Generative AI and Occupational Entry Barriers: The Labor-Supply Channel of Technological Change*

Seyed M. Hosseini† Guy Lichtinger‡

First version: January 19, 2026
This version: April 26, 2026
[Latest version]

Abstract

How will generative AI (GenAI) affect wage inequality across occupations? Beyond standard labor-demand effects, GenAI can reshape the wage distribution by lowering the expertise required to enter occupations, thereby expanding the pool of potential workers. We develop a general-equilibrium model in which changes in occupational wages reflect the interaction between these labor supply shifts and productivity gains, generating a productivity–scarcity race. Using O*NET task data and large language model (LLM) ratings, we construct novel measures of GenAI-induced changes in occupational expertise requirements and the resulting Potential Supply Shifts (PSS), alongside a companion measure of Potential Productivity Gains (PPG). We validate the LLM-based measures against occupational characteristics and real-world GenAI usage data. In our quantitative exercises, these forces operate in opposite directions: occupational productivity gains widen between-occupation wage inequality, whereas reduced entry barriers compress it.

*This research builds on the framework developed by Autor and Thompson (2025), and we are grateful for their work in providing the foundation for our analysis. We also thank David Autor, Netanel Ben-Porath, Gabriel Chodorow-Reich, Lawrence Katz, Bob Margo, Amanda Pallais, Jesse Shapiro, Peyman Shahidi, Ludwig Straub, and Neil Thompson for valuable feedback and discussions. We are grateful as well to participants in the *Economics of AI Reading Group* at MIT, the *Economics of AI Conference* at IESE Business School, the *2nd Asian Conference on Organizational Economics* at HKU, and the *AI, Health and Well-Being Conference* at University of Pittsburgh. An earlier version of this paper circulated under the title "Generative AI, Expertise, and Effective Labor Supply." All remaining errors are our own.

†Harvard University. Email: shosseinimaasoum@fas.harvard.edu.
‡Harvard University. Email: guylichtinger@g.harvard.edu.

1 Introduction

The rapid diffusion of generative AI (GenAI) tools since 2023 and the accelerating pace of model improvement have raised a central question for labor economics: how will GenAI reshape the wage distribution? Much of the existing discussion approaches this question from the labor-demand side, focusing on how GenAI raises productivity, automates subsets of tasks, and creates new ones (e.g., Eloundou et al., 2024). This emphasis continues a long tradition in labor economics of studying how technology reshapes relative labor demand, from the race between education and technology (Tinbergen, 1974; Katz and Murphy, 1992; Goldin and Katz, 2008) to task-based models in which technology substitutes for labor in some tasks and complements it in others (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018).

In this paper, we study how GenAI affects the wage distribution across occupations. We argue that, in addition to its demand-side effects, GenAI can reshape occupational outcomes through a labor-supply channel, à la Autor and Thompson (2025), by lowering expertise-based entry barriers. In the spirit of Autor and Thompson (2025), we define occupational expertise as the specialized knowledge, training, or credentials required to perform an occupation's tasks at a professional level. By lowering the human expertise required to perform occupational task bundles, GenAI can expand the pool of workers who are qualified to perform a given job. For example, GenAI systems that provide real-time guidance can enable nurse practitioners to carry out components of diagnosis that previously required physicians, and paralegals to perform components of document production that previously required attorneys and their specialized legal expertise.[1] This expansion of potential labor supply increases occupational contestability and can put downward pressure on wage premia. The overall distributional impact of GenAI therefore depends on the interaction between demand-side forces, such as productivity gains and automation, and supply-side forces operating through occupational entry barriers.

We study these mechanisms theoretically and empirically. We begin by developing a parsimonious general equilibrium model of occupations and expertise, inspired by Autor and Thompson (2025). The aim is to provide a simple and tractable framework that iso-

[1] Autor (2024) suggests that AI can "enable a larger set of workers possessing complementary knowledge to perform some of the higher-stakes decision-making tasks currently arrogated to elite experts like doctors, lawyers, coders and educators."

lates the two channels at the core of our analysis, maps naturally to the empirical objects we construct, and provides the basis for the quantitative counterfactuals that follow. In the model, occupations have heterogeneous expertise requirements and workers differ in a one-dimensional index of expertise (this one-dimensionality assumption is later relaxed in the dynamic extension). Workers can work only in occupations for which they meet the requirement, implying that workers who qualify for high-requirement occupations also qualify for all lower-requirement ones, and wages are determined competitively under CES demand. To map expertise requirements into realized employment, we embed occupational choice in a discrete-choice framework: more expert workers are more likely to sort into higher-barrier occupations, but may choose lower-barrier jobs due to idiosyncratic fit and outside options.

In this framework, technology such as GenAI affects occupational wages through two distinct channels. First, by augmenting labor productivity, GenAI shifts occupational labor demand outward, raising wages most in occupations where it delivers the largest productivity gains. Second, by changing the expertise required to perform an occupation's task bundle, GenAI can alter entry barriers and, when it lowers them, expand the pool of workers who can supply labor to that occupation, putting downward pressure on scarcity premia. The equilibrium effect on the wage distribution depends on how these two forces vary across occupations. A first-order approximation makes the resulting *race between productivity and scarcity* transparent: the change in each occupation's wage in partial equilibrium equals the net effect of a demand-side term proportional to its *Potential Productivity Gain* (PPG) and a supply-side term proportional to its *Potential Supply Shift* (PSS)—the change in the number of newly eligible workers when an occupation's expertise threshold shifts—with the labor demand elasticity governing their relative strength.[2,3]

The PSS index is the key object capturing the supply-side effect of GenAI. In the model, PSS is determined by several components. First, each occupation has a baseline expertise requirement that aggregates the expertise needed for its individual tasks. GenAI can

[2]The term *productivity–scarcity race* echoes the classical race between education and technology (Tinbergen, 1974; Goldin and Katz, 2008), where technology has been studied primarily through its effect on labor demand. Here, by contrast, we focus on how technology also shifts supply, by reducing the expertise required to enter an occupation.

[3]Under stronger assumptions, we also derive (in Appendix A.4) an exact closed-form race condition for the *expertise premium*—the elasticity of wages with respect to occupational expertise requirements. This object is related to, but distinct from, overall wage inequality.

change this requirement through two margins: an *extensive margin*, in which some tasks are fully automated and removed from the occupation's task bundle, and an *intensive margin*, in which tasks that remain human-performed become easier because workers have access to a GenAI assistant. Combining these two margins yields a post-GenAI expertise requirement. PSS then measures how many additional workers clear this new, potentially lower threshold, given the distribution of expertise in the workforce. Like common exposure measures in the literature (e.g., Eloundou et al., 2024), PSS is therefore an occupation-level index constructed from task-level assessments that can be computed without solving for equilibrium. But whereas exposure indices summarize the demand-side scope for GenAI to affect an occupation by identifying which tasks can be automated or augmented, PSS captures the supply-side consequence: how much the pool of qualified workers changes when those task-level changes alter the occupation's entry barrier.

Motivated by the model, we construct empirical counterparts to its two channels using O*NET task statements for 19,265 tasks across 923 occupations. Constructing the PSS measure requires three ingredients: baseline task expertise, GenAI-induced changes in those requirements, and a mapping from expertise changes to shifts in the qualified labor pool. For baseline expertise, we follow the definition of expertise in Autor and Thompson (2025) as a barrier to entry and use large language models (LLMs) to assign expertise scores to each task. Tasks that require more education, specialized training, or credentials receive higher scores, while tasks that most workers can perform with minimal instruction receive lower scores. The resulting occupation-level expertise measures are highly stable across frontier LLMs and strongly correlated with O*NET indicators of skill, including required education, experience, and wages. To capture the effect of GenAI on expertise requirements, we quantify the two margins: the extensive margin (task automation) and the intensive margin (reductions in task difficulty). We measure the extensive margin by updating the task-level automation exposure measure of Eloundou et al. (2024) to reflect current GenAI capabilities, and the intensive margin by re-rating task expertise under the counterfactual that workers have access to a capable GenAI assistant. Finally, we translate the resulting occupation-level expertise changes into supply shifts by approximating the workforce expertise distribution using baseline occupational requirements.[4]

[4]The full datasets are available at https://github.com/s-mahdihosseini/GenAI_Expertise.

We validate the extensive–intensive distinction using real-world AI usage data from the Anthropic Economic Index (Appel et al., 2026). The AEI classifies user–GenAI interactions across O*NET tasks into categories, including task delegation (the user offloads the task to GenAI, conceptually aligned with our extensive margin) and AI-assisted learning (the user retains the task while using GenAI for guidance, aligned with our intensive margin). Consistent with this mapping, our extensive-margin measure is positively correlated with task delegation and negatively correlated with AI-assisted learning, while our intensive-margin measure shows the reverse pattern. This pattern supports the interpretation that our LLM-based measures of the two margins indeed capture distinct mechanisms: task automation and task assistance.

On the demand side, we construct the empirical PPG measure by aggregating task-level time savings into an occupation-level productivity index using O*NET task frequency and importance weights. This empirical implementation mirrors the model's decomposition: PSS captures how GenAI changes occupational entry barriers and the potential labor supply, while PPG captures how it changes labor productivity and, hence, labor demand.

Next, we examine how PSS and PPG vary across occupations. PSS increases with occupational expertise but declines sharply in the upper tail. This pattern reflects the thin right tail of the expertise distribution: when few workers are near very high expertise thresholds, even large reductions in required expertise generate only modest expansions in the pool of eligible workers. By contrast, PPG rises almost monotonically with expertise, indicating larger GenAI-driven time savings in higher-expertise occupations. These patterns help interpret the quantitative results that follow. We also compare PSS to the dominant empirical approach for measuring GenAI's effects across occupations, namely exposure indices that classify tasks as affected or automatable and average them into occupation-level exposure levels (e.g., Eloundou et al., 2024). We find that PSS varies substantially even among occupations with nearly identical exposure levels, implying that exposure alone is an incomplete statistic for changes in entry barriers and effective labor supply. This highlights the value of PSS as a complementary, supply-side measure when assessing GenAI's labor-market impacts.

