Task-Specific Technical Change and Comparative Advantage*

Lukas Althoff† Hugo Reichardt‡

[Most recent version here]

January 13, 2026

Abstract

Artificial intelligence is changing which tasks workers do and how they do them. Predicting its labor market consequences requires understanding how technical change affects workers' productivity across tasks, how workers adapt by changing occupations and acquiring new skills, and how wages adjust in general equilibrium. We introduce a dynamic task-based model in which workers accumulate multidimensional skills that shape their comparative advantage and, in turn, their occupational choices. We then develop an estimation strategy that recovers (i) the mapping from skills to task-specific productivity, (ii) the law of motion for skill accumulation, and (iii) the determinants of occupational choice. We use the quantified model to study generative AI's impact via augmentation, automation, and a third and new channel—simplification—which captures how technologies change the skills needed to perform tasks. Our key finding is that AI substantially reduces wage inequality while raising average wages by 21 percent. AI's equalizing effect is fully driven by simplification, enabling workers across skill levels to compete for the same jobs. We show that the model's predictions line up with recent labor market data.

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

†Stanford University. lalthoff@stanford.edu

‡Centre de Recerca en Economia Internacional (CREi). hreichardt@crei.cat

1 Introduction

Technological change alters the tasks performed in production and the returns to the skills that determine workers' productivity in those tasks. Turning information on how workers' tasks change into quantitative predictions about the labor market requires an understanding of (i) how skills map into productivity across tasks (and thus govern workers' comparative advantage), (ii) how workers build those skills over their careers, and (iii) how prices, wages, and workers' occupational choices adjust in general equilibrium. This paper develops and estimates a dynamic task-based labor market model that allows researchers to estimate the effects of any task-specific technical change—observed or counterfactual—on individual workers and the overall labor market. We apply this methodology to study the labor market effects of artificial intelligence.

Our model captures key features necessary to understand the labor market effects of task-specific technical change. Workers have multidimensional skills. Their productivity in a given task follows from the match between these skills and the task's skill requirements. Workers learn on the job at a rate that depends on their occupation and their ability to learn. Each period, they choose from a menu of occupations, each consisting of a set of tasks. These occupational choices are forward-looking as workers internalize that their skill accumulation depends on those choices. In equilibrium, demand for each occupation's output equals the amount produced by workers choosing that occupation.

We model technical change as the augmentation, automation, and simplification of tasks. Augmentation increases human productivity in tasks. Automation expands the set of tasks that can be performed without human input. Augmentation and automation are standard features of task-based models (e.g., Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018). Beyond those standard forces, we also consider that technology can reduce the level of skill required to complete a task. We refer to this as *simplification*. Together, the three channels determine how each worker's productivity is affected by technical change.

To illustrate the need for a structural framework, consider radiologists, an occupation where AI is reshaping tasks in multiple ways: AI can assist with detection of abnormalities, automate screening procedures, and simplify report generation (e.g., Hosny et al., 2018; Eloundou et al., 2024; Mousa, 2025). How do these shifts in task composition affect radiologists? Do their skills become more or less valuable in the era of AI? Could they transition to related medical specialties if necessary? If retraining is required, at what pace can workers accumulate new skills? Without answering such questions—jointly and in general equilibrium—one cannot predict how workers will be impacted by AI.

To quantify how task-specific technical change affects workers using our model, we develop a strategy to identify how workers' skills determine absolute and comparative advantage across tasks and how skills accumulate. Workers' optimal time allocation yields a closed-form mapping between task-level productivity and occupational productivity (and observed wages). Combining this mapping with detailed data on skill requirements, we can recover task-level productivity from data on workers' skills, occupations, and wages. While we only observe skills at labor market entry, we identify skill accumulation from workers' occupational history and the evolution of their wages.

We estimate the model parameters that govern workers' productivity, learning, and occupational choices by maximum likelihood using data from the National Longitudinal Survey of Youth 1979 (NLSY79). We develop a procedure that is computationally feasible despite the high-dimensional nature of the model. First, we exploit that many parameters can either be estimated by linear regression or be obtained with a fast iterative routine conditional on the remaining parameters, substantially reducing the dimension of the parameter space that must be searched non-linearly. Second, we recover equilibrium prices from data directly, avoiding the need to solve the model within the estimation routine. In our quantitative application, workers' skills are five-dimensional: manual, social, math, technical, and verbal. In addition, workers differ in their "ability to learn", that is, the rate at which they accumulate skills.[1]

Having estimated the model's parameters, we provide algorithms to solve for equilibrium prices in the steady state and over the transition path. We do so by simulating workers' choices and skill accumulation and updating prices iteratively to equate demand and supply. We show how the solution to the worker's problem can be computed efficiently.

Our quantified model offers a laboratory to study how any counterfactual task-specific technical change affects the labor market.[2] The model predicts

[1]Following Heckman et al. (1998), we estimate learning ability separately by quartiles of the Armed Forces Qualification Test: a measure of ability used by the US military to determine enlistment eligibility.

[2]The estimated model can also be applied to study changes in occupational demand (since the estimation of the supply-side parameters is independent of assumptions on demand).

how technologies change workers' comparative advantage, how workers reallocate and retrain, and how prices adjust in general equilibrium.

We apply the quantified model to predict how AI will affect individual workers and the labor market as a whole. We follow Eloundou et al. (2024); Acemoglu (2025) in using large language models (LLMs) to obtain estimates of AI's capability to augment, automate, and simplify tasks. Our prompts closely follow the survey design used to collect assessments from human experts, while enabling hundreds of thousands of evaluations. We provide evidence that these LLM-generated estimates are reasonable, including validation against human expert assessments and experimental evidence.

Our first key finding is that generative AI substantially reduces wage inequality—an effect fully driven by simplification—and increases average wages by 21 percent. Simplification increases the relative productivity of lower-skill workers in tasks and occupations that were previously the territory of higher-skilled workers. This lowering of skill-based barriers is the key force reducing inequality. Automation and augmentation, on the other hand, drive most of the average effects but do not have quantitatively strong distributional implications. Simplification does not affect the average wage by much because two opposing forces roughly cancel out: it increases productivity for given skills but it reduces skill accumulation.

Second, we find that AI generates sizable welfare gains for almost all workers at labor market entry. We estimate welfare improvements equivalent to permanent wage gains of 26–34% for most workers. Consistent with the decline in wage inequality, we find that the welfare gains are largest for lower-skill workers, reducing the returns to skills. Workers with lower verbal skills see particularly large welfare increases. Math is the dimension of skill for which the returns decrease the least.

Our third key finding is that AI's impact—together with workers' responses to it—significantly alters the occupational landscape. AI generates a large reallocation of employment across occupations. For example, administrative occupations (e.g., financial clerks) see a large decline in employment, while science occupations (e.g., life scientists) expand. On average, wages rise, but some occupations—such as architects, engineers, and executives—see absolute wage declines. In many cases, the occupations that experience the largest employment gains are also those for which relative wages decrease the most. This negative relationship between wage and employment 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 to the simplification mechanism, augmentation raises wages more uniformly across occupations and does not trigger much reallocation. Automation mainly shifts employment away from highly exposed occupations without substantially changing relative wages.

Lastly, early labor market evidence provides suggestive support for our model's predictions. Using an event study design around ChatGPT's November 2022 release, we find that occupations predicted to benefit from AI show differential positive trends in wage bill shares, with approximately 8 percent of our predicted long-run effects materializing by mid-2025. A second event study shows that college major choices have begun to respond consistent with our predicted changes in the labor market returns to majors.

Related literature. To predict how technologies affect workers through their task-specific comparative advantage, we integrate three previously separate literatures—on task-based production, multidimensional skills, and dynamic occupational choice—in a single empirically tractable framework.

Our contribution to the literature on task-based production and technological change is threefold (e.g., Zeira, 1998; Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2022). First, we provide methods to estimate workers' comparative advantage across tasks, a key object in shaping how technical change affects workers. The absence of such methods has inhibited quantifying the effects of future technical change (Woessmann, 2024).[3] Second, while the literature on task-based production treats workers' skills as fixed, we allow for and estimate skill accumulation—a force shaping workers' adaptation to technical change. Third, we integrate task-based production into a general equilibrium dynamic occupational choice model (in the spirit of Keane and Wolpin, 1997; Heckman et al., 1998; Lee and Wolpin, 2006; Dix-Carneiro, 2014; Traiberman, 2019; Smeets et al., 2025), capturing workers' choices over a discrete set of task-bundled occupations (Autor and Handel, 2013; Hurst et al., 2024).

Beyond these methodological contributions, we offer a conceptual innova-

[3] Acemoglu and Restrepo (2022) show that the role of comparative advantage across groups can, to a first order, be captured by a low-dimensional propagation matrix. However, estimating this matrix relies on the identification of the technology's labor market effects, so that it cannot be used to study effects of *counterfactual* technical change.

tion by incorporating simplification into models of task-based technical change. Relatedly, Autor and Thompson (2025) show how automation's effect depends on the expertise the automated task requires compared to the remaining tasks.[4] This paper allows for the quantification of such effects and finds that they drive AI's equalizing distributional impact.

We also build on the literature emphasizing the multi-dimensionality of skills (Lindenlaub, 2017; Guvenen et al., 2020; Lise and Postel-Vinay, 2020; Baley et al., 2022). First, we integrate multi-dimensional comparative advantage and skill accumulation into a model of task-based production. Second, we overcome the empirical and computational challenges that result from this task-based approach with a new estimation strategy. Third, to enable task-level estimation and counterfactual analysis, we construct a database of task-level skill requirements. Prior work relies on O*NET's occupational aggregates (e.g., Lise and Postel-Vinay, 2020; Baley et al., 2022). We extend these to the task-level using large language models and validate the database's accuracy.

Finally, our work relates to a growing literature that quantifies the effects of AI. Freund and Mann (2025) introduce a partial equilibrium framework to understand how automation affects wages through changes in the importance of tasks within an occupation, together with a strategy to estimate the distribution of workers' task-level productivity. In comparison, our approach allows to understand how prices, wages, and workers' skills adjust dynamically to AI in general equilibrium. Hampole et al. (2025) provide a structural framework to quantify AI's effect on occupational demand through actual adoption patterns across firms. In contrast, our paper predicts how AI affects individual workers and the overall labor market by modeling and estimating workers' comparative advantage and skill accumulation.

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 a worker's skills and the skill requirements of the tasks relevant to the occupation of choice. Workers accumulate skills on the job, so that their occupa-

[4] Downey (2021); Danieli (2025) also consider mechanisms related to simplification in other contexts.

tional 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 Problem

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

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

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., $\tau \in \mathcal{N}_j$:

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

where ℓ_τ represents the share of time allocated to task τ, γ_τ and ϕ_τ capture the task-specific productivity of humans and capital respectively, $\boldsymbol{h} = (h_s)'_{s \in S}$ is a vector capturing workers' multi-dimensional skills, $\boldsymbol{r}_\tau \equiv (r_{\tau,s})'_{s \in S}$ 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 $\boldsymbol{h}$ are differentially productive in tasks depending on its skill requirements, $\boldsymbol{r}_\tau$. 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 2.4, 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 $\boldsymbol{r}_\tau$.

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.[5]

Experimental evidence has shown that AI's productivity effects tend to be stronger for less skilled workers, suggesting that simplification can be an important force in practice (e.g., Brynjolfsson et al., 2025b).

The firm's problem. Each good j is produced by a representative firm that takes the equilibrium wage $\{w_j(\boldsymbol{h})\}_{\boldsymbol{h}}$ and the costs of capital R as given. The firm chooses how many workers of each skill level $\boldsymbol{h}$ 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:

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

where $n_j(\boldsymbol{h})$ is the amount of labor employed with skill $\boldsymbol{h}$, $\ell_{j,\tau}(\boldsymbol{h})$ is the share of time allocated to task τ, and $k_{j,\tau}(\boldsymbol{h})$ denotes the capital per worker allocated to workers with skill $\boldsymbol{h}$ 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

[5] While we call this technical change "simplification", the methodology equally allows for increases in skill requirements (or "complication").

their marginal product. That is,

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

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 in so far 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 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 (Task assignment: control and contractibility)**.** The firm controls the workers' allocation of time across tasks and this allocation is not contractible.

Assumption 1 implies that the firm chooses the allocation that maximizes profits (output net of capital costs) for any given wage $w_j(\boldsymbol{h})$. In equilibrium, firms thus choose this value-added maximizing allocation and wages 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, tasks, and skills. That is, for all occupations j, skills $\boldsymbol{h}$, and tasks $\tau \in \mathcal{A}_j$,

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

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

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

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

$$\ell_{j,\tau}(\boldsymbol{h}) = \frac{\theta_{j,\tau}\gamma_\tau^{\rho-1} f(\boldsymbol{h}, \boldsymbol{r}_\tau)^{\rho-1}}{\sum_{\kappa\in\mathcal{N}_j}\theta_{j,\kappa}\gamma_\kappa^{\rho-1} f(\boldsymbol{h}, \boldsymbol{r}_\kappa)^{\rho-1}}, \tag{5}$$

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 most productive tasks, i.e., tasks for which $\gamma_\tau f(\boldsymbol{h}, \boldsymbol{r}_\tau)$ is greater. If instead tasks are complements ($\rho < 1$), workers spend more time on the 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 $\tau \in \mathcal{A}_j$

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

where $Y_j(\boldsymbol{h})$ is the profit-maximizing level of output when a worker with skill vector $\boldsymbol{h}$ 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 $Y_j(\boldsymbol{h})$ is larger.

Occupational productivity. Given the optimal allocation of time and automation technology to the production of tasks, the output per worker of type $\boldsymbol{h}$ is

$$\begin{aligned} Y_j(\boldsymbol{h}) &\equiv \left(\sum_{\tau\in\mathcal{A}_j}\theta_{j,\tau}^{\frac{1}{\rho}}(\phi_\tau k_{j,\tau}(\boldsymbol{h}))^{\frac{\rho-1}{\rho}} + \sum_{\tau\in\mathcal{N}_j}\theta_{j,\tau}^{\frac{1}{\rho}}\left(\ell_{j,\tau}(\boldsymbol{h})\gamma_\tau f(\boldsymbol{h}, \boldsymbol{r}_\tau)\right)^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}} \\ &= \Gamma_j^{\frac{\rho}{1-\rho}}\left(\sum_{\tau\in\mathcal{N}_j}\theta_{j,\tau}\gamma_\tau^{\rho-1} f(\boldsymbol{h}, \boldsymbol{r}_\tau)^{\rho-1}\right)^{\frac{1}{\rho-1}} \end{aligned} \tag{7}$$

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

Wages. Combining equation (4) and (7) yields the wage when a worker of skill $\boldsymbol{h}$ chooses occupation j:

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

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

Skill accumulation. Before entering the labor market at age $a = 1$, each worker draws an initial skill vector $\boldsymbol{h}_1$ 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 job j: $\boldsymbol{h}' = g_j(\boldsymbol{h}, \psi)$.

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. Their expected lifetime utility at age a when their previous occupation is k is represented by the value function

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

where $\mathbb{E}[\cdot]$ is the expectation over occupation-specific productivity shocks. The value after the terminal age A is zero, $V_{A+1}(\cdot, \cdot) = 0$. μ_j is the amenity value of occupation j. $g_j\left(\boldsymbol{h}, \cdot\right)$ is next period's human capital when choosing occupation 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.[6]

We assume that the log 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

[6]We assume that occupational switching costs do not apply in the first period.

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

has the closed-form solution

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

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

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

Since $V_{A+1}(\cdot,\cdot) = 0$, 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.[8] Note that we have suppressed any dependence on time in the model above. In principle, prices vary over time, so that the wage schedule $w_{j,t}(\boldsymbol{h})$, and thus the value 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 $D\left(\{p_j\}_{j=1}^{J}\right)$ that maps prices p_j into relative demand for each occupational good. In our application, we use CES demand. Having specified demand, we can now define the competitive equilibrium.

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

[8] Equation (7) solves for wages $w_j(\boldsymbol{h})$ given prices.

– The distribution over states follows from occupational choices:

$$G_{a+1,t+1}\left(\boldsymbol{h}',\psi,j\right)=\sum_{k=1}^{J}\int_{g_j(\boldsymbol{h}',\psi)\leq\boldsymbol{h}}\mathbb{P}_{a,t}\left(j\mid\boldsymbol{h},\psi,k\right)dG_{a,t}\left(\boldsymbol{h},\psi,k\right); \tag{12}$$

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

$$\mathcal{Y}_{j,t}\equiv\sum_{a=1}^{A}\sum_{k=1}^{J}\int Y_j(\boldsymbol{h})\,\mathbb{P}_{a,t}(j\mid\boldsymbol{h},\psi,k)\,\mathbb{E}[\varepsilon_j\mid j,\boldsymbol{h},\psi,k]\,dG_{a,t}(\boldsymbol{h},\psi,k);^{9} \tag{13}$$

– Demand for capital equals supply:

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

2.3 Solution Methods

We provide algorithms to solve for a stationary competitive equilibrium as well as 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 $\left(p_1^{(1)},\cdots,p_J^{(1)}\right)$.
2. For iteration r, given the prices $\left(p_1^{(r)},\cdots,p_J^{(r)}\right)$, solve the worker's problem and compute the implied output of each good $\left(\mathcal{Y}_1^{(r)},\cdots,\mathcal{Y}_J^{(r)}\right)$. Then, update prices to clear the market given supply: $\boldsymbol{p}^{(r+1)}=D^{-1}\left(\{\mathcal{Y}_j^{(r)}\}_{j=1}^{J}\right)$.
3. Repeat step 2 until $\left\|\boldsymbol{p}^{(r+1)}-\boldsymbol{p}^{(r)}\right\|<\epsilon$ for a threshold $\epsilon>0$.

Transition path. Starting from the initial stationary equilibrium, we solve for the transition path of prices $\{p_{1,t},\cdots,p_{J,t}\}_{t=1}^{\infty}$ from the moment the shock is realized. We numerically approximate this infinite sequence by solving for

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

$\{p_{1,t}, \cdots, p_{J,t}\}_{t=1}^{T}$ 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.[10]

3. For iteration r, given the sequence of prices $\left\{p_{1,t}^{(r)}, \cdots, p_{J,t}^{(r)}\right\}_{t=1}^{T}$, solve for the value function at $t = T, T-1, \ldots, 1$. Then, compute the implied output of each good at each time and the corresponding prices $\boldsymbol{p}_t^{(r+1)} = D^{-1}\left(\{\mathcal{Y}_{j,t}^{(r)}\}_{j=1}^{J}\right)$ for $t = 1, \ldots, T$.

