.
- TRAIBERMAN, S. (2019): "Occupations and Import Competition: Evidence from Denmark," *American Economic Review*, 109, 4260–4301.
- VAITHILINGAM, P., T. ZHANG, AND E. GLASSMAN (2022): "Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models," Working paper.
- VARADHAN, R. AND C. ROLAND (2008): "Simple and Globally Convergent Methods for Accelerating the Convergence of Any EM Algorithm," *Scandinavian Journal of Statistics*, 35, 335–353.
- WALTERS, C. (2024): "Empirical Bayes Methods in Labor Economics," in *Handbook of Labor Economics*, Elsevier, vol. 5, 183–260.
- WEBER, T., M. BRANDMAIER, A. SCHMIDT, AND S. MAYER (2024): "Significant Productivity Gains through Programming with Large Language Models," *Proceedings of the ACM on Human-Computer Interaction*, 8, 256:1–256:29.
- WOESSMANN, L. (2025): "Skills and Earnings: A Multidimensional Perspective on Human Capital," *Annual Review of Economics*, 17.

Appendix

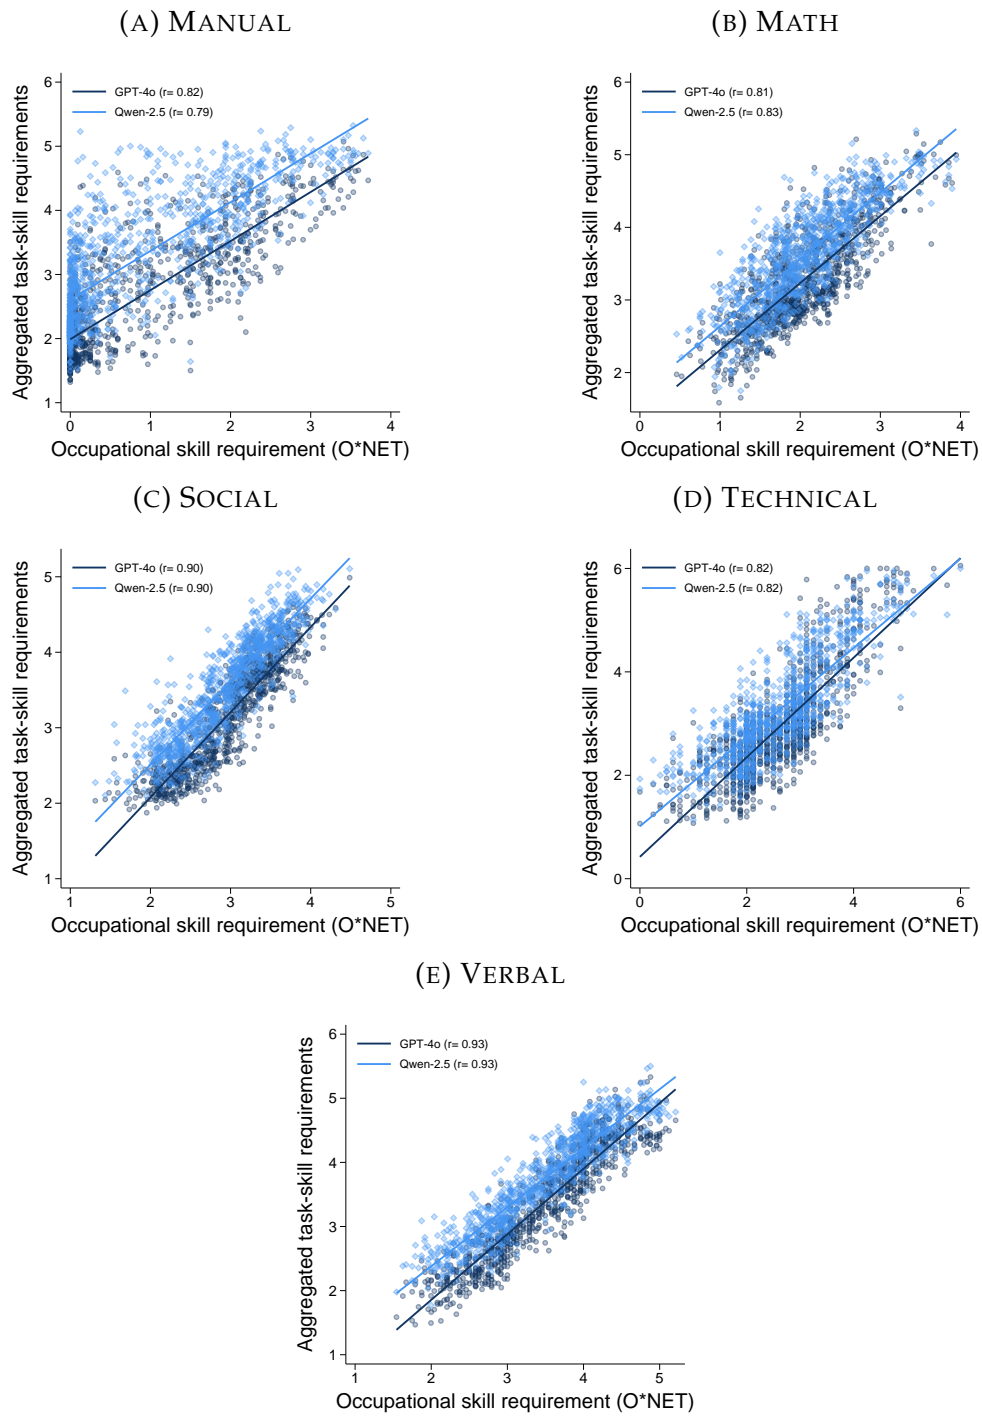
A Figures

FIGURE A.1: PRODUCTION AND SKILL ACCUMULATION: FUNCTIONAL FORMS



Notes: This figure illustrates the functional forms of the production and skill accumulation functions in equations (14) and (16), respectively. Panel A shows the production function $f(h_s, r_{\tau,s}) = h_s^\omega \exp(-\eta \min\{h_s - r_{\tau,s}, 0\}^2)$. Panel B shows the learning part of the skill accumulation function: $\tilde{g}(h_s, r_{\tau,s}, \lambda) = \max\{r_{\tau,s} - h_s, 0\} \exp(-\lambda \max\{r_{\tau,s} - h_s, 0\})$. It illustrates that maximum learning is attained when the skills are $\frac{1}{\lambda}$ below the skill requirements.

FIGURE A.2: VALIDATION OF TASK SKILL REQUIREMENT DATA WITH O*NET



Notes: This figure shows the correlation between the occupation-level skill requirement in the O*NET database versus the GPT-4o and Qwen-2.5 generated task-level skill requirements aggregated to the occupation-level for the skills used in the analysis. Each observation represents an occupation in the O*NET database.

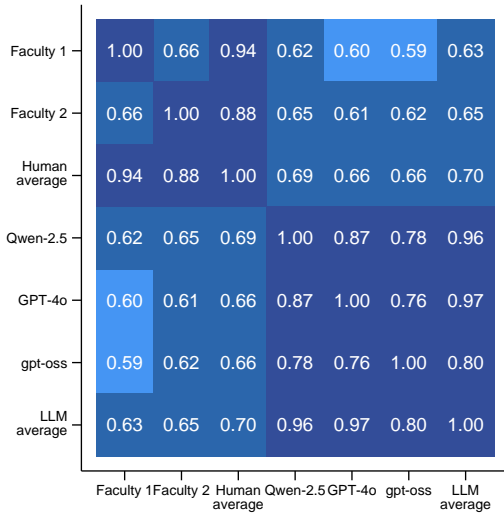
FIGURE A.3: VALIDATION OF TASK SKILL REQUIREMENT DATA WITH O*NET (35 SKILL DIMENSIONS)



Notes: This figure shows the correlations between the occupation-level skill requirement in the O*NET database versus the GPT-4o and Qwen-2.5 generated task-level skill requirements aggregated to the occupation-level for each of the 35 skills.