We then embed the empirical measures of expertise changes and productivity gains into the deliberately parsimonious calibrated model to quantify their general-equilibrium

implications for occupational wages. To isolate the model's two channels, we study three counterfactual GenAI scenarios: one in which GenAI affects occupations only through productivity gains, holding expertise requirements fixed; one in which it affects occupations only through reductions in expertise requirements, holding productivity fixed; and one in which both channels operate simultaneously. The quantitative results align with the descriptive evidence. Lower-wage occupations experience limited direct productivity gains, but their wages rise indirectly as reduced entry barriers in higher-expertise occupations draw workers upward, tightening labor supply at the bottom.[5] Occupations in the upper-middle of the wage distribution realize substantial productivity gains but also face meaningful declines in scarcity premia as entry barriers fall. At the top of the expertise distribution, wages rise sharply from productivity improvements while supply responses remain small, since few workers lie near very high expertise thresholds. Overall, the quantitative exercises confirm the model's productivity–scarcity race: productivity gains tend to widen cross-occupation wage inequality, while reductions in expertise requirements compress it. In our calibration, productivity wins this race: wage inequality rises on net, as the compression from lower entry barriers offsets only part of the divergence generated by productivity gains.

We complement the parsimonious analysis with a dynamic extension that incorporates occupation-specific expertise and worker switching costs, tracing how the entry-barrier channel unfolds over time. The model replaces the one-dimensional expertise threshold with a richer occupation-to-occupation retraining cost matrix. This captures the intuition that, for example, moving from nursing into medicine is far easier than moving from nursing into law—even when the destination occupations require similar levels of overall expertise.[6] We construct this matrix using LLM-based estimates of pairwise retraining times across all 3-digit SOC occupations, and validate it against an independent O*NET-based measure of directed skill gaps. Beyond the retraining distance, the model allows some occupations to be inherently harder to leave and others harder to enter through occupation-specific adjustment costs, which are calibrated to match observed

[5] This is related to the "ripple effects" in Acemoglu and Restrepo (2022): shocks in one part of the occupational distribution can affect wages elsewhere through worker reallocation. However, the source of reallocation differs: in their model it comes from automation-induced displacement, while here it comes from lower entry barriers in higher-expertise occupations.

[6] The pairwise switching-cost framework nests the hierarchical expertise model as a special case. If switching costs are set to zero for moves into occupations with equal or lower expertise requirements and to infinity otherwise, the retraining matrix is reduced to the hierarchical expertise case.

worker transition patterns from a large longitudinal sample of U.S. workers (Schubert et al., 2024).

We model the AI shock as a heterogeneous reduction in pairwise retraining distances, with some occupation pairs affected far more than others. We find that—consistent with the static framework—this channel reduces long-run between-occupation wage inequality by approximately 9 percent. The compression is concentrated in the upper half of the distribution, where reduced switching costs erode scarcity premia as more workers can retrain into high-wage occupations. The adjustment is gradual: at the calibrated parameters, roughly 80 percent of the long-run effect is realized within nine years, with the remainder materializing thereafter. At the occupation level, high-wage cognitive occupations absorb large net employment inflows, driving down their wages. In contrast, low-wage service occupations shed workers, and their wages rise modestly as labor supply tightens. These dynamics highlight that the equalizing effects of reduced entry barriers may take time to fully materialize and play out unevenly across occupations.

We conclude by noting several limitations. The LLM-based counterfactual assessments—particularly how GenAI changes task expertise—are inherently difficult to validate and should be interpreted cautiously. That said, related evidence is broadly consistent with our framework: Klein Teeselink and Carey (2026) construct an independent LLM-based expertise measure that is highly correlated with ours, and, using job postings across 39 countries and exploiting ChatGPT's release as a shock, find that expertise-raising automation increases advertised wages while expertise-lowering automation does not. In addition, we hold the set of tasks and occupations fixed, abstracting from endogenous task creation and task rebundling dynamics (e.g., Acemoglu and Restrepo, 2019; Autor et al., 2024; Garicano, Luis and Li, Jin and Wu, Yanhui, 2026) that may materially alter these patterns. Finally, realized effects will depend on the pace and extent of adoption, which we do not model. With these caveats, the central message is clear: beyond productivity, GenAI reshapes wage inequality by altering occupational entry barriers. The balance between productivity gains and expertise-driven supply expansions will play a decisive role in determining the distributional impact of GenAI.

The remainder of the paper proceeds as follows. Section 2 briefly discusses the related literature. Section 3 presents the static general-equilibrium model in which GenAI affects occupational wages through entry-barrier erosion and productivity gains, and derives

a first-order characterization of the productivity–scarcity race. Section 4 describes the data and constructs our key measures. Section 5 embeds these measures in the calibrated model and decomposes GenAI's inequality implications into supply-side and demand-side components. Section 6 extends the framework to a dynamic setting with pairwise switching costs. Section 7 concludes.

2 Related Literature

The economic analysis of technology-driven wage inequality has a long intellectual history. A tradition dating to Tinbergen (1974) and Freeman (1975), formalized in Katz and Murphy (1992) and synthesized by Goldin and Katz (2008), frames skill premia as the outcome of a *race between education and technology*: skill-biased technology raises relative demand for skilled labor, while the supply of educated workers determines whether premia widen or compress. This paper extends this seminal framework. Building on Autor and Thompson (2025), we argue that technology itself—specifically GenAI—can also shift the supply side of the race, by changing the expertise required to perform occupations and, with it, the set of workers who qualify.[7]

Relative to Autor and Thompson (2025), our contribution is to adapt the expertise framework to generative AI and to quantify its implications for between-occupation wage inequality. We introduce a novel occupation-level measure, the *Potential Supply Shift* (PSS), which maps GenAI-induced changes in task-level expertise requirements into changes in the share of workers newly qualified to perform each occupation. We also extend the expertise framework by allowing technology to affect expertise through an intensive margin in which tasks remain human-performed yet become easier with AI assistance. We then combine these supply-side forces with occupation-level productivity gains in a calibrated general equilibrium model to study how GenAI's demand- and supply-side effects jointly shape occupational outcomes and wage inequality. Finally, we extend the static, one-dimensional expertise framework to a dynamic setting in which entry barriers are

[7] Danieli (2024) studies a related mechanism in the context of routine-biased technological change: skill-replacing technology lowers the return to skill in routine occupations, directly compressing wage gaps within those occupations and also inducing compositional changes as higher-skill routine workers exit and routine jobs become lower-skill on average. Our focus instead is on occupation-level entry-barrier erosion under GenAI across the occupational distribution.

represented by a pairwise retraining cost matrix across occupations.

Our analysis also builds on the task-based approach to technological change (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018), which decomposes production into tasks and allows technology to substitute for labor in some tasks and complement it in others. A large body of recent work applies this approach to AI, measuring occupation-level "exposure" from task content (e.g., Eloundou et al., 2024; ILO, 2025; Hampole et al., 2025; Felten et al., 2023). These measures summarize how strongly GenAI applies to an occupation's tasks and are typically used to study demand-side channels such as automation and productivity gains. We complement this demand-side tradition with a distinct supply-side margin: technology's effect on the *entry barrier* into an occupation, defined as the expertise required to perform its task bundle. The *Potential Productivity Gain* (PPG) index we construct is in the spirit of exposure indices; the *Potential Supply Shift* (PSS) is the corresponding occupation-level object on the supply side.

A fast-growing empirical literature uses these exposure measures to estimate the labor market effects of GenAI to date by combining them with administrative employment data, job postings, and online platform outcomes. Several studies find that employment and hiring of young workers in the most exposed occupations have declined since the launch of ChatGPT (Brynjolfsson et al., 2025; Klein Teeselink, 2025; Massenkoff and McCrory, 2026), especially within adopting firms (Hosseini and Lichtinger, 2025). Complementing this evidence, Klein Teeselink and Carey (2026) analyze hundreds of millions of job postings across 39 countries and find that exposure predicts declines in job postings following the release of ChatGPT, while wage responses vary depending on whether automation lowers or raises expertise requirements. Moreover, de Souza (2026) uses administrative software registry data from Brazil and finds that AI in production settings increases employment of low-skilled and young workers by enabling them to perform complex tasks that previously required expertise, providing direct evidence of the barrier-reduction channel central to our framework. Similarly, Cruces et al. (2026) find that access to a generative-AI assistant significantly reduces the productivity gap between lower- and higher-education workers in a randomized experiment where workers were asked to complete a business problem-solving task.

Our analysis is closest to recent task-based approaches that translate GenAI's task-

level capabilities into predictions about wages. Freund and Mann (2025) study how automation reshapes the task mix within jobs and affects wages through changes in task importance and worker productivity across tasks. A key implication is that occupation-level exposure shares can conceal substantial heterogeneity because outcomes depend on which tasks are affected. In contrast, we focus on a different margin: how GenAI lowers occupational entry requirements and thereby expands the set of workers who can credibly supply labor to an occupation, which we summarize with an occupation-level Potential Supply Shift (PSS) index, designed to be a simple counterpart to exposure measures but targeted to supply-side effects.

The closest paper in spirit to ours is the recent study by Althoff and Reichardt (2026), who develop and estimate a dynamic task-based model in which workers accumulate multidimensional skills over their careers and choose occupations based on comparative advantage. In their framework, AI affects workers through augmentation, automation, and task simplification, and their richer worker-level structure allows them to study within-occupation inequality and individual transition paths. Our approach is complementary but differs along several dimensions. While Althoff and Reichardt (2026) model inequality both within and between occupations, we focus on the between-occupation margin, which allows us to cleanly isolate two channels—entry-barrier erosion and productivity gains—in a parsimonious and tractable general-equilibrium framework. This tractability also enables us to derive closed-form conditions for the productivity–scarcity race (Appendix A.4), and yields the Potential Supply Shift (PSS), a model-driven index that can be computed directly from task-level data without solving for equilibrium. Additionally, our dynamic extension features a pairwise switching-cost matrix that captures asymmetric retraining costs across occupation pairs, whereas Althoff and Reichardt (2026) use a single scalar switching cost and let differential mobility arise from skill distances. The two models also yield different quantitative implications. Althoff and Reichardt (2026) find that task simplification is strongly and monotonically equalizing while productivity gains are roughly uniform across the wage distribution. In contrast, in our quantitative exercise productivity gains disproportionately benefit higher-wage occupations, increasing between-occupation inequality, while entry-barrier reductions generate a U-shaped pattern of wage changes whose equalizing effect is realized gradually as workers face switching costs that slow reallocation.