4. Repeat step 3 until $\left\|\boldsymbol{p}_t^{(r+1)} - \boldsymbol{p}_t^{(r)}\right\| < \epsilon \; \forall t = 1, \ldots, T$ for a threshold $\epsilon > 0$.

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

First, we solve for the value function. This step is conceptually straightforward. However, when skills are multi-dimensional and there are many occupations j, the state space $(\boldsymbol{h}, \psi, j)$ is large and value function iteration costly. We exploit a convenient feature of the logit to provide relief. Since the occupational switching cost is $\kappa(j,k) = \kappa 1[j \neq k]$, the value function in equation (11) can be written as

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

where $\tilde{V}_a^j(\boldsymbol{h}, \psi) \equiv \exp\left(\frac{1}{\zeta}\left(\log w_j(\boldsymbol{h}) + \mu_j + \beta V_{a+1}\left(g_j(\boldsymbol{h}, \psi), \psi, j\right)\right)\right)$. This solution implies that it is sufficient to solve for $\tilde{V}_a^j(\boldsymbol{h}, \psi)$. Using this property, we effectively shrink the state space in the value function iteration step by a factor J (in our exercise, $J = 93$). Conditional choice probabilities can then be recovered as

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

[10] A reasonable guess is the path where prices adjust immediately to the new stationary equilibrium.

Second, after computing the value function and conditional choice probabilities, we compute the joint distribution of age a, human capital $\boldsymbol{h}$, learning ability ψ, and occupations k using the law of motion in equation (12). We do so by simulation. That is, we first draw from the initial distribution of skills and learning ability at age $a = 1$. Their states imply a conditional probability to choose each occupation. We then draw occupational choices randomly based on these probabilities to obtain the distribution at age $a = 2$. We iterate this process forward until $a = A$.[11]

Third, to obtain an update for 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 $Y_j(\boldsymbol{h})\,\mathbb{P}_{a,t}(j \mid \boldsymbol{h}, \psi, k)\,\mathbb{E}[\varepsilon_j \mid j, \boldsymbol{h}, \psi, k]$ for each worker-age and for each $j = 1, \ldots, J$. In other words, 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.

2.4 Parametrization

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

Production. We specify the task-level production function as

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

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 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.

[11] To 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 decrease.

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

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

Skill accumulation. We assume that the human capital accumulation function has the following functional form:

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

where $\ell_{j,\tau}(\boldsymbol{h})$ 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 $\ell_{j,\tau}(\boldsymbol{h})$ is greatest. Second, workers learn by performing tasks that are "hard" for them, i.e., 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.[12] 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 main data sources to estimate the model's parameters. First, we rely on O*NET to measure each occupations' tasks and skill requirements. Second, we provide and validate a new database of task-level skill requirements by extending O*NET's occupation-level survey on skill requirements to the task-level using large language models. Third, we use panel data on wages, occupational choices, and multidimensional skills from the NLSY79.

For the application of our methodology to artificial intelligence, we also re-

[12]Since $f(x) = x \exp(-\lambda x)$ 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 $1/(e\lambda)$.

quire data on AI's task capabilities. We follow the literature in using large language models to estimate these capabilities (e.g., Eloundou et al., 2024; Acemoglu, 2025).

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). O*NET contains detailed descriptions of 19,530 tasks linked to 974 occupations. We rely on these data to define both the occupations and tasks in our model. That is, we set the tasks employed across occupations in our model, $\mathcal{T}_j$, to mirror those in the O*NET data. 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.[13]

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 means being able to "read step-by-step instructions for completing a form" and 4 to "understand an email from management describing new personnel policies").[14]

We reduce O*NET's dimensions into five skill categories: manual, mathematics, social, technical, and verbal.[15] Table B.1 shows this mapping. Each skill's requirements equal the average of its related O*NET skills.

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, $\boldsymbol{r}_\tau$. While O*NET provides data on occupation-level skill requirements, it lacks task-specific data. To address this gap, we use OpenAI's GPT-4o to estimate the task-level skill requirements. To ensure consistency with O*NET and a valid survey design, we replicate O*NET's occupation-level questionnaire on the level of the task, by using their

[13] The importance weights are normalized to sum to 1 within occupations.

[14] The original O*NET questionnaire and skill level descriptions are available here.

[15] Relative to Addison et al. (2020); Baley et al. (2022); DeLoach et al. (2022), we include manual as a separate skills because we believe its interaction with AI is of particular interest.

questionnaire format, skill dimensions, and skill anchors. This process covered 19,530 task descriptions across 35 skills using 683,550 queries. See Appendix D.1 for further prompt design details.

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 0.82 to 0.93 (see Figure A.2).[16]

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. As Addison et al. (2020); Baley et al. (2022), we measure skills with the Armed Services Vocational Aptitude Battery (ASVAB): manual skills are measured as the average of scores on *auto and shop information* and *mechanical comprehension*, math skills are based on *mathematics knowledge* and *arithmetic reasoning* scores; technical skills on *general science* and *electronics information*; verbal skills on *paragraph comprehension* and *word knowledge*.[17] We standardize each of the subscores before aggregating. 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. That is, we only observe $\tilde{\boldsymbol{h}} = F(\boldsymbol{h})$ where $F(\cdot)$ is the distribution function of the initial skill distribution. We do not directly observe the cardinal measure $\boldsymbol{h}$ that is on the same scale as the skill requirements. We therefore 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

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

[17] Relative to Addison et al. (2020) and Baley et al. (2022), we add manual skills as a separate dimension.

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 which 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 presented in section 2, AI's capabilities can take three forms: augmentation, automation, and simplification. Below we describe how we estimate these capabilities by task. We acknowledge that there is substantial uncertainty surrounding these capabilities. However, we view these estimates as reasonable and provide evidence to that effect. As our baseline, we only consider the effects of generative AI (data summarized in Table B.3).[18] However, we also consider smart robots and autonomous vehicles in an alternative scenario.[19]

Automation. We also follow Eloundou et al. (2024) by eliciting automatability by task from large language models. That is, we ask for each task in the O*NET database whether AI can complete the task autonomously. From the perspective of the model, we view this as asking whether a 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.[20] We classify a task as "automatable" if it has high or full automation exposure, which is restricted to cases where the LLM indicates that generative AI can complete at least 90% of the components of the tasks. 22.2% of all tasks are classified as automatable by generative AI (see Table B.3). The prompt is documented in Appendix D.2.2.

We find high agreement rates 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 on the share of automated tasks by occupation ($\rho = 0.82$, see Figure A.4b). Table B.2 shows that agreement is also strong on the task-level.

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 time required to complete a

[18]We use Gartner's definition which can be found here.
[19]For definitions of these technologies, we again follow Gartner: see here and here.
[20]The specific prompt is documented in (Eloundou et al., 2024, Supplementary Materials).

task by at least half. We replicate their exercise with GPT-4o, except that we asked for a continuous estimate of the percentage of time saved rather than a binary measure and consider generative AI more broadly. On average, we estimate that generative AI saves 20.2% of worker time (see Table B.3). Appendix D.2.1 describes our prompt design. The prompt also describes how we extend it to smart robots and autonomous vehicles.

We validate our new data on task-level AI augmentation in two ways. First, we find that our estimates are strongly correlated with both the human rated and GPT-4 rated data from Eloundou et al. (2024), especially considering that their measures are binary (see Figure A.4a). Second, we compiled experimental estimates of average productivity effects of generative AI in various tasks and occupations and compared them with our estimates (see Table B.4). Reassuringly, our estimates closely approximate those experimental estimates.

From the perspective of the model, average productivity effects are driven both by augmentation, governed by γ_τ, and simplification, changes in task's skill requirements (see below). To estimate augmentation based on average productivity effects, we first compute the average productivity effects resulting from simplification alone based on the change in skill requirements and the pre-AI distribution of skills by occupation. We set augmentation γ_τ to the average productivity effect net of this simplification-led productivity effect.

Simplification. Lastly, we elicit the degree to which AI changes tasks' skill requirements, $\boldsymbol{r}_\tau$. In addition to our new data on pre-AI task-level skill requirements, we prompt GPT-4o to evaluate the task's skill requirements before and after workers gain access to generative AI. The prompt can be found in Appendix D.2.3. We estimate that across all tasks and skill dimensions, the average required level falls by 18.3% once workers get access to generative AI (based on O*NET's 7-step scale, see Table B.3). A one-step reduction (out of 7) is the most common change. We cannot directly validate the accuracy of these predicted changes. However, we do find that the predicted pre-AI skill requirements strongly correlate ($\rho = 0.86$) with those resulting from the prompt to elicit task-level skill requirements in section 3.1.2, which contained no reference to AI (internal consistency).

AI capabilities across occupations and skills. The degree to which an occupation is affected by these three channels is positively correlated. Automatability is highly correlated with augmentation ($\rho = 0.89$). Occupational tasks experiencing strong augmentation also see the greatest skill requirement reductions ($\rho = 0.90$).

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.

Lastly, we find that augmentation and simplification are most common for tasks with initially high skill requirements. Automation is less correlated with skill requirements: if anything, the middle-skilled tasks are most prone to automation (see Figure A.5).

4 Estimation

We jointly estimate the parameters governing productivity, skill accumulation, occupational choices, and the initial skill distribution. We provide a computationally efficient methodology to do so using direct inference on the NLSY79 estimation sample. Importantly, it recovers the equilibrium prices directly, avoiding the need to solve for the equilibrium within the estimation loop. Table B.5 shows an overview of all model parameters and their estimated values. In this section, we discuss the procedure in detail.

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. Relative to a full maximum likelihood approach, we reduce the computational burden in two main ways. First, we use a sequential approach. That is, we maximize the likelihood

$$L\left(\boldsymbol{\theta}_1, \boldsymbol{\theta}_y(\boldsymbol{\theta}_1), \boldsymbol{p}(\boldsymbol{\theta}_1), \boldsymbol{\mu}(\boldsymbol{\theta}_1)\right)$$

with respect to $\boldsymbol{\theta}_1$, where $\boldsymbol{\theta}_y = \{\eta, \{\omega_s\}_{s\in S}\}$ are task-level productivity parameters, $\boldsymbol{p} = \{p_j\}_{j=1}^J$ are equilibrium prices, and $\boldsymbol{\mu} = \{\mu_j\}_{j=1}^J$ are occupational amenities. We show how to obtain consistent estimates of $\boldsymbol{\theta}_y(\boldsymbol{\theta}_1)$, $\boldsymbol{p}(\boldsymbol{\theta}_1)$, and $\boldsymbol{\mu}(\boldsymbol{\theta}_1)$ using closed forms and fast iterative algorithms for a given $\boldsymbol{\theta}_1$. This reduces the number of parameters over which to maximize the likelihood fully non-linearly by 95%, cutting the computational cost dramatically. Second, we only maximize the likelihood of the old population's occupational choices (for whom the problem is static), avoiding repeated solution of the dynamic value function.

Inner algorithm. The first step in the inner algorithm is to compute the workers' skills given $\boldsymbol{\theta}_1$. 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 $\boldsymbol{h}$ using the marginal distribution, approximated with a Beta distribution with parameters in $\boldsymbol{\theta}_1$.[21] 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 $\boldsymbol{j}_i^{a-1} \equiv \{j_{i,1}, \ldots, j_{i,a-1}\}$, worker i's skill level at age a is

$$\boldsymbol{h}_{i,a}(\lambda(\psi_i), \delta) \equiv \left(g_{j_{i,a-1}}(\cdot, \psi_i) \circ g_{j_{i,a-2}}(\cdot, \psi_i) \circ \ldots \circ g_{j_{i,1}}(\cdot, \psi_i)\right)(\boldsymbol{h}_{i,1}).$$

where $(f \circ g)(x) \equiv f(g(x))$.[22] 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 $\boldsymbol{\theta}_1$.

We estimate the occupational wage functions using a simple linear regression given workers' skills. The derived occupational wage function in equation (15) governs how skills translate into earnings in each occupation depending on the prices, the occupation's tasks, and the parameters of the production function $\{\omega_s\}_{s\in S}$ and η. A log-linearization of this function (around no mismatch) implies that

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

where $\tilde{\theta}_{j,\tau} \equiv \theta_{j,\tau}\gamma_\tau^{\rho-1}$ (see Appendix C.2 for the proof).[23] Equation (17) shows that we can estimate the equilibrium prices p_j and the parameters of the pro-

[21]The 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.

[22]To save notation, it is left implicit above that $\boldsymbol{h}_{i,a}$ depends on $\lambda(\psi_i)$ and δ through $g_j(\cdot, \psi_i)$.

[23]For estimation, we assume that $\sum_{\tau\in\mathcal{T}_j} \tilde{\theta}_{j,\tau} = 1$ for all $j = 1, \ldots, J$ and that O*NET's task-importance weights capture $\tilde{\theta}_{j,\tau}$. Also, since we estimate the model on data before the change of interest, the equation above reflects wages when no task is automated, i.e., $\mathcal{A}_j = \varnothing$.

duction function using a simple OLS regression of wages on occupational fixed effects, skills, and skill mismatch. We use a control function approach to correct for selection on the productivity shocks ε_j (Dubin and McFadden, 1984).[24]

Given the wage function, we estimate the occupational amenities $\boldsymbol{\mu}$ using a fast iterative procedure. The worker's occupational choice problem is static at the terminal age since their choice probabilities no longer reflect occupations' differential learning value. That is, for $a = A$, equation (10) simplifies to

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

where ζ (scale of productivity shocks), κ (switching costs) are in $\boldsymbol{\theta}_1$ and thus taken as given in this step. The likelihood is maximized with respect to $\boldsymbol{\mu}$ when the observed share of workers in each occupation j, s_j, equal the model-implied share, $\tilde{s}_j(\boldsymbol{\mu})$. We solve for this $\boldsymbol{\mu}$ using the contraction mapping proposed by Berry et al. (1995):

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

where Ψ is a damping parameter. In practice, we use the SQUAREM algorithm to accelerate convergence (Varadhan and Roland, 2008; Reynaerts et al., 2012; Conlon and Gortmaker, 2020).

Outer algorithm. In the outer algorithm, we optimize over the remaining parameters governing skill accumulation function, occupational choices, and initial skills. We choose these parameters to maximize the joint likelihood of the wage function and the occupational choices (of the old population). That is,

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

where $\tilde{w}_{i,a,j}(\boldsymbol{\theta}_1)$ is the expected wage based on the inner-step given $\boldsymbol{\theta}_1$.

[24]Due to non-random occupational choice, the expected value of $\log \varepsilon_j$ conditional on choosing j is $-\zeta \log \mathbb{P}_a(j \mid \boldsymbol{h}, \psi, k)$. We control for this term in the regression. We estimate the probabilities using occupation-specific logit regressions that condition on workers' previous occupation, 10-year age bins, and each dimension of their initial skill.

In the model, the choice-relevant shocks $\log \varepsilon_j$ are Gumbel. To allow for additional *choice-irrelevant* wage noise (measurement error, idiosyncratic pay), we add an independent term ν_j so that the total wage shock is $\log \epsilon_j \equiv \log \varepsilon_j + \log \nu_j$. We take $\log \nu_j$ to be Gaussian and approximate $\log \epsilon_j$, and thus the density function $\pi(\cdot)$, by a normal distribution.[25]

In a final step, we estimate the occupational amenities $\boldsymbol{\mu}$ using the full population of workers taking all other parameters and estimated prices as given. Within the inner algorithm of the main estimation procedure, we only use the old population to avoid having to solve for the value function within the estimation routine. To match employment shares for the whole working population, we run an analogous BLP-procedure on the full population, solving the worker's dynamic problem for each iteration.[26] The estimates of the occupational amenities resulting from the inner algorithm and this dynamic BLP procedure are almost identical (correlation: 0.99).

Estimation results. Table 1 shows the results of the estimation of the production function in equation (17). The first five columns show the degree which various skills increase productivity across all skills. We find that the returns to math and social skills are highest, consistent with Deming (2017). Importantly, we also find that the cost of underqualification is substantial, yielding strong comparative advantage across tasks with different skill requirements. The coefficient on η implies that if a worker's skill is one level below the task's skill requirement (on O*NET's 1 to 7 scale) in one of the skill dimensions, their task-specific productivity is around 4.4% lower than that in tasks for which they meet the skill requirements in every dimension.

[25] Strictly speaking, the sum is not Gaussian and the distribution becomes a convolution. However, given that we find that the variance of the choice-irrelevant shocks is considerably larger than that of $\log \varepsilon_j$, the impact of this simplification is minimal.

[26] In 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.

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, the estimated occupational prices $\hat{p}_j$ are strongly correlated with the skill requirements in the respective occupations: skill requirements explain around 73% of the variance in prices across occupations (see Table B.6).

We show estimates of the parameters that determine initial skills and skill accumulation in Table 2. The depreciation rate of human capital when doing work for which one is overqualified is 0.0003. $\lambda(\psi)$ is inversely related to someone's ability to learn. The results in Table 2 thus suggest that the learning cost decreases with the AFQT score. Lastly, B_a and B_b are the shape parameters of the initial Beta distribution of skills. The implied average is $\frac{B_a}{B_a+B_b} = 0.35$. This translates to an average of 3.11 on the original O*NET scale from 1 to 7, in between the "low" and "medium" skill requirement levels. Figure A.6 plots the 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, our estimate of the scale parameter $\hat{\zeta} = 0.053$ and of 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% wage loss.

Calibrated parameters. Some parameters are set externally. In sections 3.1.1 and 3.1.2 we explain how we measure the task set $\mathcal{T}_j$, the task weights $\theta_{j,\tau}$, and the skill requirements $\boldsymbol{r}_\tau$ for each occupation j and for each task $\tau \in \mathcal{T}_j$. We set the number of periods A to 40 so that each period in the model represents a year between ages 25 and 64. Following Keane and Wolpin (1997), we set the discount factor β to 0.78 (see also Postel-Vinay and Robin (2002) for similar estimates).

We set the elasticity of substitution between occupations σ to 1.57—the mid-point between 1.81 (Burstein et al., 2019) and 1.34 (Caunedo et al., 2023)—and the substitutability between tasks ρ to 0.49, as estimated by Humlum (2019).

Lastly, we need to calibrate the share of income that will accrue to labor in each occupation, Γ_j in equation (8). In Appendix C.1, we derive how this share is identified from the share of tasks that are automated and the average cost savings by automated task:

$$\Gamma_j = 1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \left(\frac{R/\phi_\tau}{p_j} \right)^{1-\rho} = \frac{\chi^{\rho-1} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \right)}{\sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} + \chi^{\rho-1} \left(1 - \sum_{\tau \in \mathcal{A}_j} \theta_{j,\tau} \right)}. \tag{21}$$

where χ is the unit cost of producing tasks with AI relative to the unit cost of producing it with labor. Using experimental evidence, Acemoglu (2025) estimates the cost savings of AI in automatable tasks to be 27%. Hence, we set $\chi = 0.73$.