FIGURE A.4: AGREEMENT ACROSS TASK-LEVEL MEASURES


(A) SKILL REQUIREMENTS WITHOUT AI




(B) SKILL REQUIREMENTS WITH AI



(C) AUGMENTATION

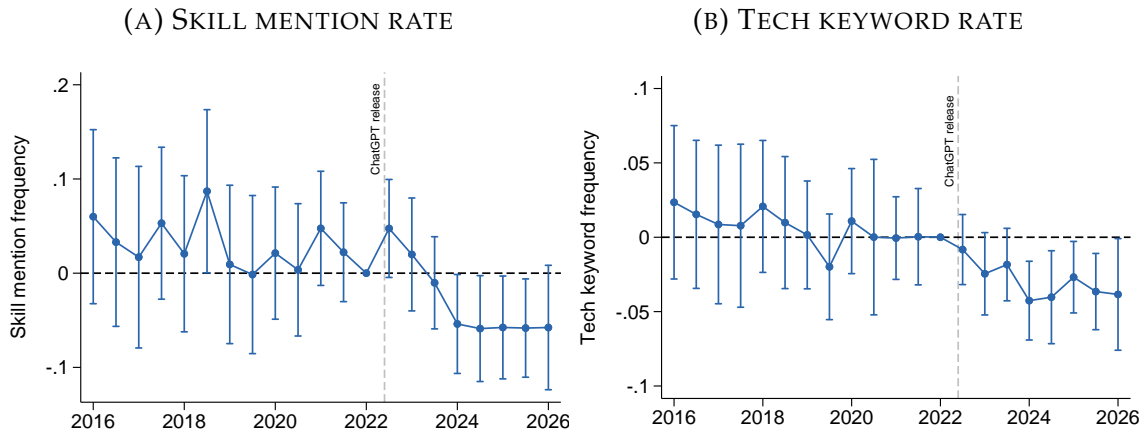


(D) AUTOMATION



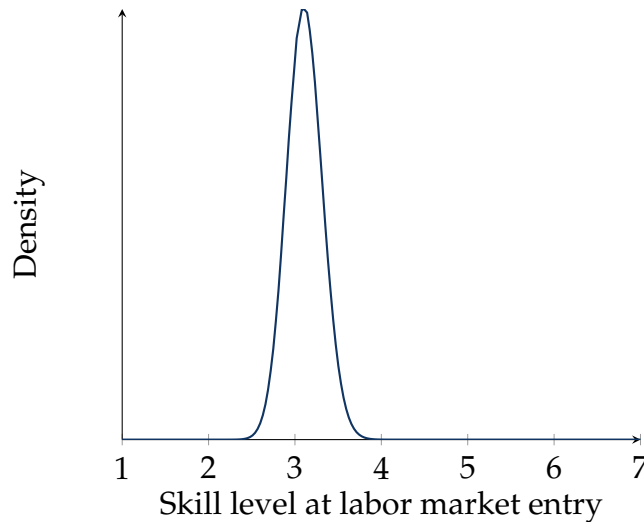
Notes: This figure shows pairwise agreement across task-level measures of skill requirements and AI capabilities. Panels A and B show pairwise correlations of task-level skill requirement ratings across two human raters, the human average, and our LLM-generated ratings for tasks in economics; Panel A shows pre-AI and Panel B post-AI skill requirement levels. All ratings are on the 1–7 scale used by O*NET. Panels C and D show pairwise agreement on task-level AI exposure across measures: Eloundou et al. (2024)’s measure based on GPT-4, our measure based on GPT-4o, and our three measures based on the open-weight Qwen model conditioned on slow, moderate, and rapid AI scenarios (Karger et al., 2026). Panel C reports pairwise correlations between measures’ task-level augmentation scores (the share of time saved to complete a task). Panel D reports the share of tasks on which the row and column measures agree on whether a task is automated, where a task is classified as automated if its exposure is “high” or “full”; Eloundou et al. (2024) provide only GPT-4-rated automation labels, so Panel D includes a single Eloundou measure. Panels C and D are computed on the 14,209 tasks (out of 19,530) present in all measures. The overall share of automated tasks is 21 percent, 22 percent, 14 percent, 22 percent, and 37 percent in Eloundou et al. (2024)’s, our, and Qwen’s slow, moderate, and rapid measures, respectively.

FIGURE A.5: ADDITIONAL MEASURES OF SIMPLIFICATION VIA JOB POSTINGS



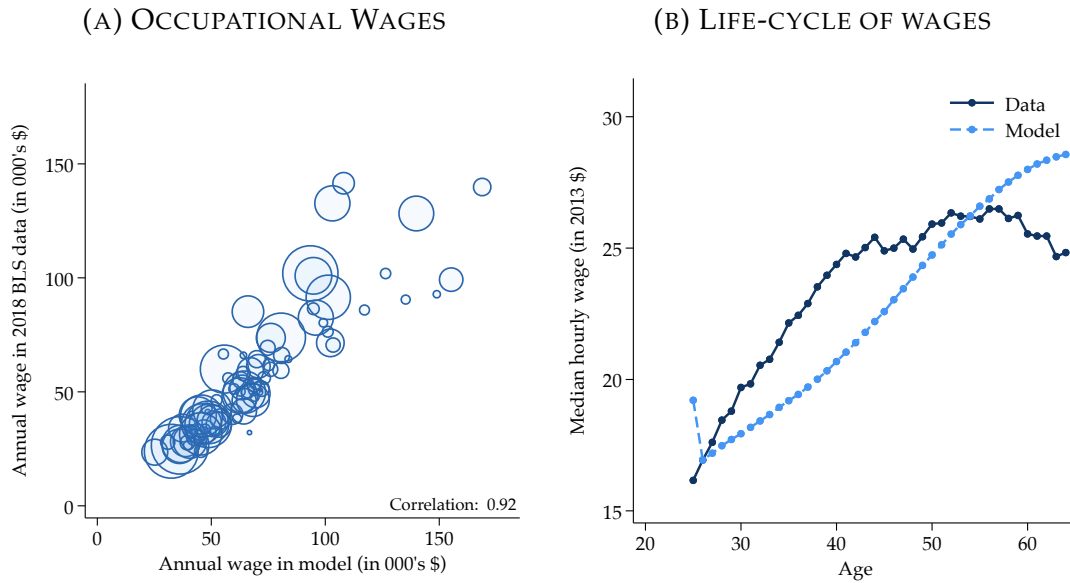
Notes: This figure shows that occupations predicted to be more simplified by AI show changes in the use of skill and technical language in real-time job postings. Continuous treatment exposure is measured using our occupation-level predicted AI-led simplification (slow AI scenario to avoid factoring in post-2026 changes in AI capabilities). Panel A uses the rate of the words “skill” or “skills” per 100 words, and Panel B uses the rate of technical keywords per 100 words (incl. terms like Python, Git, JavaScript, Excel, deep learning, LLM, or API). We use a random subsample of US job postings in Revelio Labs’ Cosmos database. Each coefficient gives the differential effect of a 1 standard deviation higher predicted simplification in t relative to 2022-H1. We include occupation and time fixed effects; standard errors are clustered at the occupation level.