Our dynamic extension also relates to a broader literature on worker reallocation with occupational switching frictions, including Artuç et al. (2010), Dix-Carneiro (2014), Caliendo et al. (2019), and Traiberman (2019). Especially relevant is Traiberman (2019), who studies worker adjustment to import competition in a dynamic occupational choice model with occupation-specific human capital and switching costs estimated from observed worker transitions. Our framework differs in two main respects. First, the shock in our setting operates through GenAI-driven reductions in occupational entry barriers and retraining requirements rather than trade-induced demand shifts. Second, rather than recover all switching frictions from realized worker flows, we measure pairwise retraining requirements directly using LLM-based assessments, while still calibrating occupation-specific adjustment costs to observed transitions.

3 Model

This section develops a parsimonious general-equilibrium model of occupational wages and inequality. The model is designed to isolate, with minimal structure, the two channels we can discipline tightly with our data: (i) *productivity gains* from GenAI that raise marginal productivity of labor and shift labor demand, and (ii) *expertise-based scarcity* that governs occupational entry and effective labor supply. The goal is to provide a simple and tractable framework that makes these mechanisms transparent and yields objects that can be brought to the data. The model also provides the basis for the quantitative counterfactuals in Section 5. The key distributional mechanism is a race between these forces: productivity gains tend to widen wage differentials across occupations, while reductions in expertise barriers increase contestability and compress scarcity premia.

3.1 Environment: Workers, Occupations, and Tasks

The economy contains a unit mass of workers. Worker i has an innate expertise level $e_i \in \mathbb{R}_+$ drawn from a continuous distribution F with density f.[8]

[8]The static model treats expertise as a one-dimensional, general index. The dynamic extension in Section 6 relaxes this by replacing the single threshold with a matrix of occupation-to-occupation retraining distances, capturing occupation-specific components of expertise in a flexible way.

There is a finite set of occupations $\mathcal{O}$. Each occupation $o \in \mathcal{O}$ consists of a finite set of tasks T_o. Each task $t \in T_o$ has an expertise requirement $r_{ot} \in \mathbb{R}_+$, representing the minimum training or knowledge needed to perform the task at a professional level, and an importance weight $\varphi_{ot} > 0$, reflecting how critical the task is to the occupation's core function.

Expertise requirement (entry barrier). The occupation-level expertise requirement aggregates task-level requirements via a generalized (power) mean:

$$R_o := \left(\frac{\sum_{t \in T_o} \varphi_{ot} \, r_{ot}^{\rho}}{\sum_{t \in T_o} \varphi_{ot}} \right)^{1/\rho}, \qquad \rho \geq 1. \tag{1}$$

A worker with expertise e can work in occupation o if and only if $e \geq R_o$. The parameter ρ controls how much the hardest tasks in an occupation determine its entry barrier. At $\rho = 1$, the barrier is the importance-weighted mean of task requirements. As $\rho \to \infty$, it converges to $\max_t r_{ot}$. In this limiting case, the entry condition in our model coincides with that of Autor and Thompson (2025), where each worker must perform all of an occupation's tasks and entry is governed by the most demanding one (the "task-bundling" assumption).[9] Autor and Thompson (2025) show that under task bundling and hierarchical expertise,[10] conditional on entry, a worker's output does not depend on how far her expertise exceeds the barrier—though her expertise remains valuable because it expands the set of occupations she can enter. The strict max is a natural theoretical benchmark but likely overstates how rigidly one task governs entry: in practice, workers can compensate for weakness in one task with strength in others. The power-mean framework nests both extremes—and any intermediate case—as special cases. Appendix A.2 shows that our main findings are robust across a wide range of ρ values.

[9] Appendix A.1 develops the hierarchical-feasibility ($\rho \to \infty$) special case in detail.

[10] In Autor and Thompson (2025), hierarchical expertise means that a worker with expertise e_i can complete task t if and only if $e_i \geq t$.

3.2 Demand: Occupational Production and Final Demand

Occupational output. Occupation o produces a differentiated intermediate input Y_o using labor:

$$Y_o \;=\; A_o\, L_o, \tag{2}$$

where $A_o > 0$ is occupation-specific labor productivity and L_o is employment in occupation o.

At the task level, producing one unit of occupational output requires h_{ot} units of labor time on task t. Total labor time per unit is $H_o = \sum_{t \in T_o} h_{ot}$, so that $A_o = 1/H_o$. Define the task's share of baseline labor time as $\lambda_{ot} := h_{ot}/H_o$. Any technology that reduces h_{ot} for some task t lowers H_o and thereby raises A_o. This formulation is in the spirit of the time-accounting approach in Acemoglu (2025), where task-level time savings map directly into occupation-level productivity gains. Formally, we model tasks within an occupation as combining through a Leontief technology,

$$Y_o \;=\; \min_{t \in T_o} \left\{ \frac{\ell_{ot}}{h_{ot}} \right\},$$

where ℓ_{ot} denotes labor time allocated to task t in occupation o and h_{ot} is the labor time required on that task per unit of occupational output. This implies perfect complementarity across tasks within an occupation—every task must be completed in fixed proportions—consistent with the inseparability principle in Autor and Thompson (2025). Under efficient task allocation, $\ell_{ot} = h_{ot} Y_o$ for all t, so total labor used in occupation o is $L_o = \sum_{t \in T_o} \ell_{ot} = Y_o \sum_{t \in T_o} h_{ot}$, which yields $Y_o = A_o L_o$ with $A_o = 1/H_o$ and $H_o = \sum_{t \in T_o} h_{ot}$.[11]

Final demand. A competitive final-good producer aggregates occupational inputs with a CES technology:

$$Y \;=\; \left(\sum_{o \in \mathcal{O}} \vartheta_o^{\frac{1}{\sigma}}\, Y_o^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 1, \tag{3}$$

[11] See Appendix A.3 for the formal mapping from task-level time savings to the occupation-level productivity shock π_o.

where $\vartheta_o > 0$ captures demand weights and σ is the elasticity of substitution across occupational outputs.[12] The final good is the numeraire.

Under perfect competition, wages equal marginal revenue product. Combining Equations (2)–(3) yields an inverse labor-demand schedule of the form

$$w_o \;=\; \underbrace{\kappa\, \vartheta_o^{\frac{1}{\sigma}}\, A_o^{1-\frac{1}{\sigma}}}_{=:B_o}\; L_o^{-\frac{1}{\sigma}}, \tag{4}$$

where κ is common across occupations (a normalization pinned down by the numeraire). Define the occupation-specific demand–productivity shifter

$$B_o \;:=\; \kappa\, \vartheta_o^{\frac{1}{\sigma}}\, A_o^{1-\frac{1}{\sigma}}. \tag{5}$$

Equation (4) highlights the demand-side source of wage dispersion: higher B_o implies higher wages for a given employment level.

3.3 Labor Supply: Feasibility and Occupational Choice

Eligibility does not mechanically imply employment because eligible workers must choose among multiple feasible occupations. We therefore microfound occupational labor supply via a discrete-choice model with idiosyncratic preferences over occupations.[13]

Feasibility. A worker with expertise e can choose occupation $o \in \mathcal{O}$ if and only if $e \geq R_o$. Define the feasible set $\mathcal{O}(e) := \{o \in \mathcal{O} :\; e \geq R_o\}$.

Occupational choice. Utility from choosing occupation o is

$$u_{io} = \log w_o + \varepsilon_{io},$$

[12] We follow Autor and Thompson (2025), who similarly impose CES gross-substitutability across occupations ($\lambda > 1$) and calibrate $\lambda = 4$ in their numerical exercises. More broadly, $\sigma > 1$ is standard in CES skill-demand models, with estimates between skilled and unskilled labor typically in the range 1.4–2 (Katz and Murphy, 1992; Acemoglu and Autor, 2011).

[13] This specification—a logit with idiosyncratic preferences over differentiated occupations—is standard in the structural occupational-choice literature (e.g., Traiberman, 2019) and is the occupational analog of the firm-level wage-setting model of Card et al. (2018).

where ε_{io} are i.i.d.[14] Type-I extreme value with scale $\tau > 0$, and $u_{io} = -\infty$ if $e_i < R_o$. The resulting choice probabilities follow a multinomial logit. In particular, the probability that a worker with expertise e chooses occupation o, given wages w and requirements R, is

$$s_o(e; w, R) = \frac{w_o^{1/\tau}}{\sum_{k \in \mathcal{O}(e)} w_k^{1/\tau}} \mathbf{1}\{e \geq R_o\}. \tag{6}$$

As $\tau \to 0$, choices concentrate on the highest-wage feasible occupation.

Aggregate labor supply. Employment in occupation $o \in \mathcal{O}$ is

$$L_o^S(w, R) = \int_0^\infty s_o(e; w, R)\, f(e)\, de. \tag{7}$$

Define the *eligible pool*

$$S_o(R) \;:=\; 1 - F(R_o), \tag{8}$$

and the *average take-up rate* among eligible workers $\bar{s}_o(w, R)$ such that $L_o^S = S_o \cdot \bar{s}_o$. This decomposition separates the scarcity margin (eligibility) from the allocation margin (sorting among feasible options).

3.4 Equilibrium and GenAI Shocks

An equilibrium is a vector of wages $\{w_o\}_{o \in \mathcal{O}}$ and employments $\{L_o\}_{o \in \mathcal{O}}$ such that, for each occupation o,

$$w_o = B_o\, L_o^{-1/\sigma} \qquad \text{and} \qquad L_o = L_o^S(w, R). \tag{9}$$

Because utility depends on $\log w_o$, multiplying all wages by a common factor does not affect occupational choice probabilities; hence the level of κ is irrelevant for allocations and can be normalized. We calibrate $\{B_o\}$ and τ in Section 5.1.1.

GenAI affects outcomes through two channels in the benchmark model. Both operate at the task level and aggregate to occupation-level objects via the structures introduced in Section 3.1–3.2.