Demand for occupations. Lastly, we estimate the demand for occupations.[27] We assume that occupational goods are substituted with a constant elasticity of substitution (CES) σ. Formally, demand for occupation j, $D_j(\{p_j\}_{j=1}) \propto \alpha_j p_j^{-\sigma}$ where α_j is the CES weight of occupation j. This demand system implies that, for two occupations i and j,

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

[27] Note that we estimate all supply-side parameters independently of demand. This is an advantage as it allows to change the demand structure without having to re-estimate any other parameters.

We compute occupational wage shares from the 2018 BLS Occupational Employment and Wage Statistics (OEWS). The occupational fixed effects in equation (17) are consistent estimates of the (log) occupational prices.[28] From those estimates, we can compute the implied weights $\{\alpha_j\}_{j=1}^{J}$ for a given σ.

4.2 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 1 reports how well the moments from this simulated panel match the data.

First, Figure 1a 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, so it is not surprising that inequality is somewhat underestimated. However, this underestimation is quite limited. For instance, the ratio between the 75$^{\text{th}}$ and the 25$^{\text{th}}$ percentile is 2.04 in the data, compared to 1.84 in the model and the top 10% wage share is 20% in the model, compared to 26% in the data. Table B.7 reports how various other measures of inequality compare between the model and the data.

The model also accurately replicates patterns of occupational sorting. Figure 1b shows the correlation between the average skill by occupation in the model and the NLSY79 data. To compute this correlation, we only use the occupational choices of the young population for which we observe the skills directly from the skill assessment scores.[29] The correlations range between 0.6 and 0.8 across skill dimensions, implying that 1) the skill assessment scores in the NLSY79 are predictive of occupational choices (see also Lise and Postel-Vinay, 2020) and 2) workers in the model select into occupations based on their skills in ways similar to that observed in the NLSY79.

Figure A.7a shows that the average wage by occupation matches the data almost perfectly. This is the most directly targeted moment, as we estimated demand based on occupational wage shares and job-specific amenities based on occupational employment shares.

[28] To reduce noise in the price estimates, we apply empirical Bayes regression to the price predicted by the skill requirements (see e.g., Walters, 2024).

[29] This makes the test as stringent as possible because it prevents the skills in each occupation to "mechanically" reflect the skills required in the occupation through the estimated learning.

The median wage by age also matches the pattern observed in the data. 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. Furthermore, 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.

The model also accurately reflects the probability that a worker changes occupation from one year to another. The probability of staying within the same 3-digit occupation is 0.86 in the model and 0.90 in the CPS data. This moment is directly targeted by the switching cost parameter κ. However, we also find that the model fits the (untargeted) probability that a worker stays within a broader 2-digit occupational group well: 0.92 (model) and 0.94 (CPS data). In other words, even though the occupational switching cost applies equally across all but one occupation, the model captures that workers are more likely to stay within a similar set of occupations.

We further 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. In other words, the model accurately predicts occupational transitions across 23 broader occupational groups. Within those groups, occupational transitions are harder to predict, because occupations are more similar in skill requirements within those groups.

FIGURE 1: MODEL FIT: COMPARING MODEL MOMENTS WITH DATA

(A) WAGE DISTRIBUTION

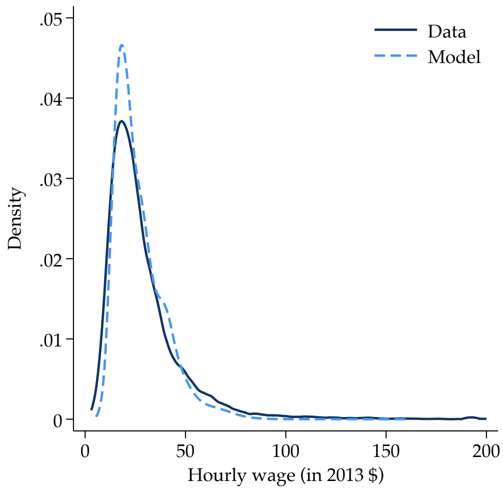

(B) SKILL SORTING BY OCCUPATION

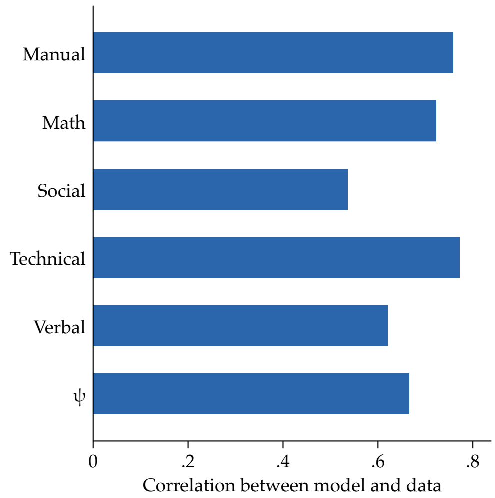

Notes: Panel A shows a kernel density plot of the wage distributions of the NLSY79 data and the model's steady state. Panel B reports the correlation between the average skill by occupation in the model's first period and the NLSY79.

5 Artificial Intelligence and the Labor Market

This section applies our model to understand how AI affects labor markets. We consider AI-induced augmentation, automation, and simplification and study its general equilibrium effects on wages, wage inequality, welfare, skill returns, and occupations.

5.1 AI's Effect on Wages and Inequality

We begin by studying AI's impact on the steady state wage distribution. Figure 2 shows sizable average wage gains of 21%. These gains are concentrated at the bottom of the distribution and are nearly zero at the 99^{th} percentile.

FIGURE 2: WAGE EFFECTS ACROSS THE DISTRIBUTION

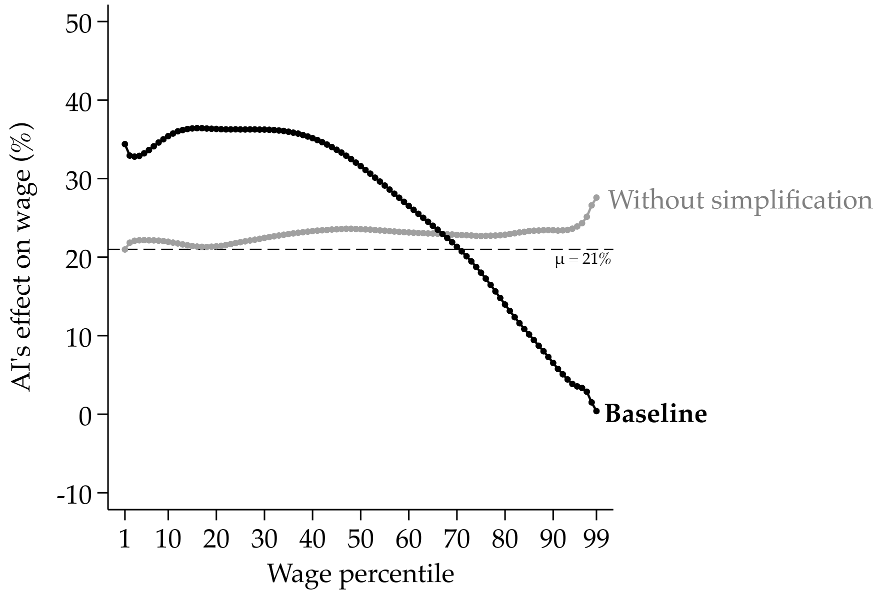

Notes: This figure shows the distribution of wage changes induced by generative AI across the wage percentile distribution. The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile. The black line shows the joint effect of AI's augmentation, automation, and simplification on each wage percentile. The gray line shows the effects if AI induced only augmentation and automation, but no simplification.

To understand these distributional effects, the figure also decomposes the technology's impact by isolating the role of simplification. Without simplification—that is, with only augmentation and automation operating—average wage gains would be moderately larger, but inequality would slightly increase rather than strongly decrease. Simplification thus emerges as AI's key mechanism for reducing inequality. Average wages, in contrast, rise mainly due to augmentation, with automation and simplification having smaller effects on average (see also Appendix Figure A.8, which isolates each mechanism).

Simplification lowers inequality in two ways. First, it reduces wage dispersion within occupations by enabling lower-skilled workers to perform tasks more productively. Second, it reduces wage differences across occupations by making occupations with high skill requirements more accessible to less skilled 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 affect tasks that require systematically different skills from the rest of the occupation's tasks. In such cases, augmentation and automation induce an indirect form of simplification by changing tasks' effective weights within an occupation. However, we find that these indirect effects are quantitatively negligible 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—such productivity differences translate only weakly into relative wage changes.

Turning from distributional to average effects, each mechanism shapes wages through competing theoretical forces. Simplification weakly increases productivity holding skill constant (see equation 14), but limits opportunities for learning; for AI, we find that the net effect is slightly negative (-3%). Automation raises productivity but displaces labor with capital (Acemoglu and Restrepo, 2018); here the productivity effect dominates, raising average wages by 6%.[30] Augmentation, in contrast, has an unambiguously positive theoretical effect on wages, and we find it to be quantitatively large.

5.2 Which Workers Gain or Lose the Most?

We next examine the implications of AI for workers' ex-ante welfare, given the sizable effects on the wage distribution. 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 (regardless of skill and occupation) that delivers the same welfare gain as the introduction of AI. Figure A.9 reports the distribution of this measure, which lie between approximately 26 and 34 percent for most workers, implying sizable ex-ante welfare gains for almost everyone. The welfare gains exceed the average wage increase because utility is concave in income and AI disproportionately raises wages at the bottom of the distribution.

Consistent with the decrease in wage inequality, we find that the ex-ante welfare gains are largest for less skilled workers. Figure 3 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

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

verbal skills (on the O*NET scale from 1 to 7) decreases the welfare gains from AI by 0.7%. Higher math skills, in contrast, have the least negative association with AI-induced welfare effects.

FIGURE 3: HOW AI'S WELFARE EFFECTS DIFFER BY SKILLS

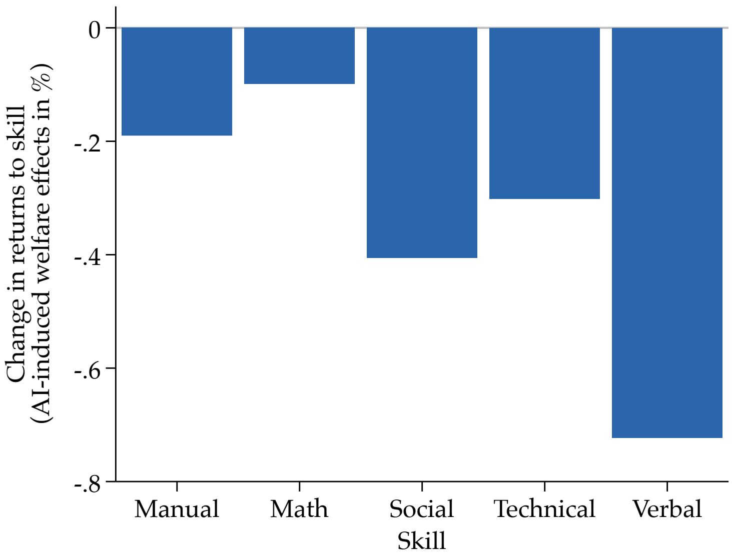

Notes: This figure shows how welfare effects differ by skills. The welfare effects are measured in equivalent permanent percentage wage increases. This figure plots the coefficient of a regression of these welfare effects on skill levels across all dimensions. For interpretability, the skills are expressed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

5.3 How Do Occupations Change?

We next ask how occupations' employment and wages are affected by AI.

First, there are strongly heterogeneous effects on occupation's total wage bills, average wages, and employment shares (see Figure A.10). While wages increase by 21.0% on average, wages in some occupation decline in absolute terms (Figure A.10b). Because our framework allows for occupational re-sorting, part of these occupational effects reflects selection. The size of this occupational reallocation is evident in Figure A.10c, which shows that some occupations lose more than 50% of their employment.

We then zoom in on individual occupations. Figure 4 shows AI's effects on the wage bills of 2-digit occupational groups. *Community and Social Service* experiences the largest wage bill increase, while *Office and Administrative Support* sees an absolute decline in its wage bill. The wage bill is the average wage times the employment. Figure A.11 shows that the effects on employment and wages often work in the opposite direction. For instance, *Architecture and Engineering* experiences the largest increase in employment share and the largest decrease in average wages. *Building Cleaning and Maintenance* experiences the largest increase in average wages and a decline in employment.

FIGURE 4: AI'S EFFECT ON OCCUPATIONS' WAGE BILLS

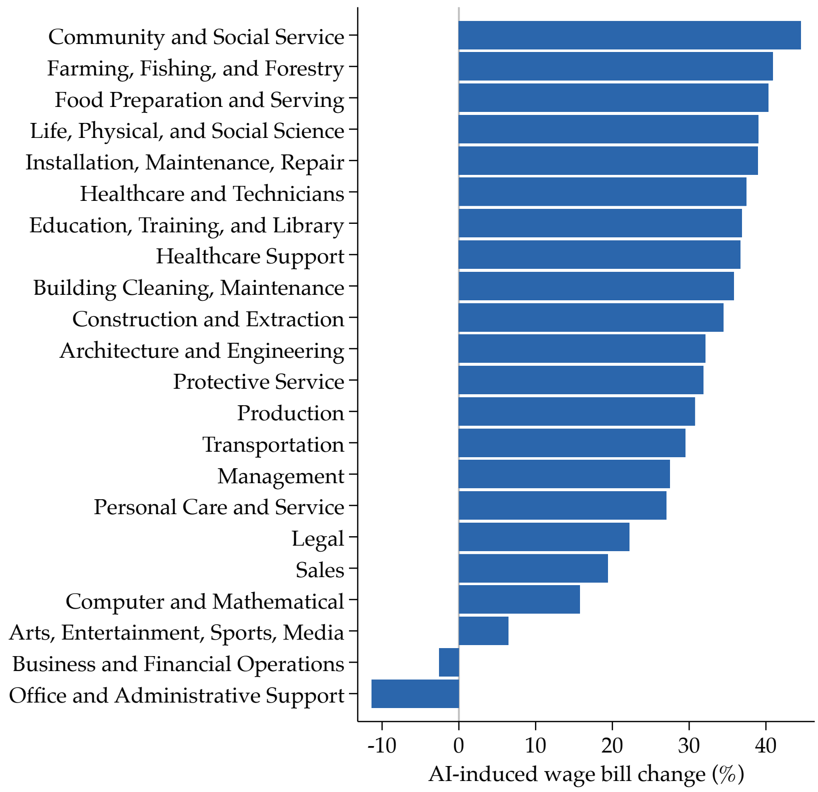

Notes: This figure shows the model predictions on AI's wage bill effects by occupational group. Appendix Figure A.12 disaggregates these effects across detailed occupations.

We assess how each AI channel—augmentation, automation, and simplification—contribute to these occupational outcomes. We recompute occupational outcomes under all possible combinations (e.g., only augmentation, only augmentation and simplification, etc.) and decompose the total effect into the contributions of the three channels. Figure 5 summarizes these results. First, augmentation generates little change in employment shares and raises average wages almost uniformly across occupations. Second, automation leads to large changes in employment, but not to substantially different wage growth across occupations. Finally, simplification generates sizable and opposing effects on employment and wages: by lowering skill requirements, it expands the pool of workers who can perform the occupation productively, which raises employment but compresses average wages. Appendix Figure A.13 shows this analysis for more detailed occupational categories.

FIGURE 5: AI'S EFFECT ON OCCUPATIONAL EMPLOYMENT & WAGES

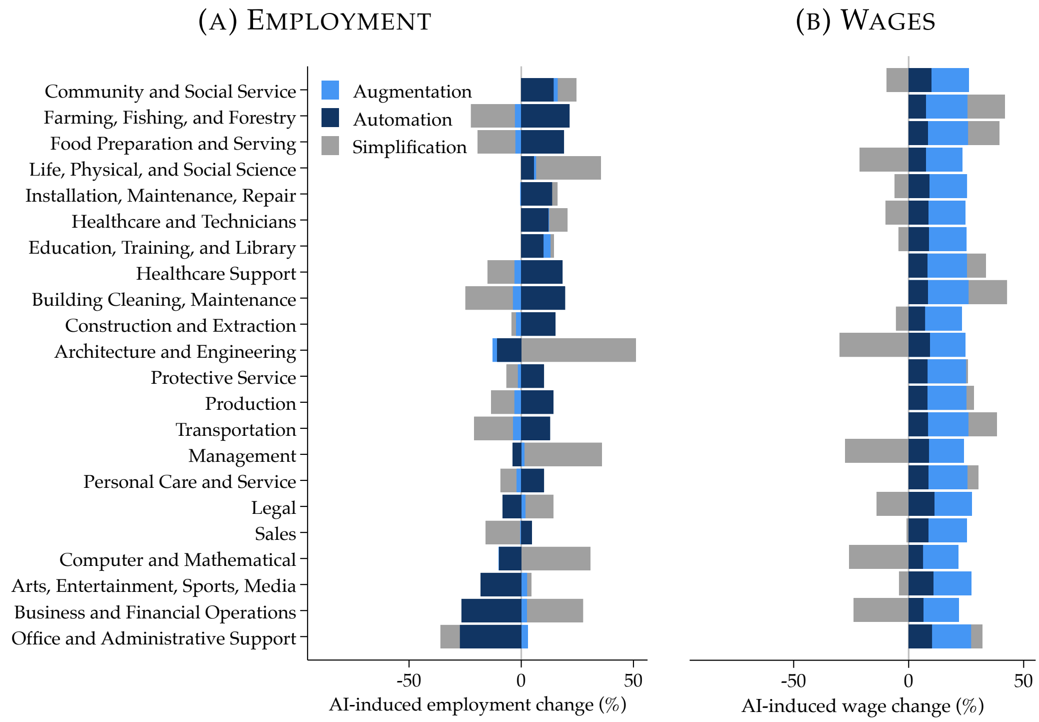

Notes: This figure shows the model's predictions on AI's employment and wage effects by occupational group. Occupations are sorted in descending order of AI's effect on the total wage bill, so that the first listed occupation experiences the largest wage bill increase. We conduct a Shapley-Owen decomposition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

The regression results in Table B.8 systematically relate occupational outcomes to augmentation, automation, and simplification exposure. It confirms that i) augmentation is not a major driver of relative employment or wage changes, ii) automation mostly reallocates employment to less exposed occupations while not having strong effects on wages, and iii) simplification leads to relative wage declines and employment growth. We further consider how augmentation and automation can indirectly induce simplification by shifting the effective weights of tasks with different skill requirements (Autor and Thompson, 2025; Freund and Mann, 2025). Such indirect simplification has effects on employment and wages similar to those of direct simplification.