FIGURE A.6: INITIAL SKILL DISTRIBUTION




Notes: This figure shows the estimated density of the skill distribution of young workers (age $a = 1$ in the model) on O*NET’s 1 to 7 scale. For comparison, for the skill “reading comprehension”, a 2 means being able to “read step-by-step instructions for completing a form”, 4 means being able to “understand an email from management describing new personnel policies”, and 6 means being able to “read a scientific journal article describing surgical procedures”.

FIGURE A.7: MODEL FIT: COMPARING MODEL MOMENTS WITH DATA




Notes: Panel A shows the correlation between the average wage in an occupation in the model's steady state and in the data as reported in the 2018 BLS OEWS data. Panel B reports the median wage by age in the NLSY79 and the model's steady state.

FIGURE A.8: WAGE EFFECTS: STANDARD VS. NESTED CES



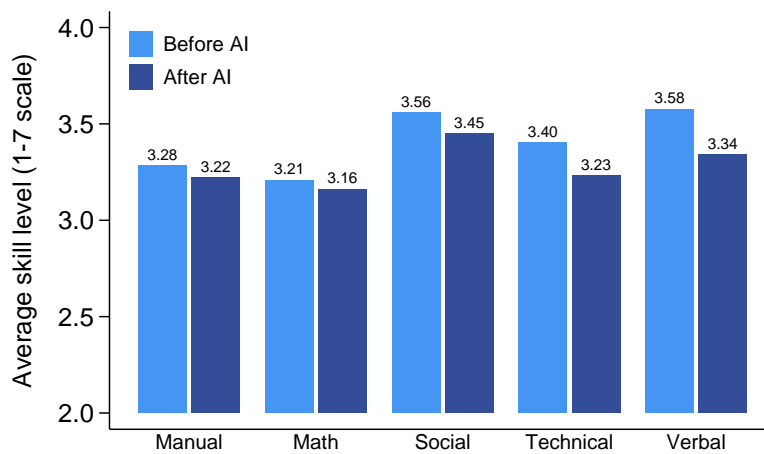
Notes: This figure shows the distribution of wage changes induced by generative AI across the wage percentile distribution under two alternative demand specifications (moderate AI scenario). The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile. The black line shows results under the baseline one-level CES demand structure, while the gray line shows results under a nested CES structure with three layers: broad industry groups, detailed industries, and occupations. Both specifications include the joint effect of AI's augmentation, automation, and simplification.

FIGURE A.9: WAGE EFFECTS BY CHANNEL



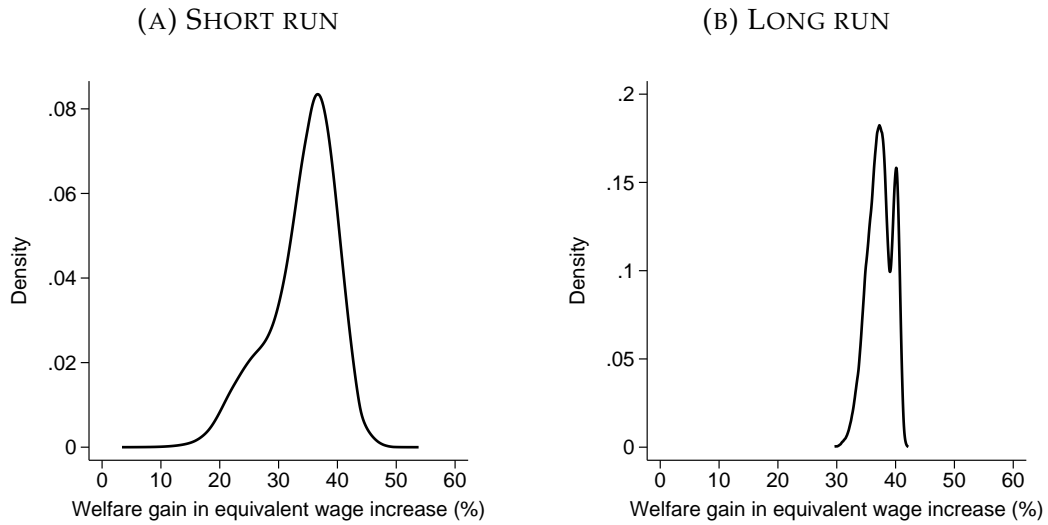
Notes: This figure shows the distribution of wage changes induced by generative AI across the wage percentile distribution (moderate AI scenario). The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile. The black line shows the joint effect of AI's augmentation, automation, and simplification on each wage percentile. Other lines show the effects when each channel operates alone.

FIGURE A.10: HOW AI CHANGES AVERAGE SKILL LEVELS



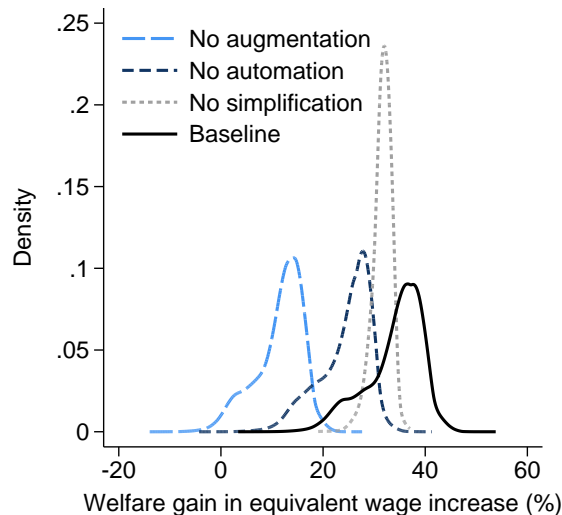
Notes: This figure shows the average skill level before and after the introduction of AI for each of the five skill dimensions (moderate AI scenario). Averages are computed across individuals in the steady state of our model prediction before versus after AI. For interpretability, skill levels are expressed on the O*NET scale from 1 to 7.

FIGURE A.11: DISTRIBUTION OF WELFARE EFFECTS



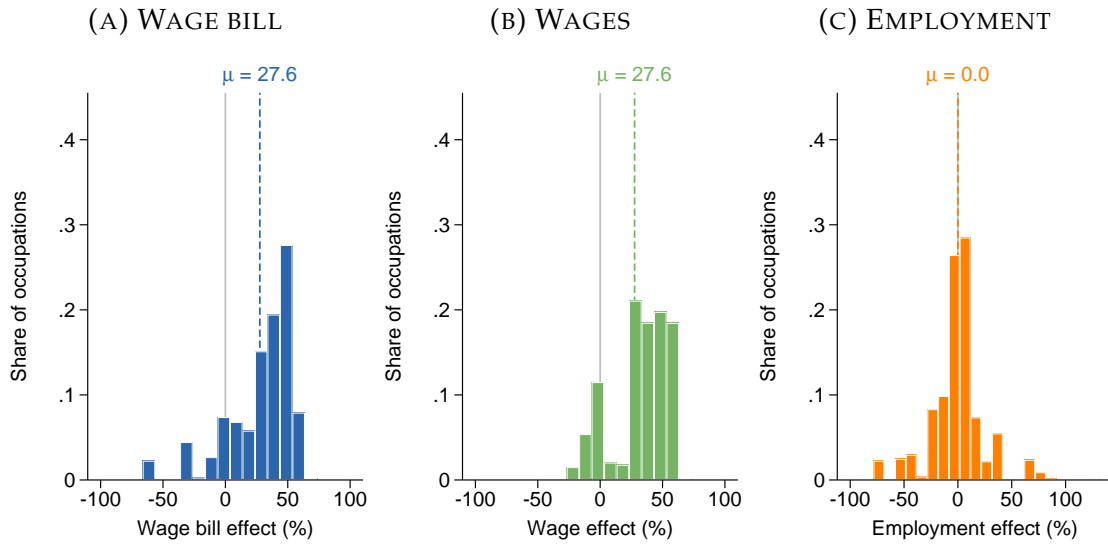
Notes: This figure shows the distribution of AI's welfare effect on individual workers (moderate AI scenario), in the short run along the transition path, at the time of the shock (Panel A) and in the long run across steady states (Panel B). The welfare effect is measured in equivalent wage variation: the permanent proportional wage increase across all occupations that yields the same welfare gain as the introduction of AI. Unlike the long-run comparison, the short-run measure accounts for the dynamic adjustment path of the economy.