[14] We abstract from non-wage amenities in the baseline specification; Appendix A.17 shows results are robust to including them.

Productivity (demand) channel: Potential Productivity Gain. GenAI reduces the labor time required for some tasks. Let $g_{ot} \in [0,1]$ denote the fraction of labor time saved on task t in occupation o when a worker uses GenAI. The occupation-level *Potential Productivity Gain* (PPG) is the work-volume-weighted average of task-level time savings:

$$PPG_o := \sum_{t \in T_o} \lambda_{ot} \, g_{ot}, \tag{10}$$

where $\lambda_{ot} = h_{ot}/H_o$ are the baseline labor-time shares defined in Section 3.2.[15] Post-GenAI, labor time per unit falls to $H_o^{AI} = (1 - PPG_o)\, H_o$, which implies a labor-augmenting productivity gain:

$$A_o^{\text{AI}} = A_o \cdot \exp(\pi_o), \qquad \pi_o = -\log(1 - PPG_o). \tag{11}$$

Under Equations (4)–(5), this shifts labor demand through B_o.

Scarcity (expertise) channel: Potential Supply Shift. GenAI relaxes expertise-based entry barriers by changing the occupation's task bundle through two margins. On the *extensive margin*, GenAI fully automates some tasks, removing them from the set performed by labor. Let $a_{ot} \in \{0,1\}$ indicate whether task t is automated. On the *intensive margin*, tasks that remain human-performed become easier: GenAI changes the expertise required from r_{ot} to $\widetilde{r}_{ot}$.[16] Combining both margins, the post-GenAI expertise requirement is

$$\widetilde{R}_o := \left(\frac{\sum_{t \in T_o} \varphi_{ot} \, (1 - a_{ot}) \, \widetilde{r}_{ot}^{\,\rho}}{\sum_{t \in T_o} \varphi_{ot} \, (1 - a_{ot})} \right)^{1/\rho}. \tag{12}$$

[15] The linear aggregation in Equation (10) follows from the time-accounting production structure: each task must be completed, so total time saved is the share-weighted sum of task-level savings. This differs from the power-mean aggregation for expertise (Equation (1)) because the two objects have different economic content—and correspondingly use different weights. Expertise requirements concern *feasibility*—whether a worker can perform the occupation at all—where harder tasks naturally receive disproportionate weight, so the aggregation uses importance weights φ_{ot}. Productivity concerns *efficiency*—how much time the occupation's full task bundle requires—where each task's contribution is proportional to its share of total labor time, so the aggregation uses time shares λ_{ot}.

[16] This extensive margin resonates with the "expertise exposure" concept of Autor and Thompson (2025), though it focuses on potential shifts induced by GenAI automation rather than on realized changes from past technologies. The intensive margin is closely related to "task simplification" in Althoff and Reichardt (2026).

In principle, GenAI could either raise ($\widetilde{R}_o \geq R_o$) or lower ($\widetilde{R}_o \leq R_o$) an occupation's expertise requirement depending on which tasks it automates and simplifies. The change in the entry barrier generates a *Potential Supply Shift*: the shift in the share of the workforce that meets the occupation's expertise threshold.

3.5 The productivity–scarcity Race

Potential Supply Shift (PSS). The key supply-side object is how many workers become newly eligible when entry barriers fall, which we term the Potential Supply Shift (PSS). Fix an occupation $o \in \mathcal{O}$. When the entry requirement falls from R_o to $\widetilde{R}_o$, the increase in the eligible pool is

$$\mathrm{PSS}_o := \big[F(R_o) - F(\widetilde{R}_o)\big]. \tag{13}$$

For a small reduction $\widetilde{R}_o = R_o - \Delta R_o$, $\mathrm{PSS}_o \approx f(R_o)\,\Delta R_o$, so potential inflows are larger when expertise density is high near the pre-AI threshold.

Outside options and marginal take-up. Newly eligible workers have feasible alternatives. Holding wages fixed, define the outside-option index just below the threshold:

$$\Omega(R_o^-) := \sum_{k:\, R_k < R_o} w_k^{1/\tau}.$$

The take-up rate of marginally newly eligible workers is

$$s_o^m := \frac{w_o^{1/\tau}}{w_o^{1/\tau} + \Omega(R_o^-)}. \tag{14}$$

At fixed wages, the induced increase in labor supply is approximately

$$\Delta L_o^S\big|_w \approx s_o^m \cdot \mathrm{PSS}_o. \tag{15}$$

Local incidence. Let $\varepsilon_o := \partial \log L_o^S / \partial \log w_o$ denote the own-wage elasticity of labor supply (holding requirements fixed). A first-order approximation implies

$$\Delta \log w_o \approx \underbrace{\frac{\sigma - 1}{\sigma + \varepsilon_o} PPG_o}_{\text{productivity}} - \underbrace{\frac{1}{\sigma + \varepsilon_o} s_o^m \frac{\text{PSS}_o}{L_o}}_{\text{expertise expansion}}, \tag{16}$$

$$\Delta \log L_o \approx \underbrace{\frac{(\sigma - 1)\varepsilon_o}{\sigma + \varepsilon_o} PPG_o}_{\text{productivity}} + \underbrace{\frac{\sigma}{\sigma + \varepsilon_o} s_o^m \frac{\text{PSS}_o}{L_o}}_{\text{expertise expansion}}, \tag{17}$$

where we use the first-order approximation $\pi_o \approx PPG_o$. These formulas make the productivity–scarcity race[17] transparent: PPG_o enters the demand side, PSS_o enters the supply side, and σ governs the relative strength of each channel. In the quantitative exercise we solve the full system in Equation (9) numerically and allow for reallocation across all occupations. The local expressions above are meant for intuition rather than measurement.

Appendix A.4 shows that, under stronger distributional assumptions, one can go further and obtain an exact closed-form characterization, not of wage inequality itself, but of the *expertise premium*, that is, how steeply wages rise with occupational expertise requirements. The expertise premium is related to, but distinct from, cross-occupation wage inequality, which also depends on how workers are distributed across occupations.[18] The analytical result yields a sharp race condition: the expertise premium steepens when productivity gains rise with expertise more strongly than the proportional expansion in the eligible labor pool does, once the former is scaled by the demand elasticity. Appendix A.5 estimates the underlying gradients directly and shows that the two terms in the race condition are close in magnitude, consistent with the quantitative counterfactuals in Section 5.

[17] The name echoes the classical race between education and technology (Tinbergen, 1974; Goldin and Katz, 2008), with one key difference: here technology shifts both sides of the race—demand through productivity gains and supply through reductions in expertise barriers.

[18] For example, if most workers are concentrated in occupations with similar expertise requirements, the expertise premium can be steep even when overall wage dispersion is low.

3.6 Discussion: Automation and Displacement

The benchmark model focuses on two channels that our data can discipline tightly: productivity gains and expertise-based scarcity. A third channel—task-level displacement, in which firms substitute from labor to automation services—is also potentially relevant but raises distinct identification challenges that motivate treating it separately. The displacement channel operates through firms' decision to substitute from labor to automation services. In task-based models of automation (Acemoglu and Restrepo, 2022; Acemoglu, 2025), the equilibrium incidence of this channel depends on the balance between a *substitution effect* (automation replaces labor in affected tasks) and a *scale effect* (lower unit costs raise output demand, partially restoring labor demand). Which force dominates depends on objects that are difficult to discipline empirically: the effective price of automation services, the within-occupation elasticity of substitution between labor and automation, and the mapping from task-level exposure to structural substitution parameters.[19]

Appendix A.7 formalizes this extension by introducing a CES task-level production function that nests the benchmark model as a special case. Appendix A.8 documents that the displacement channel's effect on inequality is highly sensitive to these additional parameters: across plausible calibrations, it can range from strongly inequality-increasing to inequality-decreasing. In contrast, the productivity and scarcity channels are disciplined entirely by observable task-level data: PPG_o is constructed from task-level exposure indicators and work-volume weights, and PSS_o is constructed from task-level expertise ratings and the empirical distribution of worker expertise (Section 4). We therefore treat the productivity–scarcity race as the benchmark and displacement as a complementary extension whose quantitative incidence remains an open empirical question.

4 Data, Key Variables, and Descriptive Findings

This section takes the model to the data. The model identifies two occupation-level objects that shape the productivity–scarcity race: the Potential Productivity Gains (PPG_o), which enters the demand side through Equation (10), and the Potential Supply Shift (PSS_o),

[19] Relatedly, Autor and Kausik (2026) show that under constant returns to scale, the wage-maximizing labor share depends only on the capital-to-labor ratio, and estimate that all twelve industrialized countries they study exceed this level—implying that further automation would raise wages.

which enters the supply side through Equation (13). Both are constructed from task-level primitives—time savings g_{ot}, expertise requirements r_{ot} and $\widetilde{r}_{ot}$, and automation indicators a_{ot}—and from the current cross-occupation distribution of workers. We measure these objects using O*NET task statements for 19,265 tasks across 923 occupations (O*NET, 2023), complemented by LLM-based evaluations. The complete task-level and occupation-level datasets are publicly available.[20]

We proceed in two parts. Section 4.1 constructs the supply-side objects: the baseline expertise requirement R_o (Equation (1)), the post-GenAI requirement $\widetilde{R}_o$ (Equation (12)), and the resulting PSS_o. Section 4.2 constructs the demand-side object PPG_o. We then document how PSS and PPG vary across the occupational distribution (Section 4.3) and benchmark PSS against standard exposure measures.

4.1 Potential Supply Shift (PSS)

The central supply-side object is the change in the fraction of workers with sufficient expertise to perform an occupation following GenAI effects, defined in the model as $PSS_o = F(R_o) - F(\widetilde{R}_o)$ (Equation (13)). We construct this in three steps. First, we measure the baseline expertise requirement R_o of each O*NET occupation. Second, we estimate how GenAI changes this requirement through the two channels formalized in Section 3.4: an extensive margin (a_{ot}), in which some tasks are automated and removed from the bundle performed by labor, and an intensive margin ($\widetilde{r}_{ot}$), in which remaining tasks become easier to perform with GenAI assistance. Third, we translate the resulting $\widetilde{R}_o$ into changes in the size of the qualified workforce using an empirical approximation of the expertise distribution F.