What characterizes occupations that gain the most from AI? Perhaps surprisingly, there is only a weak relationship between labor market gains and occupations' pre-AI skill or education levels (see Figure A.15a). Both the top and bottom deciles by wage bill increases have similar skill requirements and education levels. This weak overall relationship masks a U-shaped pattern: through the 80th percentile, occupations with larger gains tend to have progressively lower skill requirements (except manual skills), but this reverses sharply at the top, where the highest-gaining occupations skew toward higher skill intensity. Education follows the same pattern.

5.4 AI with Physical Capabilities

Lastly, we examine how these occupational effects change when considering AI technologies with physical manipulation capabilities, such as AI-powered robots and self-driving vehicles.These estimates are necessarily subject to more uncertainty, but highlight how technological capabilities determine which human skills retain value and thus which workers benefit or lose.

The addition of physical capabilities substantially amplifies AI's labor market impact on the wage distribution. Average wages rise by 39 percent in this scenario (compared to 21 percent with only generative AI). Changes in inequality, previously shown to be driven mainly by simplification, follow a similar pattern in both AI scenarios.

There are notable shifts in the patterns of occupational reallocation as AI gains physical manipulation capabilities (see Figure 6). *Community and Social Service* and *Education, Training, and Library* occupations are the main winners, more than doubling in wage bill. In contrast, a larger number of occupational groups now lose over a quarter of their wage bill, including *Office and Administrative Support*, *Transportation*, and *Production* occupations.

Several occupations that were predicted to experience large gains from generative AI are predicted to experience large losses if AI gains physical capabilities. The most striking reversals occur for occupations in *food preparation and serving*, *farming, fishing, and forestry*, *production*, and *transportation* occupations. Those are occupations requiring manual skills that AI with physical capabilities (but not generative AI) can automate (see also Figure A.16a). Overall, the pattern of returns to skills also intensifies our findings for generative AI: math skills become even more valuable, while the returns to all other skill dimensions decline further.

FIGURE 6: AI'S IMPACT WITH VS. WITHOUT PHYSICAL CAPABILITIES

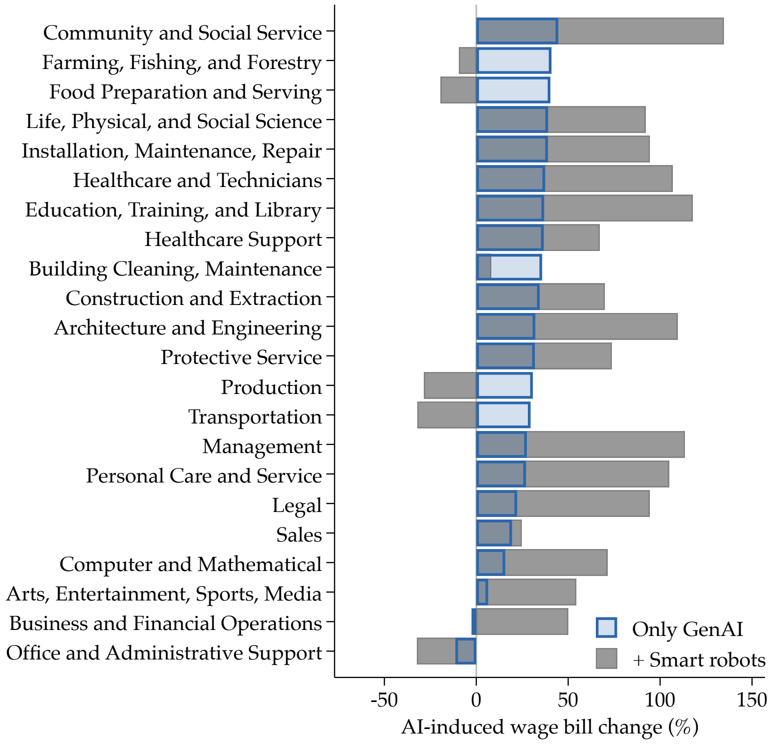

Notes: This figure shows the model's predictions on AI's wage bill effects by occupational group under two scenarios: one with only generative AI and one where AI systems also possess physical manipulation capabilities ("smart robots"). Occupations are sorted in descending order of generative AI's effect on the total wage bill, so that the first listed occupation experiences the largest wage bill increase.

Beyond skill requirements, what characterizes the occupations with the largest increases in the wage bill? There is a strong positive correlation in the education level typical to an occupation and their wage bill increase (see Figure A.16b). This contrasts sharply with the generative AI scenario, in which education showed little correlation with changes in occupational outcomes.

6 Early Signs of AI's Impact on the Labor Market

In this section, we turn from our model's theoretical predictions to empirical evidence. First, we use recent labor market data from the CPS to test whether the model's predictions on the labor market effects of generative AI are beginning to unfold. Specifically, our event study assesses whether predicted occupational outcomes correlate with observed changes since the release of ChatGPT in 2022. Second, we predict changes in the labor market returns to college majors and implement a similar event study for college major choices using data from the National Student Clearinghouse. Third, we zero in on two occupations frequently discussed in relation to AI's impact to explain and validate the model's predictions.

6.1 Event Study of Aggregate Labor Market Shifts

We implement an event study design using Current Population Survey (CPS) data from 2020 to 2025. Our continuous treatment variable is each occupation's model-predicted change in occupational outcome $\Delta\tilde{Y}_o$ (wage bill share, employment, or wages). The specification assesses whether the model predicts differential occupational trends after ChatGPT's release in late 2022, conditional on occupation fixed effects (α_o), time fixed effects (γ_t), and time fixed effects interacted with occupation-level controls ($\eta_t \cdot X_o$):

$$Y_{o,t} = \sum_{k \neq -1} \beta_k \cdot \mathbb{1}[t = k] \times \Delta\tilde{Y}_o + \alpha_o + \gamma_t + \eta_t \cdot X_o + \epsilon_{o,t}. \tag{22}$$

The coefficients β_k capture differential trends for occupations with higher predicted change, with $\beta_k = 1$ indicating that the full model-predicted effect has materialized by period k.[31]

Figure 7 presents our event study estimates. Occupations predicted to gain importance based on their wage bill's share of the overall economy indeed begin to see a relative increase starting around two years after OpenAI's first release of ChatGPT. This effect gradually increases over time. The magnitude of our estimates suggest that by late 2025, between 5 and 10 percent of the predicted wage bill share gains have materialized.

Appendix Figure A.20 shows results for the wage bill's two components: employment and wages. Employment begins to rise significantly for occupations predicted to gain employment from AI starting around one year after ChatGPT's release. In contrast, we do not observe any meaningful effects on wages, suggesting that the initial adjustment of the labor market occurs mostly through quantities not prices (this finding is consistent with evidence on young workers from Brynjolfsson et al., 2025a).

[31] To reduce sampling noise, we aggregate monthly CPS data into 6-month periods to compute the occupational outcomes $Y_{o,t}$. Our sample includes working-age individuals (18-65) in the labor force.

FIGURE 7: EARLY LABOR MARKET EFFECTS OF GENERATIVE AI

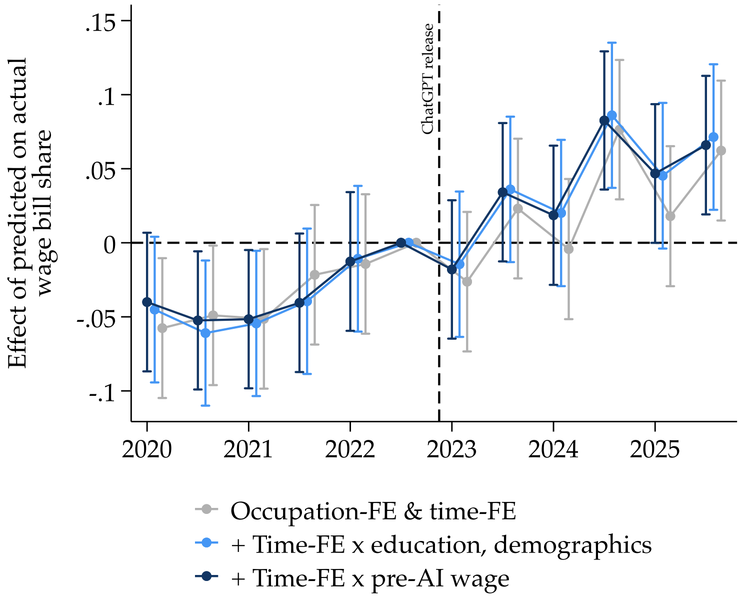

Notes: Event study estimates ($\widehat{\beta}_k$) show differential changes in occupational wage bill shares following ChatGPT's November 2022 release. Coefficients represent the effect of the model-predicted change in the outcome on the observed outcome. A coefficient of 1 would indicate complete realization of model predictions. Estimates use CPS data aggregated to 6-month periods with occupation fixed effects and time fixed effects. The specification with controls includes occupation-specific trends based on education, sectoral composition, demographics, and pre-AI wages. Error bars represent 95%-confidence intervals.

We interpret this evidence as suggestive, particularly given that some employment effects appear to predate 2022. While this timing could indicate that other factors correlated with AI exposure are driving these patterns, an alternative explanation is that occupations most exposed to AI (such as radiologists and telemarketers) had already begun adopting generative AI tools before ChatGPT's public release (see, e.g., Acemoglu et al., 2022).

6.2 Event Study of Shifts in College Major Choices

We next assess how the returns to majors change in the era of AI and whether college students have already begun adjusting their major choices. Major choices offer useful insights on early responses to technical change as young, higher-educated Americans are among the earliest AI adopters (Bick et al., 2024) and major choices can adjust quickly. In contrast, labor market adjustments require occupational switching, retraining, and equilibrium price adjustments.

To construct predictions for the changes in returns to each major, we combine our model-based estimates on AI-induced changes in the labor market returns to skills with skill intensity measures from the Course-Skill Atlas, which maps majors to ONET tasks via information from course syllabuses (Javadian Sabet et al., 2024). We link their major-task mapping to our task-skill framework, yielding skill intensities across our five dimensions for each major. As the main outcome, we use major enrollment from the National Student Clearinghouse, covering 96 percent of all US post-secondary students.

Majors intensive in skills that retain value under AI—particularly math and manual skills—see the largest predicted gains (see Appendix Figure A.17). Atmospheric sciences, astronomy, earth sciences, chemistry, and engineering rank among the top winners with returns rising over 2 percentage points more than for the average major. In contrast, majors intensive in verbal skills, where our model predicts the largest decline in returns, fare worst. French, theology, and Hebrew appear at the bottom of the distribution.

FIGURE 8: EARLY COLLEGE MAJOR EFFECTS OF GENERATIVE AI

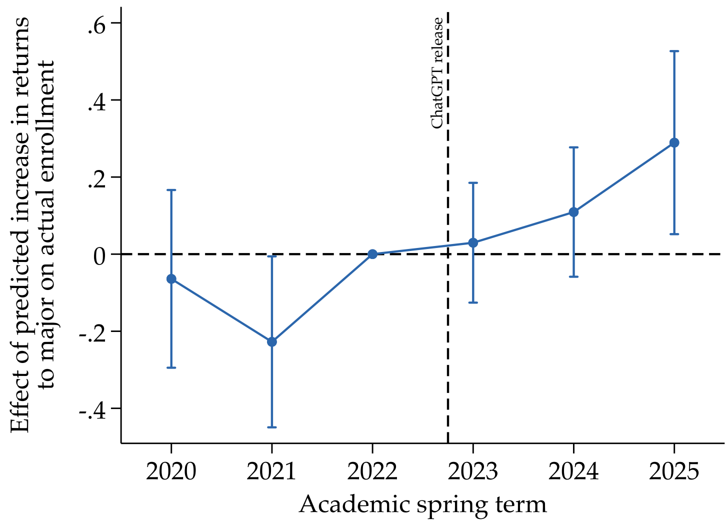

Notes: This figure shows event study estimates ($\widehat{\beta}_k$), corresponding to differential changes in college major enrollment (in logs) following ChatGPT's November 2022 release. Coefficients represent the effect of the model-predicted change in returns to each major on observed log enrollment. A coefficient of 1 would indicate that a one percentage point increase in a major's returns (relative to the average major) is associated with a doubling in enrollment. Estimates use National Student Clearinghouse data with major fixed effects and time fixed effects. Error bars represent 95%-confidence intervals clustered at the major level.

Using our predictions of each major's changing return, we implement an event study design paralleling our analysis of labor market outcomes. The specification follows equation (22), with college major enrollment as the outcome and predicted changes in returns to each major as the continuous treatment variable.

Following ChatGPT's first release, a clear positive relationship emerges: students increasingly enroll in majors our model predicts will benefit from AI-induced changes in skill returns (see Figure 8). By the spring of 2025, a one percentage point increase in a major's returns (a large change) was associated with a 30 percent increase in enrollment. Prior to ChatGPT's release, the coefficients fluctuate around zero with no discernible trend, suggesting that majors predicted to benefit from AI were not already experiencing differential enrollment growth. This evidence suggests that students are already reallocating toward fields better positioned for an AI-transformed labor market.

6.3 Case studies

Radiologists. In 2016, deep learning pioneer Geoffrey Hinton warned to "stop training radiologists," because AI would render them obsolete within five years. Since then, radiology has indeed accounted for more than 75 percent of all FDA-authorized clinical AI tools, and roughly two-thirds of US radiology departments report using AI (Mousa, 2025). However, the labor market for radiologists has grown: the wage bill share increased by 6.6 percent between 2016 and 2024. This increase is driven by strong employment growth (23.2% versus average of 9.8%) and attenuated by below-average wage growth (30.1% versus average of 36.9%).

Our model's predictions line up with these observed labor market patterns. We predict a 42 percent wage bill increase, 1.75 times larger than the average (see Appendix Figure A.12). In line with observed occupational changes, the model-predicted increase in the wage bill from AI is driven by above-average employment growth (28% versus average of 0% by construction) and below-average wage growth (11% versus average of 21%; see Appendix Figure A.13).[32]

Simplification is the key to understanding these outcomes. The occupation experiences strong simplification, pushing employment up and relative wages down. Automation, which affects other occupations far more, increases radiologists' employment further (small effects on wages); augmentation generally

[32]The model outcomes are for *Health Diagnosing and Treating Practitioners*, the 3-digit occupational group that includes radiologists.

does not strongly affect employment or relative wages.

Management Analysts. We predict that the importance of management analysts (management consultants) will shrink. The model predicts a 4 percent decline in employment and no change in wages—a wage effect well below the average of 21 percent (see Appendix Figure A.13).[33]

The reason for this prediction is that management analysts are strongly exposed to both simplification and automation. Simplification pushes average wages down and employment up. However, automation offsets the positive effects on employment, so that the occupation ends up being negatively affected in both dimensions.

The pattern of simplification in management consultancy aligns with experimental evidence. Dell'Acqua et al. (2023) find that lower-skilled consultants experience larger productivity gains from AI. Through the lens of our model, this can only be explained through (direct or indirect) simplification: the reduction in skill requirements increases the relative productivity of less skilled workers.

Telemarketers. Telemarketers represent one of the clearest cases where AI substitutes for, rather than complements, human labor. Their work consists of 12 distinct tasks, all of which can be automated by generative AI. The broader occupational group, *Other Sales and Related Workers*, also faces high automation exposure, raising wages and reducing employment, and very little simplification (which would otherwise oppose automation's positive wage effects). Indeed, our model predicts that this occupational group experiences increasing wages but ranks among the 5 percent of most negatively affected groups in terms of both employment and total wage bill (see Appendix Figures A.12 and A.13).

7 Conclusion

Technological 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

[33] The model outcomes are for *Business Operations Specialists*, the 3-digit occupational group that includes management analysts.

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 generative AI substantially raises wages, especially in the lower part of the wage distribution. A decomposition shows that simplification of tasks is the key driver behind AI's distributional effects.

This paper raises several important questions for future research. First, 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). Second, we treat workers' skills at labor market entry as exogenous. It is useful, however, to consider how technical change may affect people's educational choices (Heckman et al., 1998). Last, this paper only considers the effect of technical change on the labor market. However, technical change can also have strong distributional implications through capital income (Moll et al., 2022) and business income (Reichardt, 2025). For the latter, it is particularly pressing to understand whether AI's simplifying capabilities allow specifically small firms to benefit from its use.

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., D. AUTOR, J. HAZELL, AND P. RESTREPO (2022): "Artificial intelligence and jobs: Evidence from online vacancies," *Journal of Labor Economics*, 40, S293–S340.

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.

——— (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.

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., F. LEVY, AND R. J. MURNANE (2003): "The skill content of recent technological change: An empirical exploration," *Quarterly Journal of Economics*, 118, 1279–1333.

BALEY, I., A. FIGUEIREDO, AND R. ULBRICHT (2022): "Mismatch cycles," *Journal of Political Economy*, 130, 2943–2984.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

BICK, A., A. BLANDIN, AND D. J. DEMING (2024): "The Rapid Adoption of Generative AI," Working paper.

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., B. CHANDAR, AND R. CHEN (2025a): "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Working Paper.

BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025b): "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.

CAUNEDO, J., D. JAUME, AND E. KELLER (2023): "Occupational exposure to capital-embodied technical change," *American Economic Review*, 113, 1642–1685.

CONLON, C. AND J. GORTMAKER (2020): "Best practices for differentiated products demand estimation with pyblp," *The RAND Journal of Economics*, 51, 1108–1161.

DANIELI, O. (2025): "Skill-Replacing Technology and Bottom-Half Inequality," Tech. rep., working paper, Tel Aviv University.

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," Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013, The Wharton School Research 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.

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.

ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2024): "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 384, 1306–1308.

FREUND, L. B. AND L. F. MANN (2025): "Job Transformation, Specialization, and the Labor Market Effects of AI," Working paper.

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.

HOSNY, A., C. PARMAR, J. QUACKENBUSH, L. H. SCHWARTZ, AND H. J. W. L. AERTS (2018): "Artificial intelligence in radiology," *Nature Reviews Cancer*, 18, 500–510.

HUMLUM, A. (2019): "Robot adoption and labor market dynamics," *Princeton University*.

HURST, E., Y. RUBINSTEIN, AND K. SHIMIZU (2024): "Task-based discrimination," *American Economic Review*, 114, 1723–1768.

JAVADIAN SABET, A., S. BANA, AND R. YU (2024): "Course-Skill Atlas: a National Longitudinal Dataset of Skills Taught in U.S. Higher Education Curricula," dataset: `https://doi.org/10.1038/s41597-024-03931-8`.

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.

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).

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.

REYNAERTS, J., R. VARADHA, AND J. C. NASH (2012): "Enhancing the convergence properties of the BLP (1995) contraction mapping," 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.