FIGURE A.12: DISTRIBUTION OF SHORT-RUN WELFARE EFFECTS UNDER DIFFERENT SCENARIOS



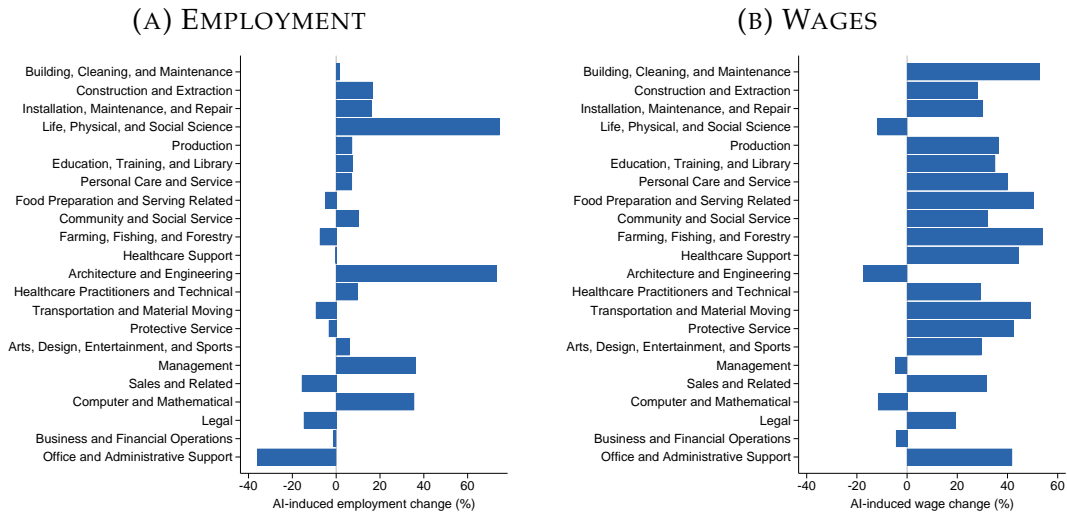
Notes: This figure shows the distribution of AI's welfare effect on individual workers along the transition path in the moderate AI scenario. The dark line represents the distribution under the baseline scenarios. The other lines show the distribution when one of the channels is not present. The welfare effect is measured in equivalent wage variation: it represents the permanent wage increase across all occupations that yields the same welfare gain as the introduction of AI. Unlike the steady-state comparison, this measure accounts for the dynamic adjustment path of the economy.

FIGURE A.13: GENERATIVE AI'S EFFECT ACROSS OCCUPATIONS




Notes: This figure shows the distribution of generative AI's predicted effects across occupations based on our structural model (moderate AI scenario). Panel (A) shows wage bill changes (wages \times employment). Panel (B) shows wage changes. Panel (C) shows employment effects, which average to zero by definition as our model does not feature unemployment.

FIGURE A.14: AI'S EFFECT ON OCCUPATIONAL EMPLOYMENT & WAGES




Notes: This figure shows the model's predictions on AI's employment and wage effects by occupational group (moderate AI scenario). Occupations are sorted in descending order of AI's effect on their wage bill, so that the first listed occupation experiences the largest wage bill increase.

FIGURE A.15: AI'S EFFECT ON DETAILED OCCUPATIONS' WAGE BILLS



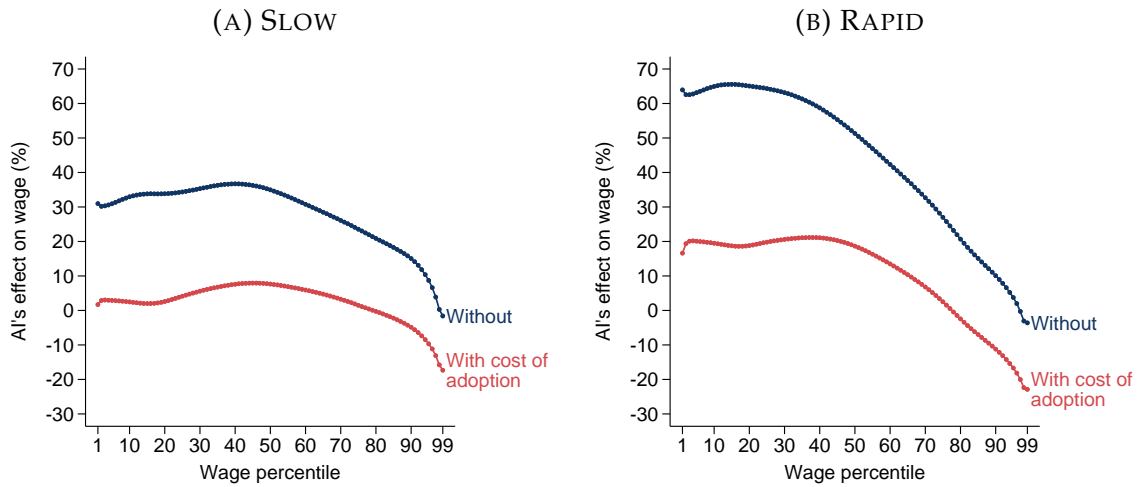
Notes: This figure shows the model predictions on AI's wage bill effects by occupation.

FIGURE A.16: AI'S EFFECT ON DETAILED OCCUPATIONS' EMPLOYMENT & WAGES



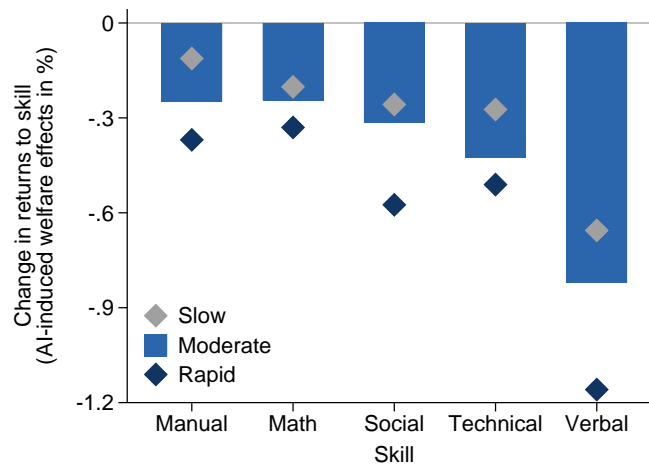
Notes: This figure shows the model's predictions on AI's employment and wage effects by occupation (moderate AI scenario). Occupations are sorted in descending order of AI's effect on the total wage bill, so that the first listed occupation experiences the largest wage bill increase. We conduct a Shapley-Owen decomposition to separate the overall change into the contribution of each channel: augmentation, automation, and simplification.