4.1.1 Step 1: Baseline Expertise of Tasks and Occupations

We begin by constructing a task-level measure of expertise following the framework of Autor and Thompson (2025). Expertise is defined as a barrier to entry: tasks that require specialized training, credentials, or substantial occupation-specific knowledge receive higher scores, while tasks that most workers can perform with minimal instruction

[20]The datasets are available at https://github.com/s-mahdihosseini/GenAI_Expertise.

receive lower scores. Using GPT-5.2, we assign each O*NET task an expertise score on a five-point scale from 1 (minimal expertise) to 5 (very high expertise). The full prompt is provided in Appendix A.12.

We then aggregate task-level expertise to the occupation level. Because expertise is measured on an ordinal scale, we map expertise categories into a continuous measure based on the task-specific training time, measured in months, required for a typical adult to perform the task at a professional level. To construct this mapping, we use an additional LLM prompt to estimate training months for each task without GenAI assistance (Appendix A.13). Rather than imposing a single global mapping, we construct an occupation-specific mapping by assigning, for each occupation and expertise category, the median training time among tasks in that category. This approach captures systematic cross-occupation differences in how expertise translates into training requirements. Table 1 reports the pooled category-to-training-time mapping for reference.

Table 1: Mapping Expertise Categories to Required Training Time

Expertise Category	Description	Training Months
1	No or minimal expertise; generic or basic tasks; learn quickly with little training	0.25
2	Low expertise; short training; limited occupation-specific knowledge	1
3	Moderate expertise; solid occupation-specific knowledge; often requires credentials, apprenticeship, or substantial on-the-job learning	4
4	High expertise; advanced specialized knowledge; significant training, degree, or certification	12
5	Very high expertise; deep specialized expertise; often advanced professional or graduate-level training	60

Notes: The table reports the mapping between the 1–5 ordinal expertise scale and the continuous measure of required training time (in months) used in our analysis. The mapping is derived by calculating the median training time for all tasks within each expertise category, based on estimates generated by GPT-5.2.

Using this mapping, we implement the model's power-mean aggregator (Equation (1)) empirically. Let x_{ot}^{noAI} denote the training-month requirement for task t in occupation o without GenAI—the empirical counterpart of the model's expertise requirement r_{ot}—and let w_{ot} denote the O*NET task weight (1 for core tasks, 0.5 for supplemental tasks, follow-

ing Eloundou et al. 2024), which serves as the empirical counterpart of the importance weight φ_{ot}. We set $\rho = 1$ (the importance-weighted arithmetic mean), following Autor and Thompson (2025), who measure occupational expertise as the average expertise across tasks within an occupation (Appendix A.2 shows that our results are robust to alternative values of ρ). The baseline occupation-level expertise requirement is then

$$X_o^{noAI} = \underbrace{\frac{\sum_{t \in T_o} w_{ot}\, x_{ot}^{noAI}}{\sum_{t \in T_o} w_{ot}}}_{\mathbb{E}[x_{ot}^{noAI}]}, \tag{18}$$

which is the empirical counterpart of R_o from Equation (1).

Appendix A.15 validates this measure, showing it is highly stable across frontier LLMs ($R^2 \geq 0.90$) and strongly correlated with O*NET-based measures of required education and experience and with wages. Figure A.8a plots the distribution of X_o^{noAI}. The mean baseline requirement is 9.61 months (SD 12.48), with a median of 5.24 months.

4.1.2 Step 2: The Effect of GenAI on Occupational Expertise Requirements

Extensive Margin: We next quantify the extensive margin formalized in Section 3.4: GenAI fully automates some tasks, removing them from the set performed by labor (the indicator a_{ot} in Equation (12)). We identify automatable tasks by updating the task-level automation exposure methodology of Eloundou et al. (2024) to reflect current GenAI capabilities. Using GPT-5.2 and an updated prompt (Appendix A.9), the model assigns each task to one of five automation categories based on the share of task components that can be performed by current GenAI systems: 1 (0 percent), 2 (0–50 percent), 3 (50–80 percent), 4 (80–100 percent), and 5 (100 percent automatable). We define a binary automation exposure indicator $\alpha_{ot} \in \{0, 1\}$ (empirical counterpart of the model's a_{ot}) that equals one for tasks in categories 4–5, i.e., at least 80 percent automatable. Under this definition, 4,107 tasks (21.3 percent) are classified as automatable. Appendix A.10 shows that occupation-level exposure shares implied by this updated measure are strongly correlated with the original Eloundou et al. (2024) measure.

Intensive Margin: We next quantify the intensive margin: tasks that remain human-performed become easier, reducing the expertise required from r_{ot} to $\widetilde{r}_{ot}$ in the model's notation (Section 3.4). To measure this channel, we re-rate each task's expertise under a counterfactual in which the worker has access to a capable GenAI assistant, using a separate prompt that combines the expertise definition with explicit GenAI capability assumptions. Importantly, this prompt is issued independently of the baseline, so the counterfactual rating is not a mechanical transformation of r_{ot}. The full prompt is reported in Appendix A.16.

Figure 1a compares the distribution of tasks across expertise categories with and without GenAI assistance, and Figure 1b summarizes the implied transitions. The share of tasks rated as high expertise (category 4) falls from 30.0 percent without GenAI to 8.4 percent with GenAI assistance, while the share in the lowest category rises from 6.0 percent to 22.9 percent. Except for category 1, more than half of tasks in each category are predicted to shift down by one expertise category; the pattern is most pronounced for category 4 tasks, of which 79.0 percent are predicted to require only category 3 expertise with GenAI assistance.

Combining the Extensive and Intensive Margins: Combining both margins yields the empirical counterpart of the model's post-GenAI expertise requirement $\widetilde{R}_o$ (Equation (12)), evaluated at $\rho = 1$:

$$X_o^{combined} = \underbrace{\frac{\sum_{t \in T_o} w_{ot}(1-\alpha_{ot}) x_{ot}^{withAI}}{\sum_{t \in T_o} w_{ot}(1-\alpha_{ot})}}_{\mathbb{E}[x_{ot}^{withAI} | \alpha_{ot}=0]}, \tag{19}$$

where x_{ot}^{withAI} is the empirical counterpart of $\widetilde{r}_{ot}$ (entry barriers with GenAI assistance).

Figure A.8 in Appendix A.14 plots the distributions of occupation-level expertise under each scenario. The combined post-GenAI distribution closely resembles the intensive-margin distribution, indicating that in our exercise the intensive margin is quantitatively dominant. This dominance is sensitive to our calibration: a higher assumed automation threshold or a more capable GenAI would each push more weight onto the extensive margin. Another important caveat is that we abstract from the creation of new tasks and occupations in response to GenAI.

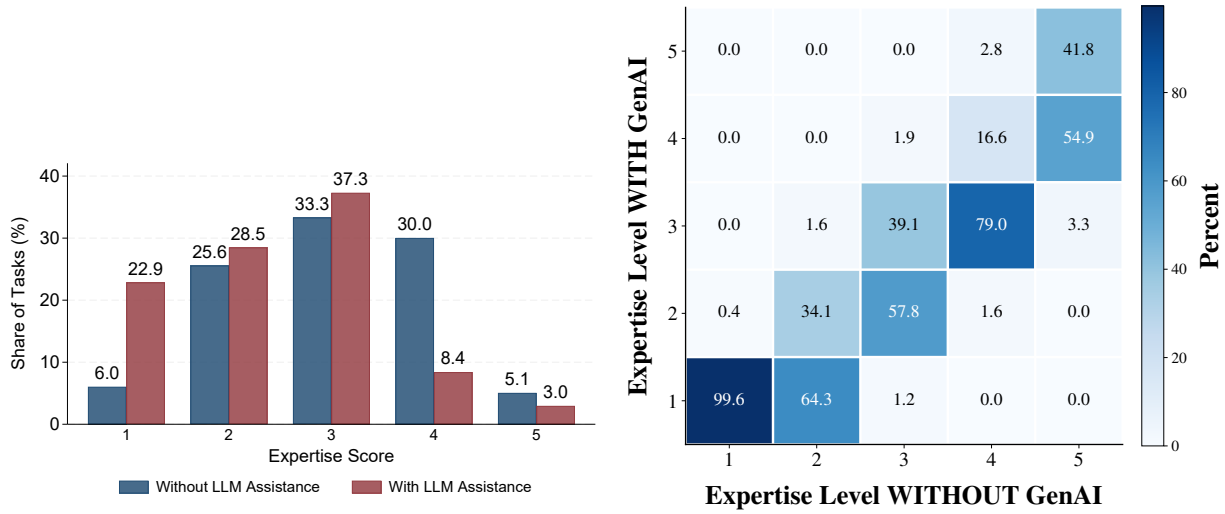

(a) Expertise Categories **(b)** Expertise Transition Matrix

Figure 1: The Expertise Distribution of Tasks With and Without GenAI Assistance

Notes: The figure displays the shift in task-level expertise requirements due to GenAI assistance. Panel (a) plots the distribution of tasks across the five expertise categories under two scenarios: "Without LLM Assistance" (blue bars) and "With LLM Assistance" (red bars). Panel (b) presents a transition matrix comparing the expertise level required without GenAI (x-axis) to the level required with GenAI (y-axis). The size of each bubble and the label above it correspond to the percentage of tasks starting in the x-axis category that transition to the y-axis category.

Validation with Real-World AI Usage Data: As an external validation of our two task-level margins, we compare them to real-world GenAI usage patterns from the Anthropic Economic Index (AEI; Appel et al., 2026). The idea is straightforward. If the extensive and intensive margins capture distinct channels through which GenAI affects work, they should be associated with different modes of AI use in practice. Tasks that are more automatable should be more likely to be delegated to AI, whereas tasks for which AI mainly lowers expertise requirements should be more likely to involve AI-assisted learning or guidance.