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.

WOESSMANN, L. (2024): "Skills and earnings: A multidimensional perspective on human capital," *Annual Review of Economics*, 17.

ZEIRA, J. (1998): "Workers, machines, and economic growth," *Quarterly Journal of Economics*, 113, 1091–1117.

Appendix

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

(A) PRODUCTION (B) SKILL ACCUMULATION

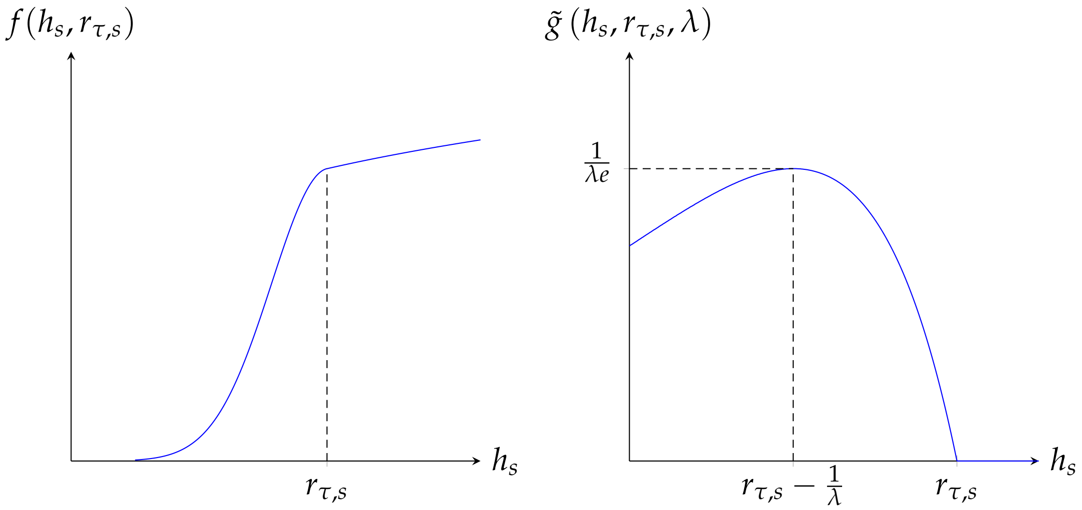

Notes: This figure illustrates the functional forms of the production and skill accumulation functions in equations (14) and (16), respectively. Panel A shows the production function $f(h_s, r_{\tau,s}) = h_s^{\omega} \exp\left(-\eta \min\{h_s - r_{\tau,s}, 0\}^2\right)$. Panel B shows the learning part 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.

A Figures

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

(A) MANUAL

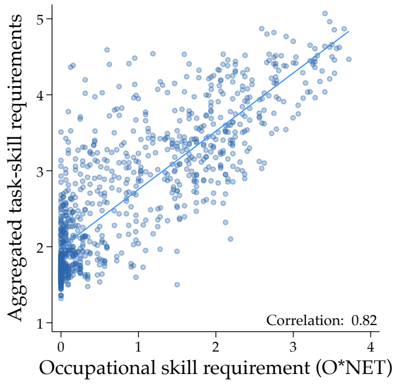

(B) MATH

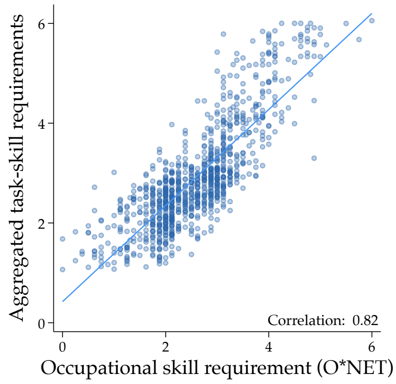

(C) SOCIAL

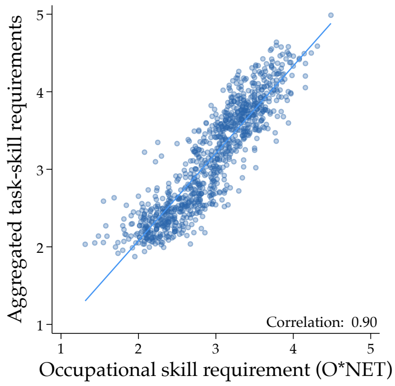

(D) TECHNICAL

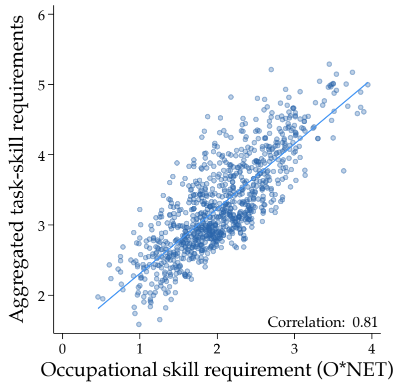

(E) VERBAL

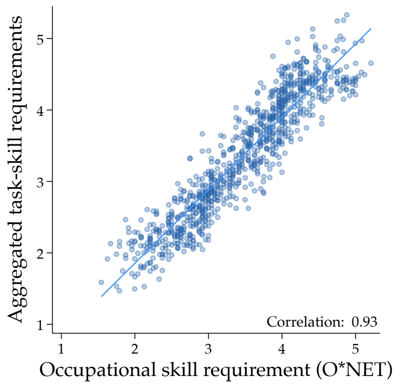

Notes: This figure shows the correlation between the occupation-level skill requirement in the O*NET database and the GPT-4o generated task-level skill requirements aggregated to the occupation-level for the skills used in the analysis. Each observation represents an occupation in the O*NET database.

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

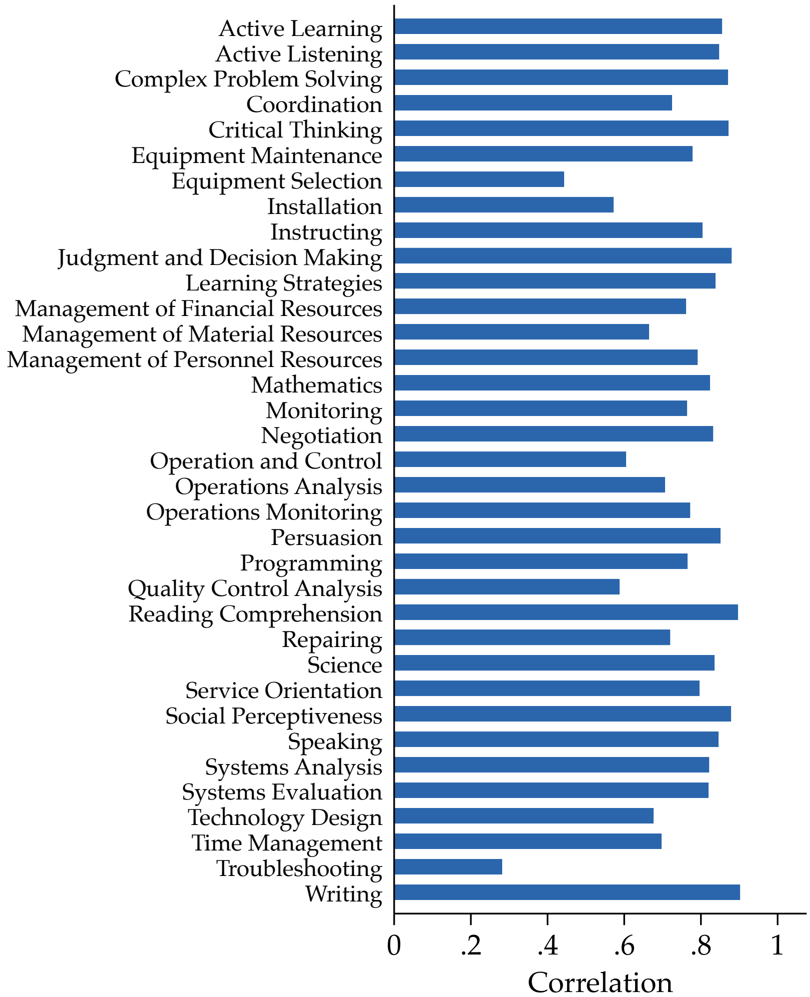

Notes: This figure shows the correlations between the occupation-level skill requirement in the O*NET database and the GPT-4o generated task-level skill requirements aggregated to the occupation-level for each of the 35 skills.

FIGURE A.4: AGREEMENT ON AI EXPOSURE WITH ELOUNDOU ET AL. (2024)

(A) AUGMENTATION

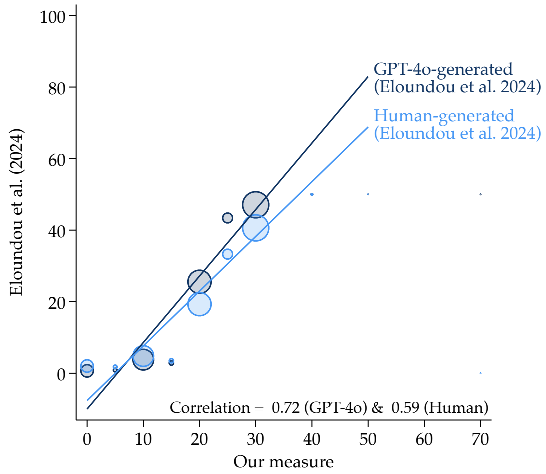

(B) AUTOMATION

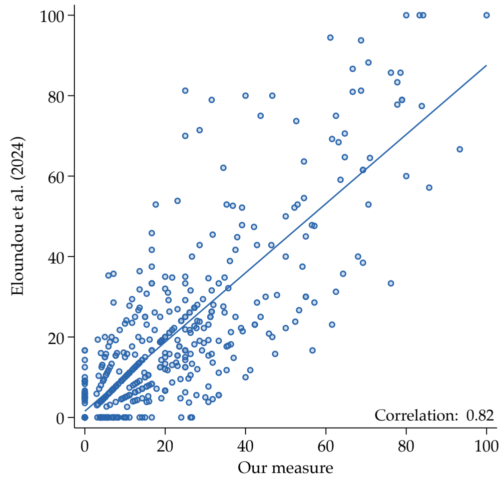

Notes: This figure compares our estimates tasks' exposure to augmentation and automation by generative AI. Panel A shows our task-level augmentation estimates (time saved to complete each task in an occupation) to those provided by Eloundou et al. (2024), measuring whether or not large language models can save at least 50 percent of time to complete a task (binary). Panel B shows the the share of tasks in each occupation that can be automated by generative AI with similar data provided by Eloundou et al. (2024).

FIGURE A.5: SKILLS & GENERATIVE AI EXPOSURE BY CHANNEL

(A) MANUAL

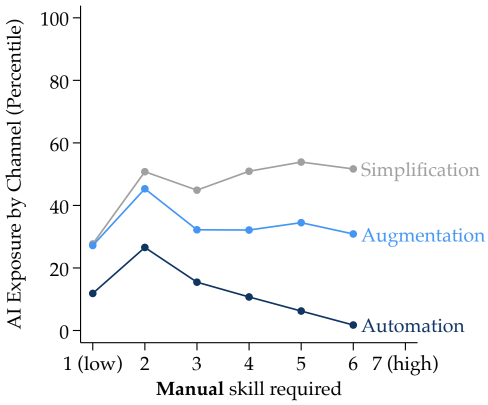

(B) MATH

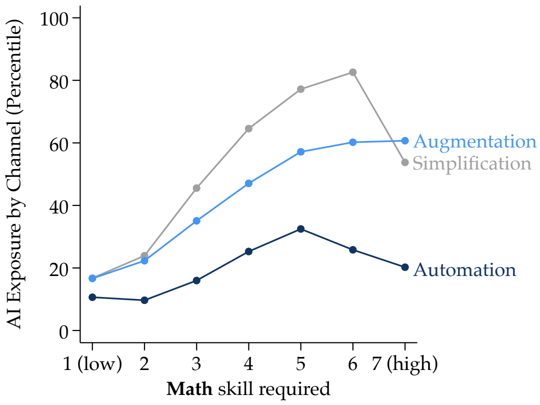

(C) SOCIAL

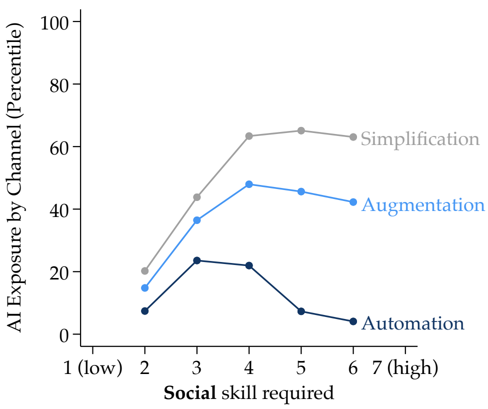

(D) TECHNICAL

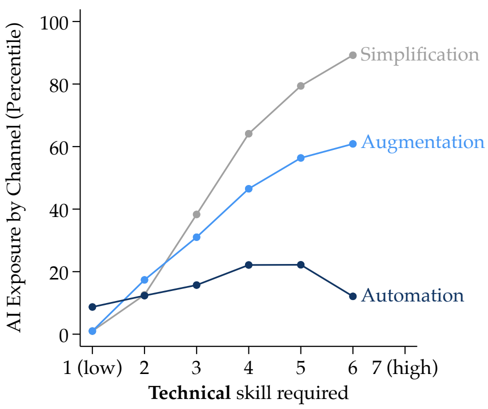

(E) VERBAL

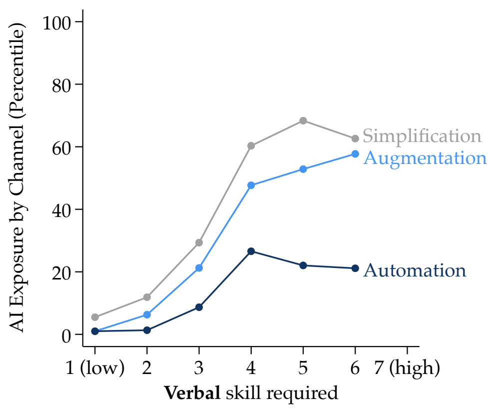

Notes: This figure shows the correlation between a task's skill requirements and its potential to be augmented, automated, or simplified by Generative AI. Each dot represents the average percentile of exposure to each channel among tasks with the same requirement in a given skill.

FIGURE A.6: INITIAL SKILL DISTRIBUTION

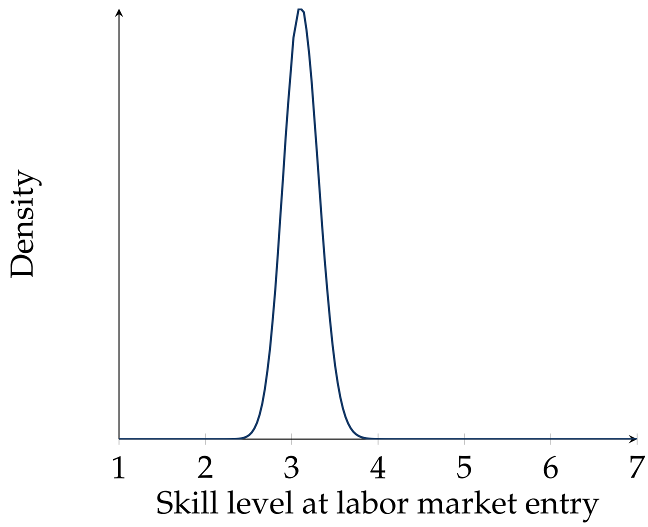

Notes: This figure shows the estimated density of the skill distribution of young workers (age $a = 1$ in the model) on O*NET's 1 to 7 scale. For comparison, for the skill "reading comprehension", a 2 means being able to "read step-by-step instructions for completing a form", 4 means being able to "understand 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

(A) OCCUPATIONAL WAGES

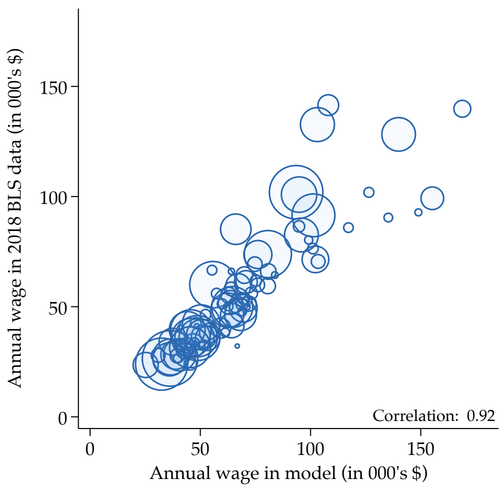

(B) LIFE-CYCLE OF WAGES

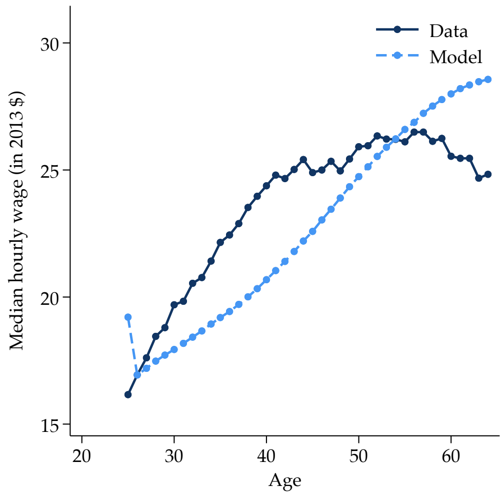

Notes: Panel A shows the correlation between the average wage in an occupation in the model's steady state and in the data as reported in the 2018 BLS OEWS data. Panel B reports the median wage by age in the NLSY79 and the model's steady state.