FIGURE A.17: ADOPTION COSTS AND AI'S WAGE EFFECTS: SLOW AND RAPID SCENARIOS



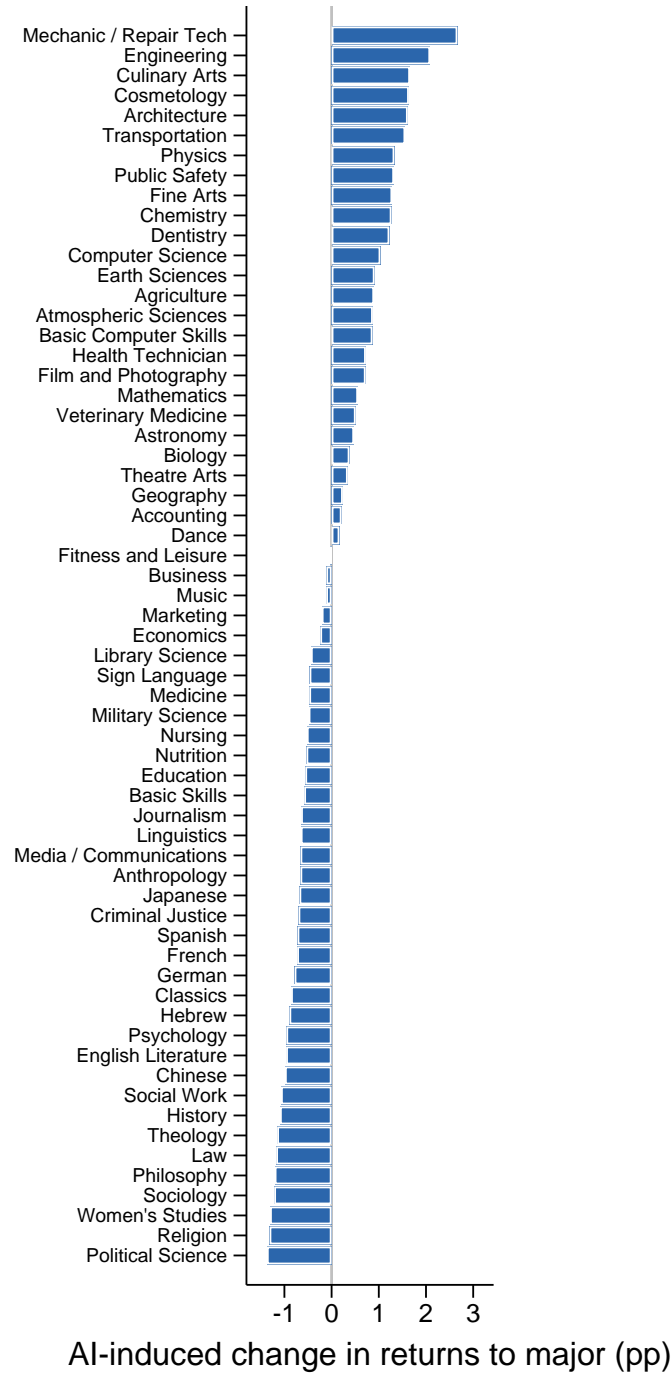
Notes: This figure shows the distribution of wage changes induced by generative AI across the wage percentile distribution with versus without adoption costs for the slow (Panel A) and rapid (Panel B) AI scenarios. Estimates without adoption costs replicate the “slow” and “rapid” lines in Panel A of Figure 3. The moderate scenario is shown in Figure 8. Red lines show estimates with adoption costs; blue lines show estimates without adoption costs. We follow Karger et al. (2026) in defining the scenarios and measure task-specific AI capabilities using Qwen 2.5. The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile.

FIGURE A.18: HOW AI CHANGES THE RETURNS TO SKILLS BY SCENARIO



Notes: This figure shows how AI's welfare effects differ by skills across three AI scenarios: slow, moderate, and rapid. The welfare effects are measured in equivalent permanent percentage wage increases. This figure plots the coefficient of a regression of these welfare effects on skill levels across all dimensions. For interpretability, the skills are expressed on the O*NET scale from 1 to 7.

FIGURE A.19: AI'S EFFECT ON RETURNS TO COLLEGE MAJORS



Notes: This figure shows the model's predictions on AI's effect on returns to college majors (moderate AI scenario). Returns are calculated by combining major-level skill intensities from the Skill Atlas with the model's predicted changes in returns to each skill dimension. Values are centered relative to the average major.

B Tables

TABLE B.1: AGGREGATION OF O*NET'S SKILL REQUIREMENTS TO 5 DIMENSIONS

Skill	O*NET skill	O*NET skill category
Manual	Equipment Maintenance	Technical
	Equipment Selection	Technical
	Installation	Technical
	Repairing	Technical
Math	Mathematics	Basic Content
Social	Active Listening	Basic Content
	Coordination	Social
	Instructing	Social
	Management of Personnel Resources	Resource Management
	Negotiation	Social
	Persuasion	Social
	Service Orientation	Social
Social Perceptiveness	Social	
Technical	Complex Problem Solving	Complex Problem Solving
	Judgment and Decision Making	Systems
	Operation and Control	Technical
	Operations Analysis	Technical
	Operations Monitoring	Technical
	Programming	Technical
	Quality Control Analysis	Technical
	Science	Content
	Systems Analysis	Systems
	Systems Evaluation	Systems
	Technology Design	Technical
Troubleshooting	Technical	
Verbal	Reading Comprehension	Basic Content
	Speaking	Basic Content
	Writing	Basic Content

Notes: This table shows the mapping of the five skill clusters—Manual, Math, Social, Verbal, and Technical—to the relevant O*NET skills and their respective O*NET's skill category. For each of the skills, we set the requirement to the average across the relevant O*NET skills. We dropped 7 out of 35 O*NET skill dimensions that could not be clearly mapped into the skills used in the analysis.

TABLE B.2: SUMMARY OF TASK-LEVEL DATA ON AI CAPABILITIES (QWEN)

	Augmentation		Automation	Simplification
	Excluding automatable tasks	Including automatable tasks		
Mean	23.1%	25.2%	22.4%	23.9%
Std. Dev.	7.8%	8.8%	41.7%	6.9%
Median	20.0%	30.0%	0.0%	24.8%
Range	0.0% - 75.0%	0.0% - 75.0%	0.0% - 100.0%	0.0% - 42.9%
Tasks	15,161	19,530	19,530	19,530

Notes: This table summarizes our new estimates of generative AI’s potential impact on tasks across three channels: augmentation (share of worker’s time saved by technology to complete the task), automation (share of tasks that can be fully automated by technology), and simplification (relative decrease in average skill requirements across all 35 O*NET skill dimensions). For augmentation, we present estimates both excluding and including tasks that can be automated. We measure task-specific AI capabilities using Qwen 2.5 under the moderate AI scenario, following [Karger et al. \(2026\)](#) in defining the scenarios.