The AEI provides task-level measures of how Anthropic's Claude model is used in practice, based on anonymized conversations matched to O*NET task statements. For each task, it reports the log count of conversations involving the task, an average speedup measure (human-only time relative to human-with-AI time), and the shares of conversations classified as "directive" (user delegates, AI executes) and "learning" (user seeks explanations to build skill). We correlate these AEI measures with our two task-level

margins: the extensive margin, measured by the automation indicator ($\alpha_{ot} \in \{0,1\}$), and the intensive margin, measured by the reduction in required training months, $d_{ot} = x_{ot}^{noAI} - x_{ot}^{withAI}$. Figure 2 reports Pearson correlations.

The results are consistent with the distinction between the two margins. Both margins are positively correlated with task usage and speedup. This indicates that tasks our measures identify as more AI-affected are also used more frequently with AI and are associated with larger time savings. But the two margins differ sharply in usage mode. The extensive margin is positively correlated with directive use and negatively correlated with learning use, while the intensive margin shows the opposite pattern. This is consistent with the interpretation that more automatable tasks are more often delegated to AI, whereas tasks for which AI lowers expertise requirements are more likely to involve workers using AI to build capability while remaining actively engaged in the task. These correlations are only suggestive, but they provide useful external evidence that the extensive and intensive margins capture empirically distinct dimensions of GenAI's effect on occupational tasks in real-world applications.

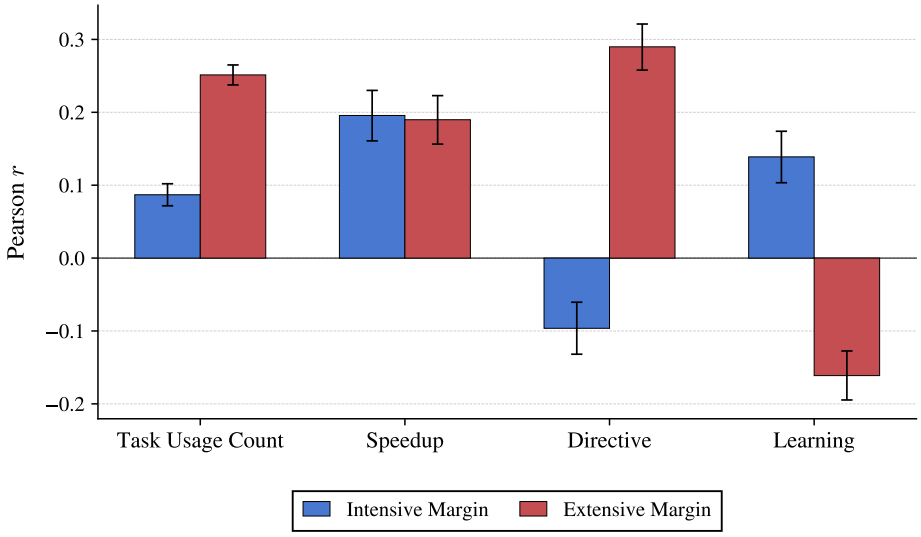

Figure 2: Correlations Between AEI Usage Metrics and Expertise Margins

Notes: The figure reports Pearson correlations between task-level AEI usage metrics and our two margins of GenAI's effect on expertise requirements: the intensive margin (reduction in required training months, blue) and the extensive margin (binary automation indicator, red). AEI metrics are averaged across the Claude.ai and first-party API platforms. Error bars show 95% confidence intervals. Data source: Anthropic Economic Index v4 (Appel et al., 2026).

4.1.3 Step 3: Mapping Expertise Changes to Potential Labor Supply Shifts

The model defines the eligible pool as $S_o(R) = 1 - F(R_o)$ (Equation (8)) and the Potential Supply Shift as $\text{PSS}_o = F(R_o) - F(\widetilde{R}_o)$ (Equation (13)). Implementing these objects requires an estimate of the expertise distribution F. We proxy this using the employment-weighted distribution of baseline occupational expertise requirements X_o^{noAI}.[21] Let $F(x)$ denote the corresponding empirical employment-weighted CDF. When we need $F(\cdot)$ at counterfactual expertise requirements, we evaluate it by linear interpolation over the empirical CDF grid. Figure 3 plots the implied PDF and CDF. The distribution is right-skewed, with substantial mass at low-to-moderate expertise levels and a thin upper tail.

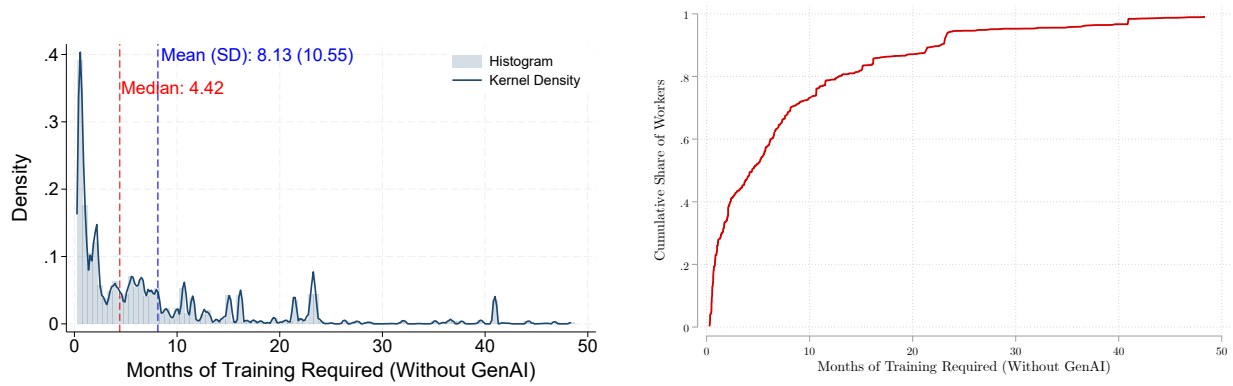

(a) PDF **(b)** CDF

Figure 3: The Distribution of Workers Across Expertise Levels

Notes: Panel (a) displays the probability density function (PDF) and Panel (b) displays the cumulative distribution function (CDF) of baseline occupational expertise, weighted by 2024 employment. This distribution serves as our proxy for the aggregate distribution of workforce expertise.

Applying the model's definitions to the data, the empirical PSS for occupation o is

$$PSS_o = F(X_o^{noAI}) - F(X_o^{combined}), \tag{20}$$

where X_o^{noAI} and $X_o^{combined}$ are the empirical counterparts of R_o and $\widetilde{R}_o$. A positive PSS_o in-

[21] This is a proxy rather than an exact inversion of the model's latent expertise distribution F: in the model, observed employment reflects both expertise-based eligibility and idiosyncratic taste shocks across feasible occupations. In Appendix A.17 we construct an alternative expertise distribution directly from the educational attainment of the CPS data and show that our counterfactual results (Section 5)—both qualitatively and in magnitude—are robust to this alternative measure.

dicates that GenAI lowers effective entry barriers, expanding the fraction of the workforce qualified to perform the occupation. Figure 4 plots the distribution of PSS_o across occupations. The mean PSS is about 0.11, implying that for the average occupation, GenAI expands the qualified labor pool by 11 percentage points of the workforce.

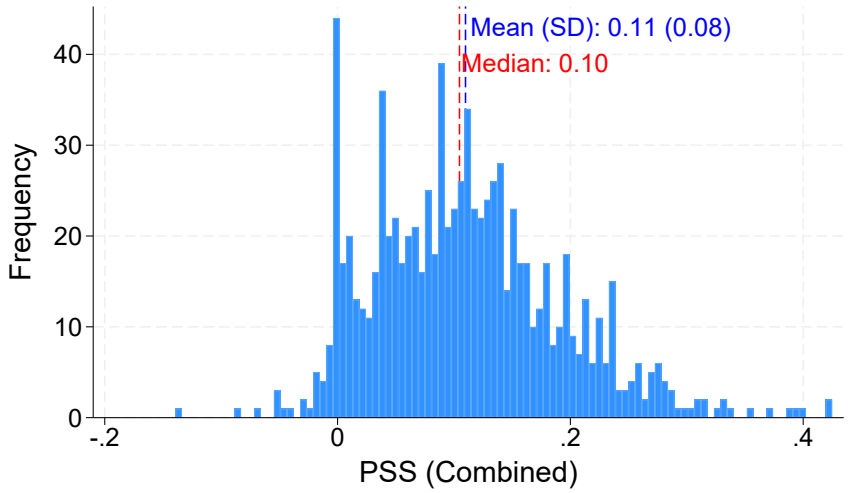

Figure 4: Potential Supply Shift (PSS) Across Occupations

Notes: The figure displays the distribution of the occupation-level Potential Supply Shift (PSS). PSS is defined as the increase in the share of the total workforce capable of performing the occupation after GenAI integration: $F(X_o^{noAI}) - F(X_o^{combined})$.

4.2 The Potential Productivity Gains (PPG) of Occupations

The model defines occupation-level potential productivity gains as $PPG_o = \sum_t \lambda_{ot} g_{ot}$ (Equation (10)), where g_{ot} is the fraction of labor time saved on task t, and λ_{ot} is task t's baseline share of labor time within occupation o. We implement this object using the task-level exposure index β_{ot} from Eloundou et al. (2024) as the empirical counterpart of g_{ot}. Specifically, $\beta_{ot} = 1$ for tasks exposed to an LLM alone, corresponding to time savings of at least 50 percent; $\beta_{ot} = 0.5$ for tasks exposed only with additional software integration; and $\beta_{ot} = 0$ otherwise. To proxy the model's labor-time shares λ_{ot}, we use O*NET measures of task frequency ($Freq_{ot}$, converted to annualized occurrences) and task importance (Imp_{ot}, measured on a 1–5 scale). Their product captures each task's share of baseline work volume within the occupation. We then scale the resulting weighted aver-

age by 0.5 to impose a conservative lower bound on implied time savings. The resulting empirical measure is:

$$PPG_o = 0.5 \times \left(\frac{\sum_{t \in T_o} Freq_{ot} \times Imp_{ot} \times \beta_{ot}}{\sum_{t \in T_o} Freq_{ot} \times Imp_{ot}} \right). \tag{21}$$

Figure 5 plots the distribution of PPG_o across occupations. The mean is 0.17 (SD 0.11), implying that for the average occupation, GenAI could reduce total work time by roughly 17 percent under this conservative calibration. In the quantitative model (Section 5), we map PPG_o into the occupation-level productivity shock using Equation (11), so that $\pi_o = -\log(1 - PPG_o)$. Note that for small values of PPG_o, $\pi_o \approx PPG_o$.