FIGURE A.8: WAGE EFFECTS ACROSS THE DISTRIBUTION FOR ALL CHANNELS

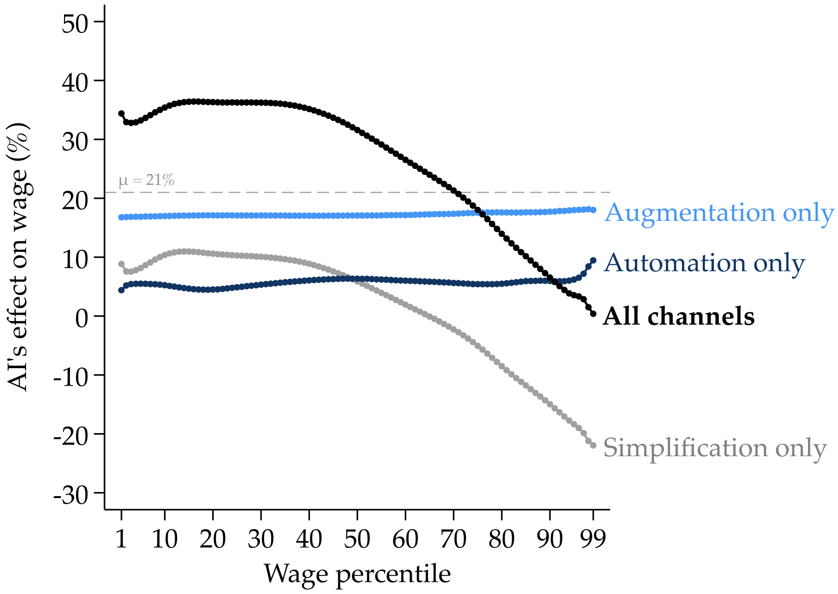

Notes: This figure shows the distribution of wage changes induced by generative AI across the wage percentile distribution. The horizontal axis represents wage percentiles weighted by pre-AI employment, and the vertical axis shows the percentage change in wages for each percentile. The black line shows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

FIGURE A.9: DISTRIBUTION OF WELFARE EFFECTS

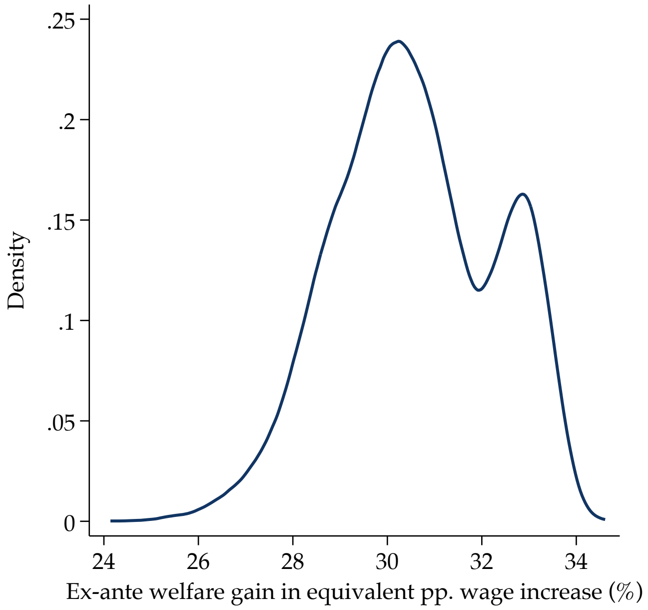

Notes: This figure shows the distribution of AI's welfare effect on individual workers. The welfare effect is measured in equivalent wage variation: it represents the permanent wage increase across all occupations that yields the same welfare gain as the introduction of AI.

FIGURE A.10: GENERATIVE AI'S EFFECT ACROSS OCCUPATIONS

(A) WAGE BILL (B) WAGES (C) EMPLOYMENT

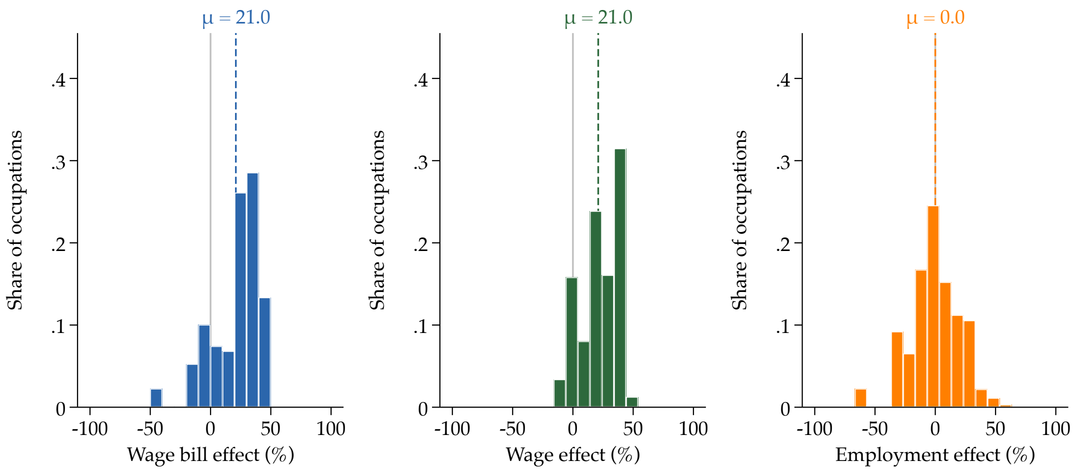

Notes: This figure shows the distribution of generative AI's predicted effects across occupations based on our structural model. Panel (A) shows wage bill changes (wages × employment). Panel (B) shows wage changes. Panel (C) shows employment effects, which are symmetric around zero by definition as our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

FIGURE A.11: AI'S EFFECT ON OCCUPATIONAL EMPLOYMENT & WAGES

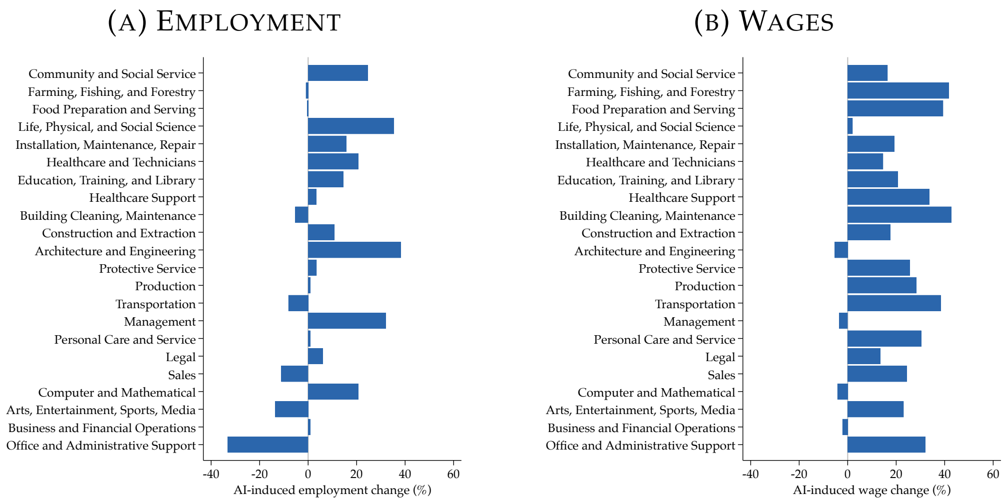

Notes: This figure shows the model's predictions on AI's employment and wage effects by occupational group. Occupations are sorted in descending order of AI's effect on their wage bill, so that the first listed occupation experiences the largest wage bill increase.

FIGURE A.12: AI'S EFFECT ON DETAILED OCCUPATIONS' WAGE BILLS

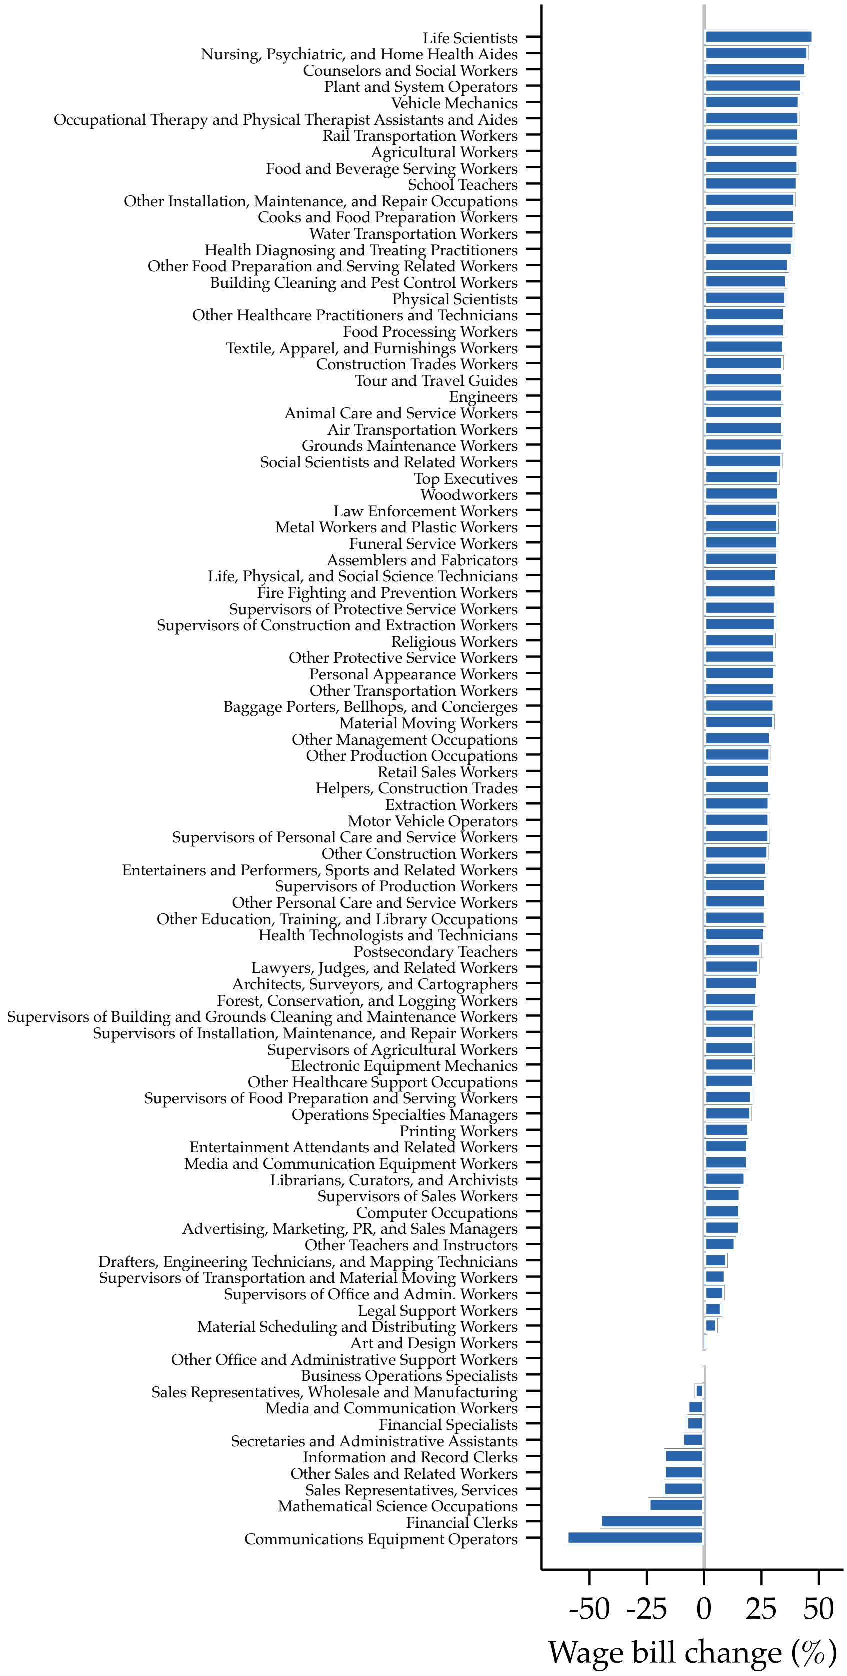

Notes: This figure shows the model predictions on AI's wage bill effects by occupation.

FIGURE A.13: AI'S EFFECT ON DETAILED OCCUPATIONS' EMPLOYMENT & WAGES

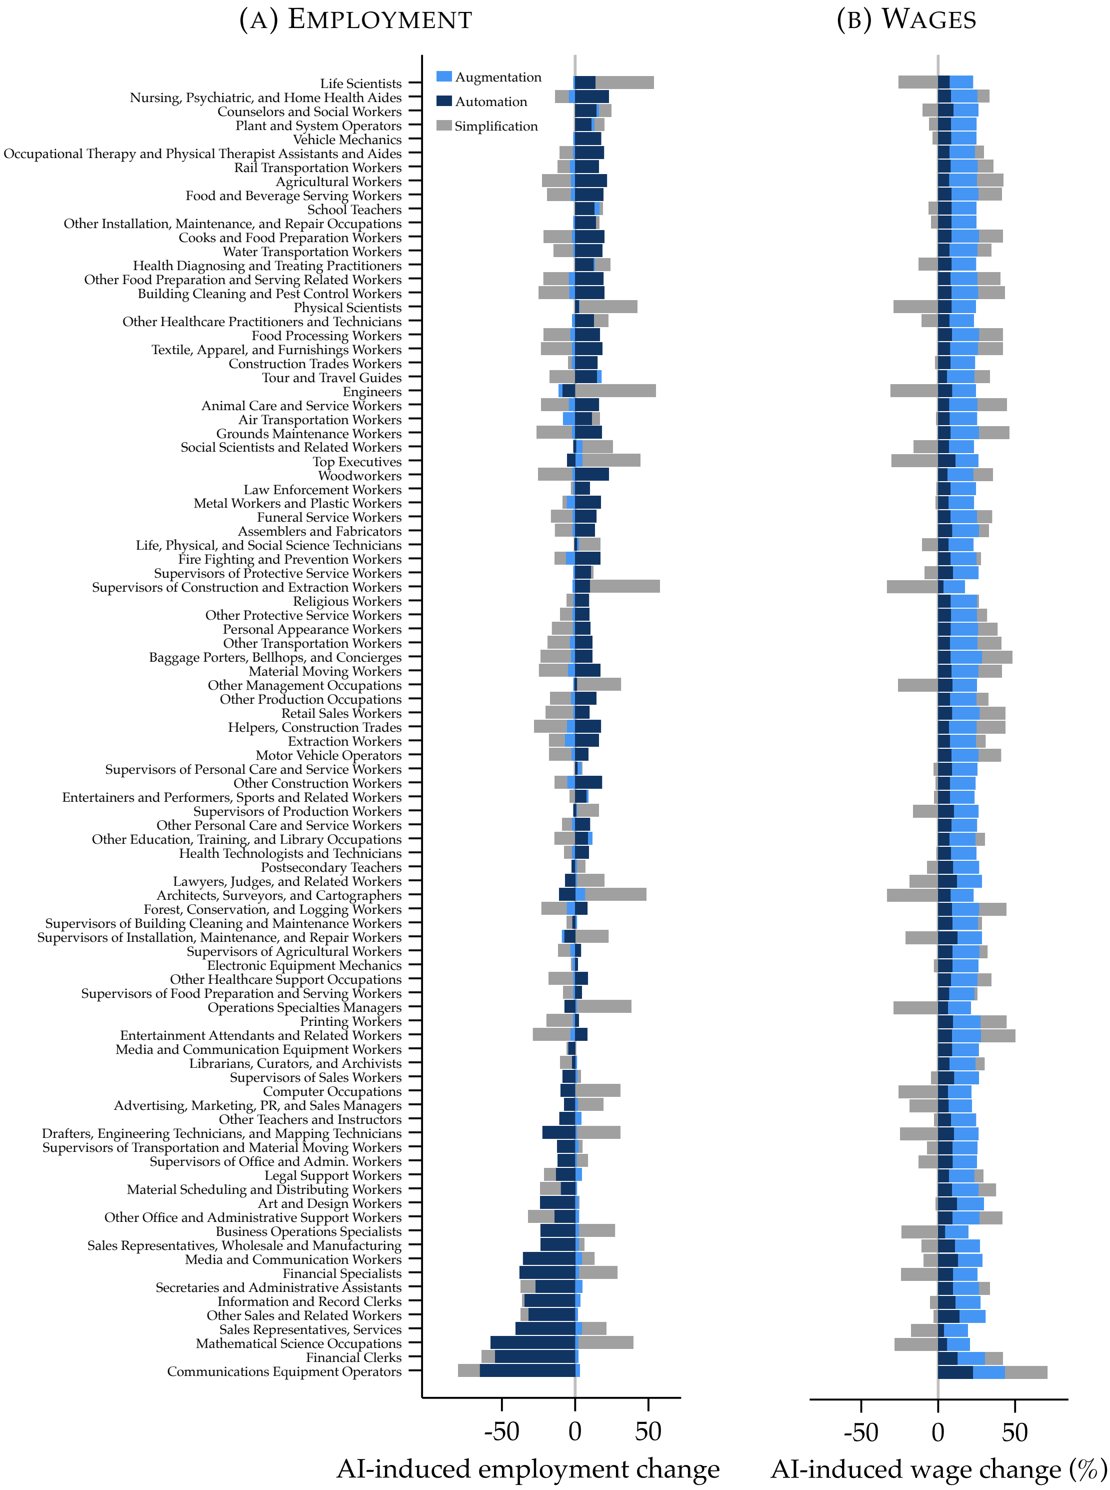

Notes: This figure shows the model's predictions on AI's employment and wage effects by occupation. Occupations are sorted in descending order of AI's effect on the total wage bill, so that the first listed occupation experiences the largest wage bill increase. We conduct a Shapley-Owen decomposition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

Figure A.14: Automation exposure most predictive of labor market losses

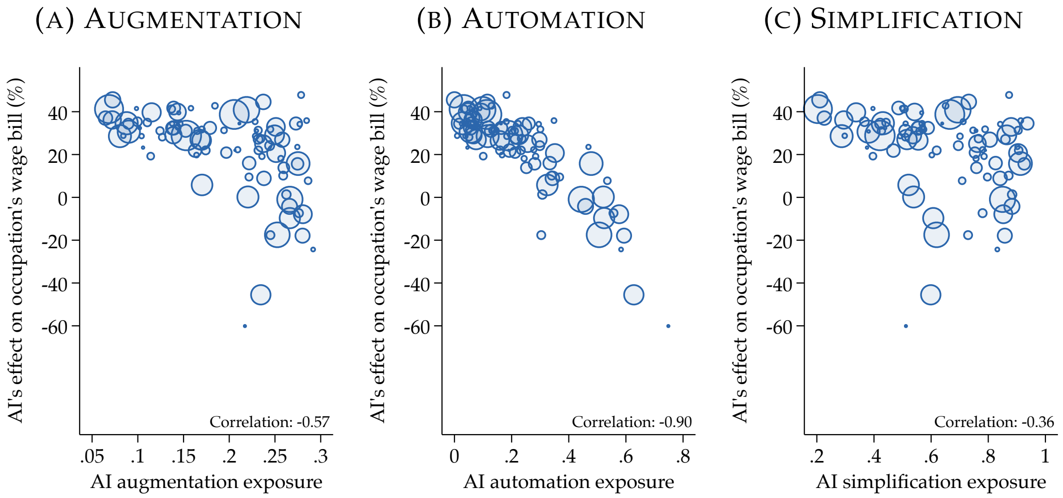

Notes: This figure shows the relationship between three dimensions of AI exposure and model-predicted changes in occupational wage bills. Each bubble represents an occupation, with size proportional to pre-AI employment. Panel (A) shows augmentation exposure, measured as the share of time access to generative AI can save in completing an occupation's tasks. Panel (B) shows automation exposure, measured as the share of an occupation's tasks than generative AI can complete autonomously. Panel (C) shows simplification exposure, measured as the (negative) relative change of skill levels required to complete an occupation's tasks (averaged across all 35 O*NET skills).