TABLE B.3: EXPERIMENTAL ESTIMATES COMPARED TO OUR TASK AUGMENTATION DATA

Occupation	Activity	Tool	Estimate		N	Notes	Source
			RCT	LLM			
Software developer	Coding	GitHub Copilot	26%	30%	4,867	+26.1% number of completed tasks in lab; new developers higher adoption rates & higher productivity gains	[1]
Software developer	Coding	GitHub Copilot	56%	30%	95	55.8% time saved, quality ↑	[2]
Software developer	Coding	GitHub Copilot	36%	30%	23	36% time saved for familiar tasks; no change for unfamiliar tasks; 48% fewer issues	[3]
Programmer	Coding	GPT-3	27%	30%	100	27% time saved among 100 expert programmers; 50 non-programmers perform tasks similarly well with LLM	[4]
Programmer	Coding	GitHub Copilot	0%	30%	24	No time saved; however, most participants still preferred using LLM	[5]
Management consultant	Consulting	GPT-4	25%	30%	758	25.1% time saved, +12.2% tasks completed, +40% quality (decreased for tasks beyond AI frontier); lower-skilled consultants benefited more	[6]
Customer support	Resolution	GPT-4	14%	30%	5,179	+14% productivity (issues resolved per hour), +34% for low-skill workers; minimal for high-skill workers	[7]
–	Writing	GPT-3.5	40%	30%	453	40% time saved, +18% output quality; inequality between workers ↓; low-skill workers benefited most	[8]
Taxi driver	Selecting routes	AI Navi	14%	10%	520	Shorter cruising time; gains only among low-skill drivers	[9]
Lawyer	Legal writing	Vincent, o1-prev.	20%	30%	127	19.9% time saved across different legal writing tasks, quality ↑, LLM “Vincent” slightly higher gains	[10]
Product designer	Product marketing, dev.	GPT-4o	13-16%	30%	776	+0.37 SD quality and 16.4% time saved for individuals; +0.39 SD quality and 12.7% time saved for teams	[11]
Software developer	Coding	GitHub Copilot	65%	30%	24	Developers implemented ~65% more requirements with AI assistance	[12]
Software developer	Coding	Google AI	21%	30%	96	AI users finished an enterprise-grade task 21% faster. Results stronger for senior developers.	[13]
Programmer	Coding	CodeFuse	55%	30%	1,219	Lines of code produced ↑ 55%, gains concentrated among junior staff	[14]
Knowledge workers	E-mail	MS Copilot	365 11%	30%	7,137	Treated spent 12% less time on email each week; did not significantly change time spent in meetings.	[15]
Train commissioning technician	Troubleshooting	GPT-3.5 + RAG	20%	20%	173	+1.14 SD quality score; 20% increase in tasks completed not significant; less-experienced benefit more.	[16]
Knowledge workers	E-mail	GPT-4.1 assistant	9-15%	30%	1,174	Duration falls by 0.96 min (low edu) and 1.51 min (high edu) on ~10.4–10.7 min baseline (~9–15%); closes ~75% of edu performance gap.	[17]
Software developer	Coding	Frontier AI tools	-19%	30%	16	Experienced developers in 246 real repository tasks	[18]
Knowledge worker	Office tasks	GPT-4o	21%	40%	101	Randomized GPT-4o access reduces time by 21%; quality rose 33-44%	[19]
Security professional	Reporting	Security Copilot	23%	30%	147	7% higher accuracy; reports have 49% more key facts	[20]

Notes: Sources correspond to [1] Cui et al. (2024), [2] Peng et al. (2024), [3] Clarke and Hanrahan (2024), [4] Campero et al. (2022), [5] Vaithilingam et al. (2022), [6] Dell’Acqua et al. (2023), [7] Brynjolfsson et al. (2025), [8] Noy and Zhang (2023), [9] Kanazawa et al. (2022), [10] Schwarcz et al. (2024), [11] Dell’Acqua et al. (2025), [12] Weber et al. (2024), [13] Paradis et al. (2024), [14] Gambacorta et al. (2024), [15] Dillon et al. (2025), [16] Lowhagen et al. (2025), [17] Cruces et al. (2026). [18] Becker et al. (2025), [19] Maršál and Perkowski (2025), [20] Edelman et al. (2024). To construct our own estimates of task augmentation by generative AI at the level of work activities, we aggregate our task-level estimates within the relevant occupation as equally weighted averages for tasks we judge to be relevant to the work activity covered in each experiment.

TABLE B.4: OVERVIEW OF MODEL PARAMETERS

Model object	Symbol	Value	How it is set
Elast. of substitution: Occupations	σ	1.57	Burstein et al. (2019); Caunedo et al. (2023)
Elast. of substitution: Tasks	ρ	0.49	Humlum (2021)
Number of occupations	J	93	3-digit BLS SOC occupations
Number of periods	A	40	Years between 25 and 64
Discount factor	β	0.78	Following Keane and Wolpin (1997)
Skill dimensions	S		Addison et al. (2020); Baley et al. (2022), plus manual
Occupational task sets	\mathcal{T}_j		O*NET tasks
Occupational task weights	$\theta_{j,\tau}$		O*NET task importance
Task-level skill requirements	r_τ		Large language model
Task-level AI augmentation	γ_τ		"
Task-level AI automation	\mathcal{A}_j		"
Learning cost: 1 st AFQT quartile	$\lambda(1)$	3.50	Maximum likelihood
Learning cost: 2 nd AFQT quartile	$\lambda(2)$	2.97	"
Learning cost: 3 rd AFQT quartile	$\lambda(3)$	2.81	"
Learning cost: 4 th AFQT quartile	$\lambda(4)$	2.67	"
Human capital depreciation	δ	0.0003	"
Scale of productivity shocks	ζ	0.053	"
Occupational switching cost	κ	0.340	"
Skill distribution (Beta)	(B_a, B_b)	(62,114)	"
Cost of underqualification	η	0.04	OLS within MLE routine
Skill productivity: Manual	ω_{Mn}	0.33	"
Math	ω_{Mt}	0.79	"
Social	ω_S	0.55	"
Technical	ω_T	0.22	"
Verbal	ω_V	0.36	"
Occupational amenities	$\{\mu_j\}_{j=1}^J$		Match employment shares à la (Berry et al., 1995)
Occupational demand	$\{\alpha_j\}_{j=1}^J$		Using estimated prices and wage bills

Notes: This table provides an overview of the parameters of the model, their mathematical symbols, the value at which they are set, and the procedure with which we arrived at the value.

TABLE B.5: PRODUCTION FUNCTION: ROBUSTNESS OF PARAMETER ESTIMATES

	General skill					Mismatch
	ω_{Mn}	ω_{Mt}	ω_S	ω_T	ω_V	η
Baseline	0.330	0.786	0.545	0.220	0.361	0.044
NLS	0.330	0.785	0.547	0.219	0.362	0.043
Spline	0.304	0.816	0.533	0.223	0.328	0.042

Notes: This table shows robustness of the estimates of the task-level productivity parameters to changes in the estimation routine. Subscripts Mn , Mt , S , T , V refer to manual, math, social, technical, and verbal, respectively. The first row shows the baseline estimates in Table 1. The second row shows estimates resulting from non-linear least squares on the exact wage function in equation (15), instead of from OLS on the linear approximation as in the baseline. The third row shows the estimates resulting from controlling for the selection term more flexibly using a cubic spline.

TABLE B.6: OCCUPATIONAL SKILL REQUIREMENTS PREDICT OCCUPATIONAL PRICES

Skill requirements	Dependent variable: Log occupational price \hat{p}_j			
Manual	0.228 (0.207)	0.230 (0.214)	0.226 (0.201)	0.227 (0.208)
Math	-0.070 (0.277)	-0.073 (0.286)	-0.077 (0.269)	-0.080 (0.278)
Social	-0.083 (0.352)	-0.078 (0.363)	-0.077 (0.341)	-0.073 (0.353)
Technical	1.342** (0.599)	1.341** (0.619)	1.353** (0.581)	1.352** (0.601)
Verbal	0.554 (0.376)	0.554 (0.388)	0.550 (0.364)	0.550 (0.377)
Sample occupations	All	50+	All	50+
Empirical Bayes applied	No	No	Yes	Yes
Observations	93	87	93	87
R^2	0.73	0.73	0.74	0.74

Notes: This table shows the coefficients and R^2 of a regression of the estimated occupational prices (in logs) on occupational skill requirements. Each observation represents one occupation and is weighted by the number of worker-year observations in that occupation. Columns where sample occupations indicates “50+” only include occupations with at least 50 observations. The third and fourth column show the results with fixed effects on which empirical Bayes regression has been applied. Occupational prices are estimated as the occupational fixed effects in regression equation (17). Occupational skill requirements refer to $\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} r_{\tau,s}$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.7: MODEL FIT: WAGE INEQUALITY IN DATA AND MODEL

	Gini	Ratios			Top shares		
		$\frac{p_{90}}{p_{10}}$	$\frac{p_{90}}{p_{50}}$	$\frac{p_{75}}{p_{25}}$	10%	5%	1%
Data	0.32	4.01	2.13	2.04	0.26	0.16	0.05
Model	0.24	2.96	1.83	1.84	0.20	0.11	0.02

Notes: This table reports measures of inequality in the unconditional wage distribution in the NLSY79 data and in the model’s steady state. The unit of observation is a worker-age pair in the data and in the model. We only included workers who remain in the NLSY79 and work until 2020 without interruptions over 18 months. Sample weights are applied in the NLSY79 data.