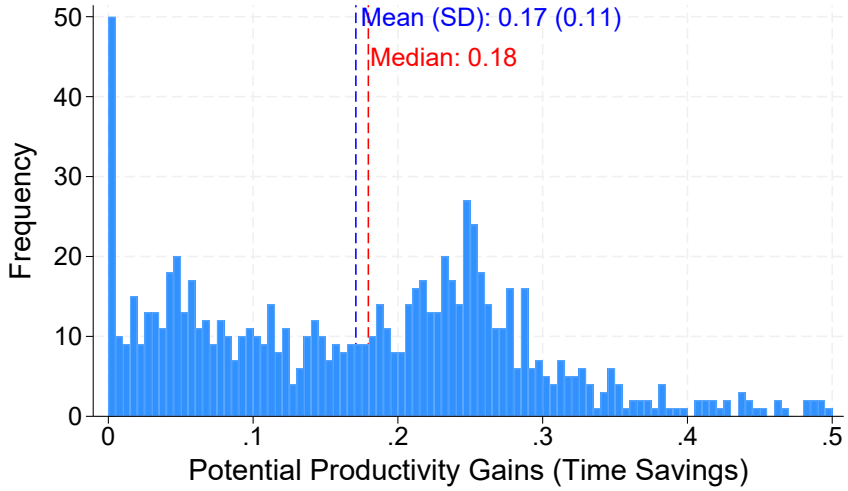

Figure 5: Potential Productivity Gains (PPG) Across Occupations

Notes: The figure shows the distribution of the occupation-level Potential Productivity Gains (PPG_o). PPG_o is defined in Equation (21) as 0.5 times the work-volume-weighted average of the task-level augmentation indicator β_{ot} from Eloundou et al. (2024), where task work volume is $Freq_{ot} \times Imp_{ot}$.

4.3 PSS and PPG Across the Occupational Distribution

This section summarizes how GenAI-induced changes in occupational entry barriers and productivity vary across occupations. Figure 6 reports bin-scatter plots of PSS and PPG against baseline expertise. Panel 6a shows that PSS increases with baseline expertise, measured as $(1 - S_o) \times 100$, across most of the distribution, implying that GenAI tends to

relax entry barriers more strongly for many initially higher-expertise occupations. However, the relationship is sharply non-monotone: above roughly the 90th percentile, PSS declines steeply, indicating comparatively small supply expansions in the extreme upper tail. This upper-tail decline reflects the thin right tail of the workforce expertise distribution. When an occupation's entry threshold lies far out in the upper tail, few workers sit close to the cutoff, so even sizable reductions in required expertise translate into limited expansions in the eligible pool.

Panel 6b plots PPG against baseline expertise. PPG rises strongly with baseline expertise over most of the distribution: occupations that sit higher in the expertise distribution are predicted to experience larger time savings from GenAI augmentation. The contrast between the two panels is striking: while both PSS and PPG rise with expertise across most of the distribution, the supply-side effect (PSS) reverses sharply in the upper tail whereas the demand-side effect (PPG) continues to increase.

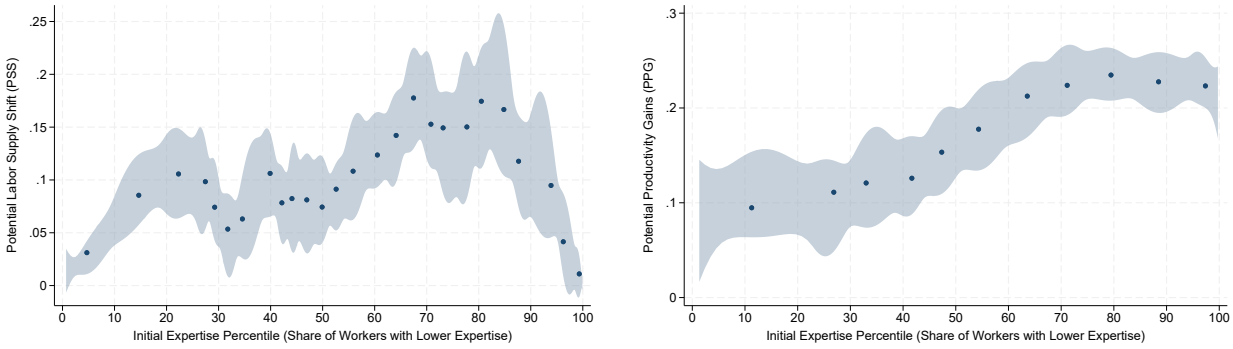

(a) PSS vs. Initial Expertise Percentile

(b) PPG vs. Initial Expertise Percentile

Figure 6: PSS and PPG vs. Baseline Expertise—Bin-scatter Plots

Notes: The figure reports bin-scatter plots at the occupation level. Occupations are sorted by the baseline expertise percentile, defined as $(1 - S_o) \times 100$, and grouped into equal-sized bins; each point plots the bin mean of the y-axis variable against the bin mean of baseline expertise. Panel (a) plots the Potential Supply Shift (PSS), defined as the change in the share of the workforce qualified to perform the occupation after GenAI integration. Panel (b) plots the Potential Productivity Gains (PPG_o), defined in Section 4.2.

These patterns imply heterogeneous wage effects across the occupational distribution. Lower-skilled occupations are predicted to experience relatively modest direct productivity gains, but may still be affected indirectly as reduced entry barriers expand access to more skilled work. Occupations in the middle and upper-middle of the expertise distribu-

tion combine substantial productivity gains with meaningful increases in contestability, as relaxed entry barriers erode scarcity premia. By contrast, the highest-expertise occupations are predicted to benefit strongly from productivity improvements while experiencing comparatively small supply responses, because reductions in expertise requirements admit few additional workers from the thin upper tail of the expertise distribution. This heterogeneity motivates the general equilibrium analysis that follows, which quantifies how productivity gains and expertise-driven supply expansions jointly shape wages and employment once workers reallocate across feasible occupations.

4.4 Benchmarking PSS Against Exposure Measures

A leading approach in the GenAI labor-market literature is to measure occupational "exposure" by classifying tasks as automatable and then computing the share of exposed tasks (e.g., ILO, 2025; Gmyrek et al., 2023; Eloundou et al., 2024; Felten et al., 2023). These exposure indices play a central role in both research and policy discussions of GenAI (e.g., Goldman Sachs, 2023; ILO, 2025; OECD, 2024; IMF, 2025; Brynjolfsson et al., 2025; Johnston and Makridis, 2025; The Budget Lab, 2025; Hosseini and Lichtinger, 2025). Because exposure remains the dominant benchmark, it is useful to benchmark it against PSS, even though this comparison is somewhat tangential to our main focus on distributional patterns in expertise and productivity.

Exposure measures are useful for summarizing how strongly GenAI applies to an occupation's tasks, but they are less well suited to our question, which concerns how GenAI changes occupational entry barriers and effective labor supply. For that purpose, exposure measures are incomplete for two reasons. First, they typically do not distinguish between impacts on high-expertise and low-expertise tasks, even though these have different implications for entry barriers and scarcity premia (see Autor and Thompson, 2025). Second, they often focus on extensive-margin automation and therefore miss intensive-margin effects through which GenAI reduces the expertise required to perform tasks that remain human-performed. Our PSS measure is designed to capture this supply-side channel by incorporating task expertise and combining extensive-margin automation with intensive-margin task simplification.

Figure 7 illustrates the relationship between PSS and an exposure measure, defined as

the share of automatable tasks within an occupation (see Section 4.1.2 for details on the automation measure). Panel (a) shows a positive association between exposure and PSS, but also substantial dispersion: even among occupations with similar exposure shares, PSS varies widely. Panel (b) shows that PSS rises with exposure at low to moderate exposure levels, but the relationship flattens among highly exposed occupations. Together, these patterns underscore that exposure alone is an incomplete statistic for characterizing GenAI-induced changes in occupational entry barriers and effective labor supply. To make the dispersion concrete, Appendix Figure A.7 reports the PSS of occupations with nearly identical aggregate exposure shares—between 29 and 31 percent (approximately the 75th percentile of the exposure distribution). Despite their nearly identical exposure levels, these occupations exhibit large differences in PSS. Treating them as equally exposed, as is common in the literature, obscures economically meaningful heterogeneity in how GenAI affects entry barriers and labor supply.

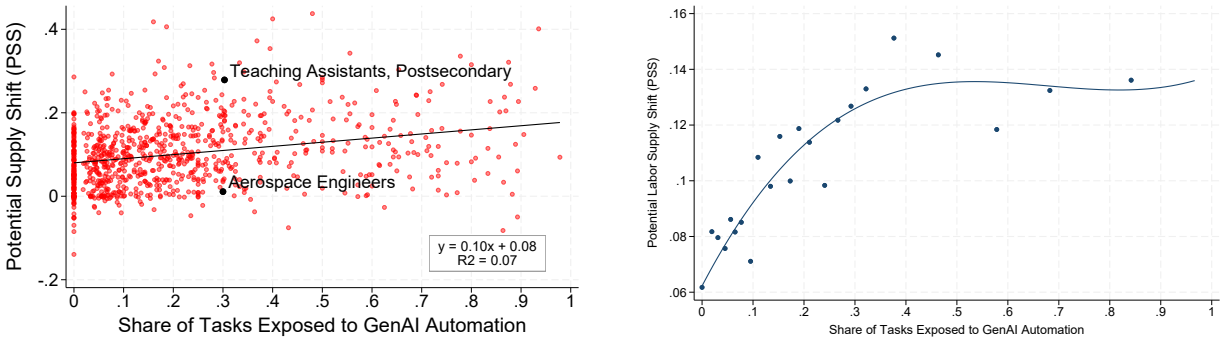

(a) Scatter plot

(b) Bin-scatter plot

Figure 7: Automation Exposure vs. PSS

Notes: The figure plots the occupation-level relationship between the share of exposed tasks and the potential supply shift (PSS) index. Panel (a) shows a scatter plot in which each dot represents a single occupation. Panel (b) presents a bin-scatter plot that groups occupations into equal-sized bins based on aggregate exposure and plots the mean PSS within each bin.