Figure A.15: Skills and education by GenAI's wage bill effect

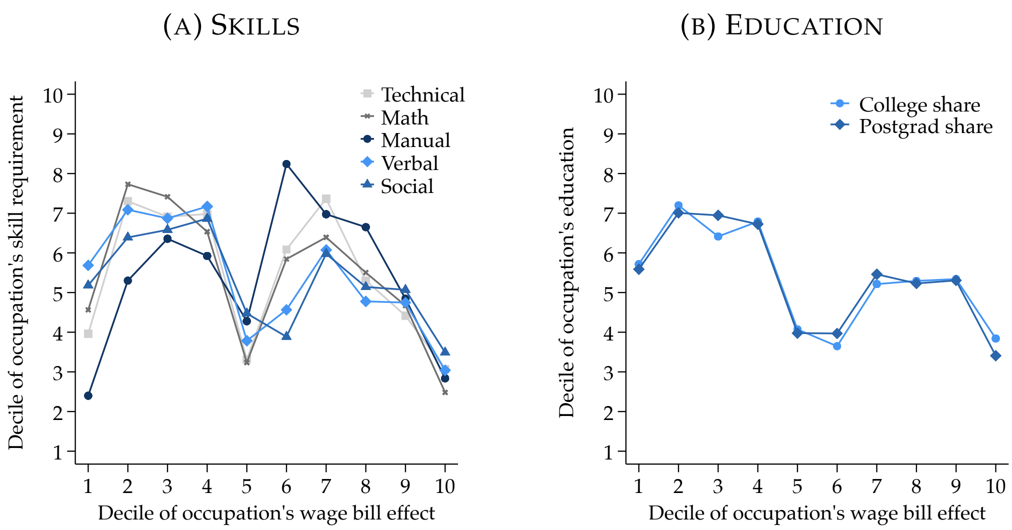

Notes: This figure shows the relationship between occupational characteristics and generative AI's wage bill effects. Panel A plots average skill requirement deciles against wage bill effect deciles. Panel B plots education levels against wage bill effect deciles. Each point represents a decile of occupations ranked by their predicted wage bill change, weighted by pre-AI employment.

FIGURE A.16: SKILLS AND EDUCATION BY WAGE BILL EFFECT: PHYSICAL AI SCENARIO

(A) SKILLS (B) EDUCATION

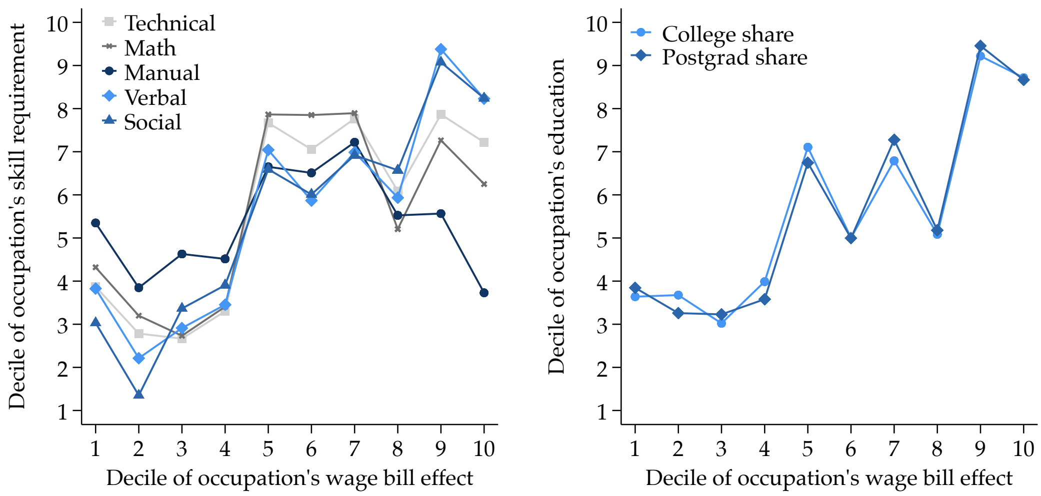

Notes: This figure shows the relationship between occupational characteristics and AI's wage bill effects when physical manipulation capabilities are available. Panel A plots average skill requirement deciles against wage bill effect deciles. Panel B plots education levels against wage bill effect deciles. Each point represents a decile of occupations ranked by their predicted wage bill change, weighted by pre-AI employment.

FIGURE A.17: AI'S EFFECT ON COLLEGE MAJOR ENROLLMENT

Notes: This figure shows the model's predictions on AI's effect on returns to college majors. 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.

FIGURE A.18: RADIOLOGISTS' TASK-BASED AUTOMATION EXPOSURE

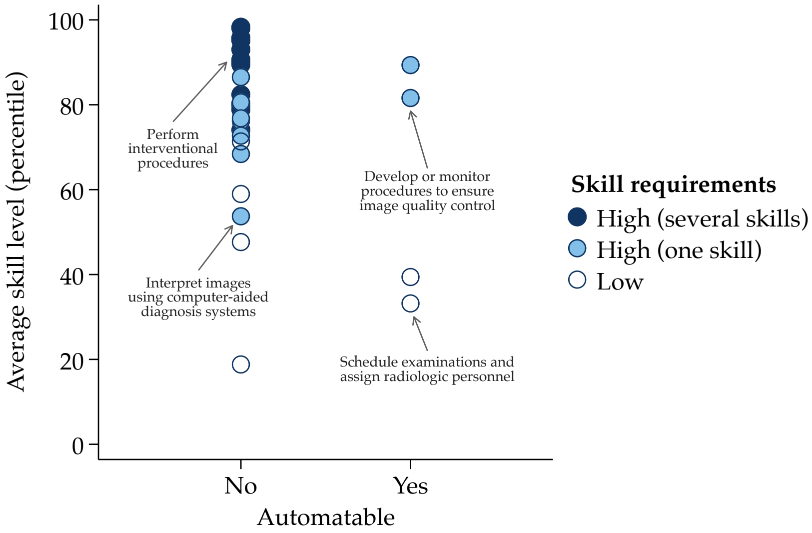

Notes: This figure shows radiologists' automation exposure across all tasks they engage in according to O*NET. We classify a task as 'automatable' if it has high or full automation exposure, which is restricted to cases where the LLM indicates that generative AI can complete at least 90% of the components of the tasks. We classify skill requirements as "high" if they exceed level 3.5 in O*NET's scale from 1 to 7.

FIGURE A.19: MANAGEMENT ANALYSTS' TASK-BASED AUTOMATION EXPOSURE

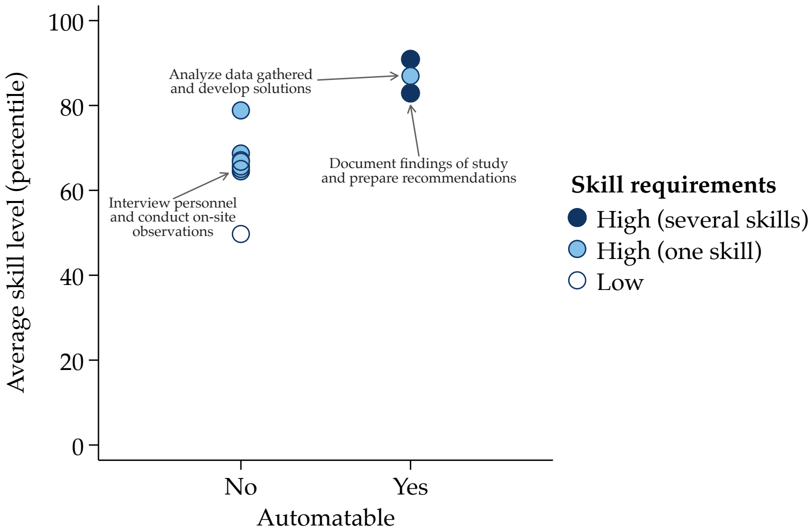

Notes: This figure shows management analysts' automation exposure across all tasks they engage in according to O*NET. We classify a task as 'automatable' if it has high or full automation exposure, which is restricted to cases where the LLM indicates that generative AI can complete at least 90% of the components of the tasks. We classify skill requirements as "high" if they exceed level 3.5 in O*NET's scale from 1 to 7.

FIGURE A.20: EARLY LABOR MARKET EFFECTS OF GENERATIVE AI

(A) EMPLOYMENT

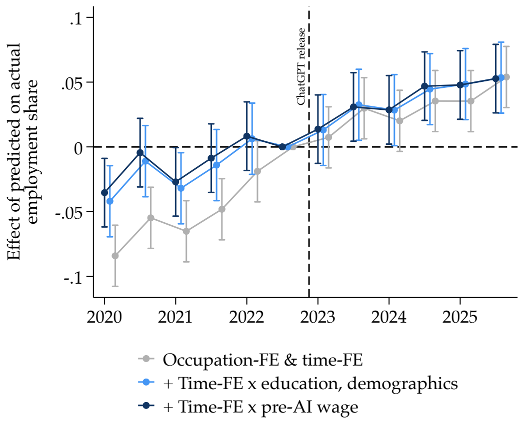

(B) WAGES

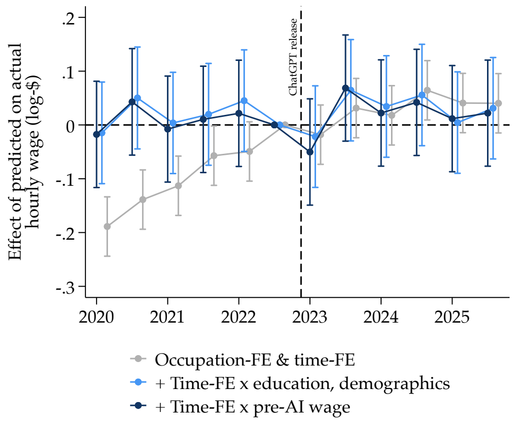

Notes: Event study estimates ($\widehat{\beta}_k$) show differential changes in occupational employment and wages following ChatGPT's November 2022 release. Coefficients represent the effect of the model-predicted change in the outcome on the observed outcome. A coefficient of 1.0 would indicate complete realization of model predictions. Estimates use CPS data aggregated to 6-month periods with occupation fixed effects and time fixed effects. The specification with controls includes occupation-specific trends based on education, sectoral composition, demographics, and pre-AI wages. Error bars represent 95%-confidence intervals.

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: AGREEMENT ON AUTOMATION EXPOSURE WITH ELOUNDOU ET AL. (2024)

	Eloundou et al. (2024)				
Our measure	None	Low	Medium	High	Full
None	26.86	1.65	0.01	0.34	0.14
Low	6.12	21.34	1.98	2.18	0.12
Medium	0.17	7.27	7.61	3.65	0.06
High	0.04	1.77	5.61	11.53	0.18
Full	0.02	0.07	0.02	1.02	0.24

Notes: This table shows the agreement rates between our measure of automation and that of Eloundou et al. (2024) on the task-level. The table is computed based on 14,209 tasks (out of 19,530) that are in both databases. We classify a task as automated if the exposure is "high" or "full". The share of automated tasks is 21% and 22% in Eloundou et al. (2024)'s and our measure, respectively.

TABLE B.3: SUMMARY OF TASK-LEVEL DATA ON AI CAPABILITIES

	Augmentation		Automation	Simplification
	Excluding automatable tasks	Including automatable tasks		
Mean	17.9%	20.2%	22.2%	18.3%
Std. Dev.	9.4%	9.6%	41.6%	6.6%
Median	20.0%	20.0%	0.0%	20.4%
Range	0.0% - 70.0%	0.0% - 70.0%	0.0% - 100.0%	0.0% - 32.0%
Tasks	15,192	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. Augmentation and simplification estimates are generated by GPT-4o; automation estimates are generated by GPT-5 with low to medium reasoning effort.

TABLE B.4: EXPERIMENTAL ESTIMATES COMPARED TO OUR TASK AUGMENTATION DATA

Occupation	Task	Tool	Estimate		N	Notes	Source
			Theirs	Ours			
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%	27%	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%	28%	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 new & low-skill workers; minimal impact on experienced & 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; likelihood of using AI after experiment ↑	[8]
Taxi driver	Selecting routes	AI Navi	14%	9%	520	Shorter cruising time; gains only among low-skill drivers	[9]
Lawyer	Legal writing	Vincent & o1-preview	20%	30%	127	19.9% time saved across different legal writing tasks, quality ↑, LLM "Vincent" slightly higher gains	[10]
Product designer	Product marketing & development	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%	27%	24	Developers implemented ∼65% more requirements with AI assistance	[12]
Software developer	Coding	Google AI Tools	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 365 Copilot	11%	26%	7,137	Treated spent 12% less time on email each week; did not significantly change time spent in meetings.	[15]
Train commissioning technician	Trouble-shooting	GPT-3.5 + RAG	20%	20%	173	+1.14 SD quality score; 20% increase in tasks completed not significant; less-experienced benefit more.	[16]

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). To construct our own estimates of task augmentation (share of time saved to complete task) 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.5: 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 (2019).
Number of occupations	J	93	3-digit BLS SOC occupations.
Number of periods	A	40	Years between 25 and 65.
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	$\boldsymbol{r}_\tau$		Large language model.
Task-level AI augmentation	γ_τ		"
Task-level AI automation	$\mathcal{A}_j$		"
Learning cost: 1st AFQT quartile	$\lambda(1)$	3.50	Maximum likelihood.
Learning cost: 2nd AFQT quartile	$\lambda(2)$	2.97	"
Learning cost: 3rd AFQT quartile	$\lambda(3)$	2.81	"
Learning cost: 4th 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.6: OCCUPATIONAL SKILL REQUIREMENTS PREDICT OCCUPATIONAL PRICES

Skill requirements	Dependent variable: Log occupational price $\hat{p}_j$			
Manual	0.228	0.230	0.226	0.227
	(0.207)	(0.214)	(0.201)	(0.208)
Math	-0.070	-0.073	-0.077	-0.080
	(0.277)	(0.286)	(0.269)	(0.278)
Social	-0.083	-0.078	-0.077	-0.073
	(0.352)	(0.363)	(0.341)	(0.353)
Technical	1.342**	1.341**	1.353**	1.352**
	(0.599)	(0.619)	(0.581)	(0.601)
Verbal	0.554	0.554	0.550	0.550
	(0.376)	(0.388)	(0.364)	(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 $\frac{p90}{p10}$	$\frac{p90}{p50}$	$\frac{p75}{p25}$	Top shares 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 data (NLSY79) 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.

TABLE B.8: SKILL CHANNELS AND LABOR MARKET OUTCOMES

	Wage growth (%)		Employment growth (%)	
	(1)	(2)	(3)	(4)
Augmentation	1.48	-2.89	2.11	6.79**
	(4.79)	(4.18)	(3.18)	(2.91)
Automation	4.74**	19.04***	-23.86***	-39.03***
	(1.80)	(3.94)	(1.39)	(3.67)
Simplification	-24.68***	-15.19***	24.81***	15.25***
	(4.25)	(3.38)	(2.93)	(2.94)
Indirect simplification				
Augmentation-led		3.52*		-2.50**
		(1.78)		(1.24)
Automation-led		-15.88***		17.05***
		(4.34)		(4.22)
Observations	93	93	93	93

Notes: This table presents weighted OLS regressions of occupation-level wage and employment growth on AI exposure measures and skill requirement changes. The dependent variable is wage growth in columns (1)–(2) and employment growth in columns (3)–(4). All independent variables are standardized to have standard deviation 1. Columns (1) and (3) include only direct AI exposure measures (augmentation, automation, simplification). Columns (2) and (4) add indirect skill requirement changes induced by AI. All regressions are weighted by pre-AI occupation employment shares. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Estimation

C.1 Cost savings and AI's income share

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

$$c^l_\tau(\boldsymbol{h}) \equiv \frac{\Lambda_j(\boldsymbol{h})}{\gamma_\tau f(\boldsymbol{h}, \boldsymbol{r}_\tau)}$$

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

$$c^k_\tau \equiv \frac{R}{\phi_\tau}.$$

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

$$\chi_\tau \equiv \frac{c^k_\tau}{c^l_\tau(\boldsymbol{h})} = c^k_\tau \cdot \frac{\gamma_\tau f(\boldsymbol{h}, \boldsymbol{r}_\tau)}{\Lambda_j(\boldsymbol{h})}.$$

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

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

For our quantification, we need an estimate of $\sum_{\tau\in\mathcal{A}_j}\theta_{j,\tau}\left(\frac{c^k_\tau}{p_j}\right)^{1-\rho}$ 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 $\forall j = 1, \ldots, J$. Second, we assume that the automated tasks are not different in productivity and skill requirements from the non-automatable tasks, i.e., $\gamma_\tau^{\rho-1} f\left(\boldsymbol{h}, \boldsymbol{r}_\tau\right)^{\rho-1} \approx \sum_{\tau\in\mathcal{N}_j}\theta_{j,\tau}\gamma_\tau^{\rho-1} f(\boldsymbol{h}, \boldsymbol{r}_\tau)^{\rho-1}$ for all $\tau \in \mathcal{A}_j$ and $\forall j = 1, \ldots, J$. Under those two assumptions, the cost savings simplify to

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

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.

C.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 $\mathcal{A}_j = \varnothing$ (so that $\Gamma_j = 1$ and $\mathcal{N}_j = \mathcal{T}_j$) for all $j = 1, \ldots, J$, we obtain

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

Now define the variable $m_\tau \equiv \sum_{s \in S} \min\{h_s - r_{\tau,s}, 0\}^2$ 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).

D Data

D.1 Task-level Skill Requirements

We elicit a task's skill requirements by replicating O*NET's occupation-level questionnaire on the task-level using OpenAI's GPT-4o. We requested 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. Below, we present the full text of the prompt for the skill *Reading comprehension*, the occupation *Chief Executives*, and the task *"Prepare budgets for approval, including those for funding or implementation of programs."*

> *The occupation [Chief Executives] contains the task: [Prepare budgets for approval, including those for funding or implementation of programs].*
>
> *What level of skill in [reading comprehension] is needed to perform the task in this occupation well?*
>
> *Provide the answers on a scale from 1 to 7, where 2 means [Read step-by-step instructions for completing a form], 4 means [Understand an email from management describing new personnel policies], and 6 means [Read a scientific journal article describing surgical procedures].*
>
> *Output only a single integer, valued between 1 and 7. Do not output anything else.*

D.2 AI and Task Augmentation, Automation, Simplification

We model technologies' 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 OpenAI's large language models. In our baseline scenario, we only consider Generative AI. However, we also consider automation by Autonomous Vehicles, and Smart Robots. For automation assessments, we use GPT-5; for augmentation and simplification channels, we use GPT-4o (completed before GPT-5's release).