C Solution Algorithm

Below we describe in detail the algorithms used to solve for a stationary competitive equilibrium and the transition path after an unexpected one-off technological shock.

Stationary equilibrium. To solve for a stationary equilibrium, we use the following algorithm:

1. Guess an initial vector of relative prices $(p_1^{(1)}, \dots, p_J^{(1)})$.
2. For iteration r , given the prices $(p_1^{(r)}, \dots, p_J^{(r)})$, solve the worker's problem and compute the implied output of each good $(\mathcal{Y}_1^{(r)}, \dots, \mathcal{Y}_J^{(r)})$. Then, update prices to clear the market given supply: $\mathbf{p}^{(r+1)} = D^{-1}(\{\mathcal{Y}_j^{(r)}\}_{j=1}^J)$.³¹
3. Repeat step 2 until $\|\mathbf{p}^{(r+1)} - \mathbf{p}^{(r)}\| < \epsilon$ for a threshold $\epsilon > 0$.

Transition path. Starting from the initial stationary equilibrium, we solve for the transition path of prices $\{p_{1,t}, \dots, p_{J,t}\}_{t=1}^{\infty}$ from the moment the shock is realized. We numerically approximate this infinite sequence by solving for $\{p_{1,t}, \dots, p_{J,t}\}_{t=1}^T$ for a large enough T such that prices are constant after period T . The solution algorithm is based on [Boppart et al. \(2018\)](#):

1. Compute the stationary equilibrium before ($t = 0$) and after the change ($t = T$).
2. Guess a path for the sequence of prices.³²
3. For iteration r , given the sequence of prices $\{p_{1,t}^{(r)}, \dots, p_{J,t}^{(r)}\}_{t=1}^T$, solve for the value function at $t = T, T - 1, \dots, 1$. Then, compute the implied output of each good at each time and the corresponding prices $\mathbf{p}_t^{(r+1)} = D^{-1}(\{\mathcal{Y}_{j,t}^{(r)}\}_{j=1}^J)$ for $t = 1, \dots, T$.
4. Repeat step 3 until $\|\mathbf{p}_t^{(r+1)} - \mathbf{p}_t^{(r)}\| < \epsilon \forall t = 1, \dots, T$ for a threshold $\epsilon > 0$.

³¹In case of simple CES demand for occupations as in our baseline, there is a simple closed-form solution for the updated price. We also consider a nested CES structure, in which case the price update can be computed using a Newton tâtonnement algorithm.

³²A reasonable guess is the path where prices adjust immediately to the new stationary equilibrium.

Computing implied output given the sequence of prices is the main computational challenge in these algorithms. It consists of three main steps for which we provide more detail below.

First, we solve for the value function. While conceptually straightforward, the large state space (\mathbf{h}, ψ, j) makes brute force value function iteration costly. We exploit a convenient feature of the value function in equation (11) to provide relief. Since the occupational switching cost is $\kappa(j, k) = \kappa 1[j \neq k]$, the value function can be written as

$$V_a(\mathbf{h}, \psi, k) = \zeta \log \left[e^{-\frac{\kappa}{\zeta}} \sum_{j=1}^J \tilde{V}_a^j(\mathbf{h}, \psi) + \left(1 - e^{-\frac{\kappa}{\zeta}}\right) \tilde{V}_a^k(\mathbf{h}, \psi) \right]$$

where $\tilde{V}_a^j(\mathbf{h}, \psi) \equiv \exp\left(\frac{1}{\zeta} (\log w_j(\mathbf{h}) + \mu_j + \beta V_{a+1}(g_j(\mathbf{h}, \psi), \psi, j))\right)$. This solution implies that it is sufficient to solve for $\tilde{V}_a^j(\mathbf{h}, \psi)$, shrinking the state space in the value function iteration step by a factor equal to the number of occupations J (in our exercise, $J = 93$).³³

Second, after computing the value function and conditional choice probabilities, we compute the joint distribution of state variables using the law of motion in equation (12). We do so by simulation. We first draw from the initial distribution of skills and learning ability at age $a = 1$. For each of these individuals, we then draw an occupation based on the choice probability conditional on their initial states and update state variables accordingly. We iterate this process forward until $a = A$.³⁴

Third, to update the relative prices, we compute total implied production of each good given the previous price iteration. The previous steps yield a sample of workers with a given occupation and set of skills. From there, we approximate the integral in equation (13) for each occupation j . That is, we evaluate the term $Y_j(\mathbf{h}) \mathbb{P}_{a,t}(j | \mathbf{h}, \psi, k) \mathbb{E}[\varepsilon_j | j, \mathbf{h}, \psi, k]$ for each worker-age and for each $j = 1, \dots, J$. We do not condition on the occupational draw in the computation of production so that sampling noise only affects workers' states, not production conditional on those states.

³³Conditional choice probabilities can then be recovered as

$$\mathbb{P}_a(j | \mathbf{h}, \psi, k) = \frac{\tilde{V}_a^j(\mathbf{h}, \psi) e^{-\frac{\kappa}{\zeta} 1[j \neq k]}}{e^{-\frac{\kappa}{\zeta}} \sum_{j=1}^J \tilde{V}_a^j(\mathbf{h}, \psi) + \left(1 - e^{-\frac{\kappa}{\zeta}}\right) \tilde{V}_a^k(\mathbf{h}, \psi)}.$$

³⁴To save computational costs at early price iterations, we begin with a small number of simulations, and increase sample sizes as the difference between subsequent price iterations decreases.