D.2.1 Augmentation of Tasks

For augmentation assessment, we ask GPT-4o to estimate time savings when workers get access to these technologies. We assess all 19,530 O*NET tasks and the technologies "generative AI," "smart robots," and "autonomous vehicles," resulting in 58,590 prompts that are evaluated independently. We use OpenAI's GPT-4o with temperature between 0.05 and 0.1.

The prompt structure is consistent across technologies, varying only in the technology description. Here we show the full prompt for **Generative AI**:

> *We are conducting a rigorous assessment of the time a worker can save on specific tasks by using Generative AI.*
>
> *1. Description of technology: Generative AI*
>
> *Generative AI refers to AI techniques that learn a representation of artifacts from data, and use it to generate brand-new, unique artifacts that resemble but don't repeat the original data. These artifacts can serve benign or nefarious purposes. Generative AI can produce totally novel content (including text, images, video, audio, structures), computer code, synthetic data, workflows and models of physical objects. Generative AI also can be used in art, drug discovery or material design.*
>
> *2. Description of the worker and the task:*
>
> *Worker's role: [Occupation] with an average level of expertise.*
>
> *Worker's access to tools: Has all the standard tools available to someone in this position. In addition, this worker now gains access to a Generative AI.*
>
> *Worker's task: [Occupational task]*
>
> *3. Question:*
>
> *Estimate the percentage of time that the worker can save by using the described Generative AI to assist with the task.*
>
> *4. Output Format:*
>
> *Provide your answer as a percentage (numeric value between 0 and 100). Do not output an explanation or any additional information. The answer should be a single number representing the estimate.*

For **Smart Robots**, the technology description changes to:

A smart robot is an AI-powered, often-mobile machine designed to autonomously execute one or more physical tasks. These tasks may rely on, or generate, machine learning, which can be incorporated into future activities or support unprecedented conditions. Smart robots can be split into different types based on the tasks/use cases, such as personal, logistics and industrial.

For **Autonomous Vehicles**, the description is:

An autonomous vehicle is one that can drive itself from a starting point to a predetermined destination in "autopilot" mode using various in-vehicle technologies and sensors, including adaptive cruise control, active steering (steer by wire), anti-lock braking systems (brake by wire), GPS navigation technology, lasers and radar.

D.2.2 Automation of Tasks

To measure technologies' 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 assess all 19,530 O*NET tasks and the technologies "generative AI", "smart robots", and "autonomous vehicles", resulting in 58,590 prompts that are evaluated independently. We use OpenAI's GPT-5 with low to medium reasoning effort and temperature between 0.05 and 0.1.

The automation prompt follows Eloundou et al. (2024)'s format, with technology-specific definitions and examples. The prompt for **Generative AI** is:

T Automation Rubric

1. Determine if the occupation/task pair meets the definition of T0 No-Automation Exposure. If it does, label it as T0 and stop.

2. If the occupation/task pair does not meet the definition of T0 No-Automation Exposure, determine if the occupation/task pair meets one of the other definitions and select the label that applies:

- *T4: Full automation exposure*
- *T3: High automation exposure*
- *T2: Moderate automation exposure*

- *T1: Low automation exposure*

Rubric

Generative AI refers to AI techniques that learn a representation of artifacts from data, and use it to generate brand-new, unique artifacts that resemble but don't repeat the original data. These artifacts can serve benign or nefarious purposes. Generative AI can produce totally novel content (including text, images, video, audio, structures), computer code, synthetic data, workflows and models of physical objects. Generative AI also can be used in art, drug discovery or material design.

Assume you are a worker with an average level of expertise in your role trying to complete the given task. You have access to Generative AI as well as any other existing software or computer hardware tools mentioned in the task. You also have access to any commonly available technical tools accessible via a laptop (e.g. a microphone, speakers, etc.). You do not have access to any other physical tools or materials.

Please label the given task according to the rubric below.

T0 No-Automation Exposure *A class of tasks for which Generative AI cannot conceivably perform any aspect of the task in any manner.*

T4 Full Automation Exposure *A class of tasks where, in most contexts in which this task is currently performed by a human, Generative AI can complete all aspects of this task with high quality when prompted by a human. The output does not normally require oversight by a human. Oversight is not normally required for tasks labeled T4 because the consequences for failure or inaccuracy are small for this task, human judgment is not necessary to complete this task, and generative models can consistently perform this task with very high quality.*

T3 High Automation Exposure *A class of tasks where, in most contexts in which this task is currently performed by a human, Generative AI could complete 90-100% of the components of the task when prompted, but the output requires oversight from a human. Oversight is normally required because the consequences for failure or inaccuracy are significant for this task, human judgment is necessary to complete this task, and/or generative models cannot perform all aspects of this task with high quality consistently. These tasks rely almost exclusively on the processing of digital* *information, but human judgment is needed to ensure that any digital outputs from Generative AI are high enough quality to be acceptable for the particular context.*

T2 Moderate Automation Exposure *A class of tasks where, in most contexts in which this task is performed by a human, Generative AI could complete between 50%-90% of the components of the task at high quality. These tasks normally rely heavily on the processing of digital information, but a significant portion of the task also involves actions that Generative AI cannot perform with high quality. These tasks require at least some human action beyond just double-checking generative model outputs (such as interpretation, judgment, human-to-human communication, or physical actions).*

T1 Low Automation Exposure *A class of tasks where, in most contexts in which this task is performed by a human, Generative AI could complete between 0%-50% of the components of the task at high quality. These tasks normally rely only partially on the processing of digital information, while the majority of the task involves actions that Generative AI cannot perform with high quality. A majority of the actions that need to be taken to complete this task require a human to perform the action.*

Definitions

High quality *means someone receiving or reviewing the output would not be able to tell the difference between whether it came from Generative AI or a human. For tasks that require a lot of interaction during the completion of the task (e.g. meetings, negotiations), high quality means the people you were interacting with either would not know or would not care that they were interacting with Generative AI.*

Digital information *or* ***information that can easily be expressed digitally*** *includes but is not limited to text, audio, images, video, PDFs, books, code, and data.*

Annotation Examples

Occupation: Special Education Teachers, Preschool

Task: Develop individual educational plans (IEPs) designed to promote students' educational, physical, or social development.

Automation score (T0/T1/T2/T3/T4): T1

Explanation: GenAI drafts SMART goals, accommodations, progress-monitoring templates, and meeting summaries well. But "develop" in practice includes assessments, legal compliance under IDEA, multi-party negotiation, and parent/team consensus—high-stakes, non-digital work that goes far beyond checking model output. The human does a majority of the task through judgment and human-to-human interaction → <50% automatable at high quality.

Occupation: Construction and Building Inspectors

Task: Inspect and monitor construction sites to ensure adherence to safety standards, building codes, or specifications.

Automation score (T0/T1/T2/T3/T4): T1

Explanation: T1 is right because authority, on-site judgment, and safety liability are inherently human/embodied. Modern AI (CV on photos/video, drone logs, code lookups) can pre-screen and draft reports, but the core task is physical inspection plus enforcement. The given explanation is dated—codes are digital; the real blockers are embodiment, accountability, and legal sign-off.

Occupation: Food and Tobacco Roasting, Baking, and Drying Machine Operators and Tenders

Task: Dump sugar dust from collectors into melting tanks and add water to reclaim sugar lost during processing.

Automation score (T0/T1/T2/T3/T4): T0

Explanation: The task as phrased is purely physical and you've restricted tools to what's on a laptop. GenAI can't perform any part of this task (not "write SOPs" but do the dumping). If the task were broadened to "optimize reclaim procedure / generate checklists," exposure would rise, but for the literal action it's T0.

Occupation: Interpreters and Translators

Task: Refer to reference materials (dictionaries, lexicons, encyclopedias, computerized terminology banks) as needed to ensure translation accuracy.

Automation score (T0/T1/T2/T3/T4): T4

Explanation: GenAI can autonomously retrieve, disambiguate, and compile authoritative references and recommended terms with citations and *confidence scores across most high-resource language pairs. The subtask is fully digital, consistent, and does not normally require human oversight.*

Occupation: Electronics Engineer

Task: Prepare documentation containing information such as confidential descriptions or specifications of proprietary hardware or software, product development or introduction schedules, product costs, or information about product performance weaknesses.

Automation score (T0/T1/T2/T3/T4): T3

Explanation: Given structured inputs (requirements, BOMs, test data), GenAI can generate nearly all documentation variants (design specs, DFMEA summaries, release notes, red-team risk write-ups), maintain traceability tables, and enforce style guides. Consequences and confidentiality mean human review is required, but the model can complete 90–100% of components in digital form.

Occupation: Advanced Practice Psychiatric Nurse

Task: Write prescriptions for psychotropic medications as allowed by state regulations and collaborative practice agreements.

Automation score (T0/T1/T2/T3/T4): T3

Explanation: If the medication choice is already clinically decided, GenAI can populate eRx fields, check interactions/contraindications, map to formularies, and draft prior-auth notes—essentially all digital components. High-stakes legality and safety mandate human sign-off, but execution is 90–100% automatable with oversight. (If the task were "evaluate and decide what to prescribe," exposure would drop; the annotation's verb is "write.")

Now, apply the above rubric to the example below:

Occupation: [Occupation]

Task: [Occupational task]

Automation score (T0/T1/T2/T3/T4):

Explanation:

Below, we highlight the substantive changes for prompts related to **Smart Robots** (other than pure wording changes that replace "generative AI" with "smart robots"):

T Automation Rubric

[...]

Rubric

A smart robot is an AI-powered, often-mobile machine designed to autonomously execute one or more physical tasks. These tasks may rely on, or generate, machine learning, which can be incorporated into future activities or support unprecedented conditions. Smart robots can be split into different types based on the tasks/use cases, such as personal, logistics and industrial.

[...]

Definitions

[...]

Annotation Examples

Occupation: Special Education Teachers, Preschool

Task: Develop individual educational plans (IEPs) designed to promote students' educational, physical, or social development.

Automation score (T0/T1/T2/T3/T4): T1

Explanation: A robot can assist with data capture (sensor-based observations), simple assessments, and pre-filling forms, but developing IEPs requires pedagogical judgment, legal compliance, and multi-party collaboration—most of which remains human.

Occupation: Construction and Building Inspectors

Task: Inspect and monitor construction sites to ensure adherence to safety standards, building codes, or specifications.

Automation score (T0/T1/T2/T3/T4): T3

Explanation: Robots (drones/UGVs) with CV/LiDAR can navigate, capture, measure, compare against BIM/specs, and draft reports (90–100% of components). Human oversight is needed for code interpretation, contractor communication, and legal sign-off.

Occupation: Food and Tobacco Roasting, Baking, and Drying Machine Operators and Tenders

Task: Dump sugar dust from collectors into melting tanks and add water to reclaim sugar lost during processing.

Automation score (T0/T1/T2/T3/T4): T4

Explanation: Repetitive material handling and dosing in a controlled plant are fully automatable with robotic manipulation, sensing, and safety interlocks.

Occupation: Interpreters and Translators

Task: Refer to reference materials (dictionaries, lexicons, encyclopedias, computerized terminology banks) as needed to ensure translation accuracy.

Automation score (T0/T1/T2/T3/T4): T4

Explanation: Purely digital retrieval/matching. A robot running GenAI can autonomously consult termbases, disambiguate senses, enforce glossaries, and return citations without routine human oversight.

Occupation: Electronics Engineer

Task: Prepare documentation containing information such as confidential descriptions or specifications of proprietary hardware or software, product development or introduction schedules, product costs, or information about product performance weaknesses.

Automation score (T0/T1/T2/T3/T4): T3

Explanation: With structured inputs (requirements, BOMs, test data), a robot+GenAI stack can draft nearly all documents and maintain traceability. Human review remains for accuracy, confidentiality, and compliance.

Occupation: Advanced Practice Psychiatric Nurse

Task: Write prescriptions for psychotropic medications as allowed by state regulations and collaborative practice agreements.

Automation score (T0/T1/T2/T3/T4): T3

Explanation: A robot can complete the eRx workflow (populate fields, check interactions, format to payer formularies, draft prior auth), but human authorization/clinical judgment is required; oversight is routine.

Occupation: Warehouse Workers

Task: Move inventory from receiving dock to storage locations using hand trucks or pallet jacks.

Automation score (T0/T1/T2/T3/T4): T4

Explanation: AMRs/AGVs integrated with WMS can autonomously transport pallets/totes end-to-end in structured warehouses.

Occupation: Assembly Line Workers

Task: Attach components to products moving along assembly line according to specifications.

Automation score (T0/T1/T2/T3/T4): T4

Explanation: In the typical modern assembly context, robots can complete all attachment steps with consistently high quality. Errors are caught by automated fail-safes/poka-yoke and do not require routine human oversight; technicians intervene only on rare exceptions or maintenance, which is outside the task scope.

Now, apply the above rubric to the example below:

[...]

Below, we highlight the substantive changes for prompts related to **Autonomous Vehicles** (other than pure wording changes that replace "generative AI" with "autonomous vehicles"):

T Automation Rubric

[...]

Rubric

An autonomous vehicle is one that can drive itself from a starting point to a predetermined destination in "autopilot" mode using various in-vehicle technologies and sensors, including adaptive cruise control, active steering (steer by wire), anti-lock braking systems (brake by wire), GPS navigation technology, lasers and radar.

[...]

Definitions

[...]

Annotation Examples

Occupation: Special Education Teachers, Preschool

Task: Develop individual educational plans (IEPs) designed to promote students' educational, physical, or social development.

Automation score (T0/T1/T2/T3/T4): T0

Explanation: *Purely cognitive/interpersonal; no driving component for an AV to perform.*

Occupation: Construction and Building Inspectors

Task: Inspect and monitor construction sites to ensure adherence to safety standards, building codes, or specifications.

Automation score (T0/T1/T2/T3/T4): *T1*

Explanation: *An AV can transport the inspector to/around sites, but the inspection, judgments, and sign-off remain human; AV contributes a minority transport component.*

Occupation: Veterinarians

Task: Drive mobile clinic vans to farms so that health problems can be treated or prevented.

Automation score (T0/T1/T2/T3/T4): *T2*

Explanation: *AVs can perform most road driving to rural sites, but last-meters access (gates, unmarked farm roads, ad-hoc parking/turnarounds) and dynamic on-site constraints often require human intervention. Overall, the AV covers a large portion of the task, but not reliably* $\geq$*90% across most contexts.*

Occupation: Correctional Officers and Jailers

Task: Drive passenger vehicles and trucks used to transport inmates to other institutions, courtrooms, hospitals, and work sites.

Automation score (T0/T1/T2/T3/T4): *T3*

Explanation: *On-road transport between facilities is highly automatable; AVs can execute routing and vehicle control end-to-end. However, the context is high-stakes (security protocols, perimeter handoffs, incident response), so human oversight remains standard even if the driving component is largely automated.*

Occupation: Taxi Drivers and Chauffeurs

Task: Test vehicle equipment, such as lights, brakes, horns, or windshield wipers, to ensure proper operation.

Automation score (T0/T1/T2/T3/T4): *T4*

Explanation: AV platforms can autonomously run pre-trip self-checks and diagnostics (actuate systems, read sensors/OBD, verify via cameras), producing pass/fail results without routine human oversight in most contexts.

Now, apply the above rubric to the example below:

[...]

D.2.3 Simplification of Tasks

The simplification channel assesses how technologies change the skill requirements for performing tasks. We ask GPT-4o to evaluate skill levels both without and with technology access, allowing us to measure the change in required skills. The prompt asks for both values simultaneously. For example, with **Generative AI**:

The occupation [Occupation] contains the task: [Occupational task].

Technology name: Generative AI

Technology description: Generative AI refers to AI techniques that learn a representation of artifacts from data, and use it to generate brand-new, unique artifacts that resemble but don't repeat the original data. These artifacts can serve benign or nefarious purposes. Generative AI can produce totally novel content (including text, images, video, audio, structures), computer code, synthetic data, workflows and models of physical objects. Generative AI also can be used in art, drug discovery or material design.

What level of skill in [Skill] is needed to perform the task in this occupation well WITHOUT access to Generative AI? What level of skill in [Skill] is needed to perform the task in this occupation well WITH access to Generative AI?

Provide the answers on a scale from 1 to 7, where 2 means [Skill Level 2 Anchor], 4 means [Skill Level 4 Anchor], and 6 means [Skill Level 6 Anchor].

Output only two integers separated by a comma, valued between 1 and 7. The first integer is the skill level WITHOUT access to Generative AI, the second integer is the skill level WITH access to Generative AI. Do not output anything else.

All 35 O*NET skills and 19,530 O*NET tasks are evaluated independently. Skill level anchors and task descriptions are drawn from O*NET as discussed in Appendix D.1.

This approach allows us to measure both the baseline skill requirements $\boldsymbol{r}_\tau$ and the technology-adjusted requirements $\boldsymbol{r}'_\tau$ 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 the simplification effect of the technology on task skill requirements.

References

ADDISON, J. T., L. CHEN, AND O. D. OZTURK (2020): "Occupational skill mismatch: Differences by gender and cohort," *ILR Review*, 73, 730–767.

BALEY, I., A. FIGUEIREDO, AND R. ULBRICHT (2022): "Mismatch cycles," *Journal of Political Economy*, 130, 2943–2984.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

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.

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.

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," NBER Working Paper No. 33641.

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," Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013, The Wharton School Research Paper.

DILLON, E. W., S. JAFFE, N. IMMORLICA, AND C. T. STANTON (2025): "Shifting Work Patterns with Generative AI," 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.

GAMBACORTA, L., H. QIU, S. SHAN, AND D. M. REES (2024): "Generative AI and Labour Productivity: A Field Experiment on Coding," Working paper.

HUMLUM, A. (2019): "Robot adoption and labor market dynamics," *Princeton University*.

KANAZAWA, K., D. KAWAGUCHI, H. SHIGEOKA, AND Y. WATANABE (2022): "AI, Skill, and Productivity: The Case of Taxi Drivers," Working Paper.

KEANE, M. P. AND K. I. WOLPIN (1997): "The career decisions of young men," *Journal of Political Economy*, 105, 473–522.

LOWHAGEN, N., P. SCHWENDENER, AND T. NETLAND (2025): "Can a Troubleshooting AI Assistant Improve Task Performance in Industrial Contexts?" *International Journal of Production Research*.

NOY, S. AND W. ZHANG (2023): "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381, 187–192.

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.

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.

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.

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.