D Estimation

D.1 Cost savings and AI's income share

The cost of performing a task τ with the worker's unit of time equals

$$c_\tau^l(\mathbf{h}) \equiv \frac{\Lambda_j(\mathbf{h})}{\gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau)}$$

where $\Lambda_j(\mathbf{h})$ is the shadow value of a unit of time in occupation j given skills \mathbf{h} . Similarly, let c_τ^k be the unit cost of producing task τ with capital, i.e.,

$$c_\tau^k \equiv \frac{R}{\phi_\tau}.$$

For any automated task $\tau \in \mathcal{A}_j$, the cost of producing the task with capital relative to performing the task by labor thus equals

$$\chi_\tau \equiv \frac{c_\tau^k}{c_\tau^l(\mathbf{h})} = c_\tau^k \cdot \frac{\gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau)}{\Lambda_j(\mathbf{h})}.$$

Since the shadow value of a unit of time is the wage, i.e., $\Lambda_j(\mathbf{h}) = w_j(\mathbf{h})$, equation (8) implies that cost savings are equal to

$$\chi_\tau = \frac{c_\tau^k}{p_j} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \left(\frac{c_\tau^k}{p_j} \right)^{1-\rho} \right)^{\frac{1}{\rho-1}} \frac{\gamma_\tau f(\mathbf{h}, \mathbf{r}_\tau)}{\left(\sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} f(\mathbf{h}, \mathbf{r}_\tau)^{\rho-1} \right)^{\frac{1}{\rho-1}}}.$$

For our quantification, we need an estimate of $\sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \left(\frac{c_\tau^k}{p_j} \right)^{1-\rho}$ for each occupation. To obtain this, we make two simplifying assumptions. First, we assume that the cost savings do not vary across automatable tasks, i.e., $\chi_\tau = \chi$ for all $\tau \in \mathcal{A}_j$ and $\forall j = 1, \dots, J$. Second, we assume that the automated tasks are not different in productivity and skill requirements from the non-automatable tasks, i.e., $\gamma_\tau^{\rho-1} f(\mathbf{h}, \mathbf{r}_\tau)^{\rho-1} \approx \frac{1}{\sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau}} \sum_{\tau \in \mathcal{N}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} f(\mathbf{h}, \mathbf{r}_\tau)^{\rho-1}$ for all $\tau \in \mathcal{A}_j$ and $\forall j = 1, \dots, J$. Under those two assumptions, the cost savings simplify to

$$\chi = \frac{c_\tau^k}{p_j} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \left(\frac{c_\tau^k}{p_j} \right)^{1-\rho} \right)^{\frac{1}{\rho-1}} \frac{1}{\left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \right)^{\frac{1}{\rho-1}}} \quad \forall j = 1, \dots, J, \forall \tau \in \mathcal{A}_j$$

so that

$$c_\tau^k / p_j = \left(\chi^{\rho-1} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \right) + \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \right)^{\frac{1}{\rho-1}},$$

from which equation (21) follows.

D.2 Log-linearization of the production function

In this paragraph, we derive the log-linear wage regression equation in (17). Starting from the wage equation (15), and imposing $\mathcal{A}_j = \emptyset$ (so that $\Gamma_j = 1$ and $\mathcal{N}_j = \mathcal{T}_j$) for all $j = 1, \dots, J$, we obtain

$$\begin{aligned} \log w_j(\mathbf{h}) &= \log p_j + \sum_{s \in \mathcal{S}} \omega_s \log(h_s) \\ &+ \frac{1}{\rho-1} \log \left(\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} \exp \left(-\eta \sum_{s \in \mathcal{S}} \min \{h_s - r_{\tau,s}, 0\}^2 \right)^{\rho-1} \right). \end{aligned}$$

Now define the variable $m_\tau \equiv \sum_{s \in \mathcal{S}} \min \{h_s - r_{\tau,s}, 0\}^2$ and log-linearize the wage function around $m_\tau = 0$ for all $\tau \in \mathcal{T}_j$. That is, we linearize the wage function around the perfectly matched worker.

A first-order Taylor expansion around this point yields

$$\begin{aligned} &\log \left(\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} \exp \left((1-\rho) \eta m_\tau \right) \right) \\ &\approx \log \left(\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} \right) + \eta(1-\rho) \frac{\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} m_\tau}{\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1}} \\ &= (1-\rho) \left(\eta \sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} m_\tau \right) \end{aligned}$$

where the second equality follows from $\sum_{\tau \in \mathcal{T}_j} \theta_{j,\tau} \gamma_\tau^{\rho-1} = 1$. Combining the equations above with the definition of m_τ yields equation (17).

E Data

E.1 AI Capabilities: Augmentation, Automation, Simplification

We model AI’s task-level impact on workers through three distinct channels: augmentation, automation, and simplification. We leverage O*NET’s assessment framework and descriptions of occupations, tasks, and skills to generate new data using a variety of LLMs, including two open-weight models (Alibaba’s Qwen 2.5-72B-Instruct, a 72.7-billion-parameter dense model; and OpenAI’s gpt-oss-120b, a mixture-of-experts model with 117 billion total and 5.1 billion active parameters) and a closed-weight model (OpenAI’s GPT-4o, whose parameter counts are not publicly disclosed).

We treat our main estimates as those from Qwen 2.5-72B-Instruct under the *moderate* AI-progress scenario; the slow and rapid scenarios, together with the gpt-oss-120b and GPT-4o estimates, serve as robustness checks. Our prompt structure is consistent across AI scenarios and LLM models, varying only the AI scenario. All prompts can be found in our GitHub repository: https://github.com/lukasalthoff/ai_labor_markets.

TABLE E.8: AI-PROGRESS SCENARIOS

Slow	Moderate	Rapid
AI is a capable assisting technology for humans: writing literature reviews at the level of a capable PhD student, handling half of all freelance software-engineering jobs that would take an experienced human a day to complete, topping up online grocery carts, and physically being able to unload dishwashers in some homes.	AI is an effective collaborator across domains: autonomous lab systems can make rapid advances in solar-cell technology; almost all freelance software-engineering jobs requiring 5 days of effort from an experienced human are automatable; robots can do dishes as quickly as humans; robo-taxis can drive anywhere that humans can.	AI systems surpass humans in most cognitive and physical tasks. Autonomous researchers can collapse years-long research timelines into months or even days. AI systems surpass all freelance software engineers, customer service agents, paralegals, and clerical workers. Models can write 2025-Pulitzer-caliber books—and negotiate the resulting book contract. Robots can assist in an arbitrary home or factory anywhere in the world.

Notes: This table shows the three scenarios borrowed from [Karger et al. \(2026\)](#) to anchor the LLM’s prediction on AI’s task-specific capabilities. In addition to these three scenarios, we also generate estimates using a no-scenario condition (i.e., the prompts include no explicit scenario statement, which we use as an additional robustness check).

We consider multiple scenarios of AI capabilities: slow, moderate, and rapid AI progress as defined in [Karger et al. \(2026\)](#)—see Table E.8. We additionally provide all of our prompts without outlining a specific scenario as a further robustness check.

E.1.1 Augmentation of Tasks

For augmentation assessment, we ask the language model to estimate time savings when workers get access to generative AI. We evaluate one prompt per O*NET task, for a total of 19,530 independent prompts.

E.1.2 Automation of Tasks

To measure AI’s potential to automate occupational tasks, we follow [Eloundou et al. \(2024\)](#) in using a five-tier rubric ranging from no automation (T0) to full automation (T4) exposure. We evaluate one prompt per O*NET task, for a total of 19,530 independent prompts. The automation prompt follows [Eloundou et al. \(2024\)](#)’s format, with generative-AI-specific definitions and examples.

E.1.3 Simplification of Tasks

The simplification channel assesses how AI changes the skill requirements for performing tasks. We evaluate skill levels both without and with AI access, allowing us to measure the change in required skills.

We elicit a task’s skill requirements by replicating O*NET’s occupation-level questionnaire on the task-level. We request the skill requirement for each of the 19,530 tasks for each of the 35 skill dimensions, resulting in 683,550 independent prompts. As in O*NET, the skill requirements are rated from 1 to 7 and each of the 35 skills have different “level anchors” to indicate the meaning of levels 2, 4, and 6. These anchors, as well as the set of tasks in each occupation, and their descriptions, are taken from the O*NET database.

The prompt asks for both values simultaneously. This approach allows us to measure both the baseline skill requirements r_τ and the AI-adjusted requirements r'_τ in a single API call for each skill and task, improving consistency and reducing potential discrepancies from separate queries. The difference between these two values captures AI’s simplification effect on task skill requirements.