Consider, for example, Postsecondary Teaching Assistants and Aerospace Engineers. Both occupations have similar shares of exposed tasks, yet their predicted PSS differs sharply: Teaching Assistants exhibit a large PSS, close to 30 percentage points, while the PSS for Aerospace Engineers is close to zero. Several forces contribute to this difference. On the intensive margin, GenAI simplifies tasks in both occupations, reducing required expertise. On the extensive margin, however, GenAI is predicted to automate relatively

high-expertise tasks for Teaching Assistants, such as grading and course material preparation, which lowers the occupation's effective expertise requirement. For Aerospace Engineers, by contrast, exposure is concentrated in lower- and moderate-expertise tasks, such as documentation and record keeping, which raises the average expertise of the remaining task bundle and partially offsets intensive-margin effects. Initial expertise levels also matter: Aerospace Engineers sit far out in the upper tail of the expertise distribution, where worker density is low, so comparable reductions in required expertise translate into smaller increases in effective labor supply than for Teaching Assistants, who are closer to the center of the expertise distribution where density is higher. This interaction between task composition and the underlying expertise distribution explains why PSS can differ sharply even among occupations with similar exposure.

5 Inequality Implications

This section quantifies how GenAI affects the distribution of occupational wages through the two forces emphasized by the model: (i) *productivity gains*, which raise marginal revenue products and shift labor demand, and (ii) *expertise reductions*, which relax entry-barrier constraints and expand the pool of workers eligible for each occupation (the scarcity channel).

Section 5.1 embeds the empirically constructed PSS and PPG measures into the calibrated equilibrium system from Section 3.4 and solves for counterfactual wages numerically, allowing the full non-parametric structure of the occupation distribution to shape outcomes.[22] Appendix A.5 complements this analysis by taking the closed-form race condition for the expertise premium (Appendix A.4) directly to the data, estimating how productivity gains and supply-pool expansions vary with expertise across occupations and finding that the two race-condition forces are close in magnitude, with productivity slightly dominating.

[22]The local comparative statics in Section 3.5 provide intuition but are not designed for the potentially large and heterogeneous changes induced by GenAI. Large shocks can generate non-linear responses, and expertise reductions can shift the relevant scarcity margins by expanding feasible sets non-uniformly across the expertise distribution. For these reasons, we solve the full wage–employment system (Section 3.4) numerically.

5.1 Quantitative Counterfactual Exercise

5.1.1 Calibration and Implementation

We take the baseline (pre-AI) occupational wage and employment vectors from the O*NET data. To construct the distribution of worker expertise, we assign each occupation an expertise level as described in Section 4.1.1 and map all workers employed in that occupation to that expertise level. This yields an employment-weighted distribution of occupational expertise, which we use as a proxy for the underlying distribution of worker expertise in the baseline economy. While this construction abstracts from within-occupation heterogeneity in expertise, it preserves the cross-occupation structure that governs feasibility and scarcity in the model. It is worth noting that this proxy does not invert the model's latent expertise distribution F: in the model, observed employment reflects both expertise-based eligibility and idiosyncratic taste shocks across feasible occupations. As an alternative, Appendix A.17 constructs an expertise distribution from CPS ASEC 2024 educational attainment data. The two measures are highly correlated at the occupation level (Pearson correlation of 0.826 across 889 occupations), and our counterfactual results are robust, both qualitatively and quantitatively, to using this alternative measure.

We then solve for the model objects that rationalize these baseline wages and employment levels as an equilibrium under baseline feasibility. The two key parameters governing equilibrium sensitivity are the CES elasticity of substitution across occupational inputs, σ, and the dispersion parameter in occupational choice, τ. The vector of occupation-specific wedges $\{B_o\}$ is calibrated so that the observed baseline wage vector is an *exact* fixed point of the model given baseline employment and feasibility. This calibration ensures that counterfactual changes are interpreted as deviations from a correctly matched baseline equilibrium, rather than reflecting baseline misfit.[23] Table 2 summarizes the parameters used in the quantitative exercise, their roles in the model, and the moments or conditions used for calibration.

[23] Appendix A.17 shows that the main results are robust to enriching the model with occupation-specific amenities that match observed employment shares via a BLP-style contraction mapping.

Table 2: Calibration Targets and Parameter Values

Object	Role	Target (baseline)	Value / outcome
σ	Substitution across occupations	Standard value	5.0
τ	Choice dispersion (sorting / reallocation friction)	Baseline employment dispersion: $\mathrm{Var}(\log L_o)$	0.71
$\{B_o\}_{o\in\mathcal{O}}$	Occupation-specific wedges	GE fixed-point condition: $w_o = B_o\, L_o^{-1/\sigma}$	Calibrated occupation-by-occupation

Notes: B_o is chosen so that the observed baseline wage vector is an exact fixed point of the model under baseline feasibility.

5.1.2 The Effect of GenAI on the Wage Distribution

In this section, we evaluate three counterfactual environments: (i) expertise reductions only, in which GenAI relaxes feasibility constraints through intensive and/or extensive margins without affecting productivity; (ii) productivity gains only, in which occupation-specific productivity shifts $A_o \mapsto A_o e^{\pi_o}$ occur with feasibility held fixed; and (iii) combined effects, in which expertise reductions and productivity gains operate simultaneously. For each environment, we compute partial-equilibrium (PE) outcomes—which hold the broader wage structure fixed—and general-equilibrium (GE) outcomes—which allow workers to reallocate across occupations subject to feasibility constraints (i.e., workers can enter only occupations whose expertise requirements they meet) and idiosyncratic fit.

We begin by presenting aggregate changes in the wage distribution, and then turn to the incidence of wage changes across the pre-AI wage distribution to clarify the mechanisms behind these aggregate patterns.

Aggregate wage inequality. Table 3 reports standard inequality statistics, and Figure A.13 in Appendix A.18 plots the corresponding employment-weighted densities of occupational log wages.

Expertise reductions only. Relaxing expertise constraints is equalizing. In general equilibrium, the variance of log wages falls from 0.201 to 0.187, and upper-tail inequality declines substantially: the $p90-p50$ gap falls from 0.733 to 0.670. These effects reflect

reduced scarcity in high-expertise occupations as lower entry barriers expand the pool of eligible workers. General-equilibrium reallocation amplifies the compression relative to partial equilibrium, as newly eligible workers sort into previously scarce occupations.

Productivity gains only. When GenAI raises productivity without changing feasibility, wage dispersion increases. In general equilibrium, the variance of log wages rises to 0.253, and the $p90-p50$ gap increases to 0.857. Although worker reallocation dampens these effects relative to partial equilibrium, productivity gains remain strongly disequalizing, with larger wage increases concentrated in higher-wage occupations.

Joint counterfactual. When both expertise reductions and productivity gains operate simultaneously, the two forces partially offset. In general equilibrium, overall wage dispersion rises modestly relative to baseline: the variance of log wages increases from 0.201 to 0.230, while the $p90-p50$ gap rises from 0.733 to 0.789. Lower-tail inequality increases slightly, with the $p50-p10$ gap rising from 0.421 to 0.451. Overall, the joint outcome highlights the productivity–scarcity tradeoff: productivity gains push wages apart, while expertise reductions compress wage premia, yielding a net effect on inequality that is modest in aggregate but uneven across the distribution.[24]

Table 3: Wage Inequality Under Baseline and GenAI Counterfactuals

Scenario	$\mathrm{Var}(\log w)$	$p10$	$p50$	$p90$	$p90-p10$	$p50-p10$	$p90-p50$
Baseline	0.201	10.495	10.916	11.648	1.154	0.421	0.733
Supply PE	0.192	10.540	10.964	11.634	1.094	0.424	0.670
Supply GE	0.187	10.596	10.990	11.660	1.064	0.394	0.670
Prod PE	0.264	10.538	11.070	11.928	1.389	0.531	0.858
Prod GE	0.253	10.549	11.049	11.906	1.356	0.500	0.857
Both PE	0.247	10.631	11.115	11.952	1.322	0.485	0.837
Both GE	0.230	10.686	11.136	11.925	1.240	0.451	0.789

Notes: Employment-weighted wage moments. "Supply" corresponds to expertise reductions (scarcity channel); "Prod" corresponds to productivity gains (demand channel); "Both" combines both forces. PE holds the broader wage structure fixed; GE allows reallocation across feasible occupations.

Wage changes by initial wage percentile. Aggregate distributions summarize how inequality changes overall, but they do not show *where in the pre-AI wage distribution* gains

[24]These patterns are robust to both an alternative expertise distribution derived from CPS educational attainment and to the inclusion of occupation-specific amenities (Appendix A.17).

and losses occur. To make the model mechanism transparent—and to connect directly to the descriptive evidence on how PSS and PPG vary across the expertise and wage distributions (Section 4.3)—we plot wage changes against an occupation's position in the baseline wage distribution.

Figure 8 plots, on the x-axis, the occupation's percentile in the pre-AI wage distribution and, on the y-axis, the percentage change in wages in general equilibrium. The figure is shown separately for expertise reductions only and productivity gains only.

Expertise reductions (scarcity channel). Wage changes under expertise reductions are U-shaped across the initial wage distribution. Occupations in the lower tail experience wage gains as reduced entry barriers expand access to higher-paying jobs. In the middle of the distribution, wages fall as contestability increases and scarcity premia are eroded. At the top of the distribution, downward wage pressures diminish sharply and wages rebound toward zero, consistent with the descriptive finding that PSS is small in the extreme upper tail (Figure 6a): because the workforce expertise distribution is thin at very high thresholds, even sizable reductions in required expertise admit relatively few newly eligible workers. This U-shaped pattern is robust to the choice of the aggregation exponent ρ (Appendix A.2, Figure A.3).

Productivity gains. In contrast, productivity gains generate an increasing profile of wage changes across the initial wage distribution. This pattern mirrors the descriptive evidence that PPG is higher for more expert—and hence higher-wage—occupations (Figure 6b), so demand shifts are concentrated toward the upper part of the distribution. The productivity channel is therefore strongly disequalizing in isolation.

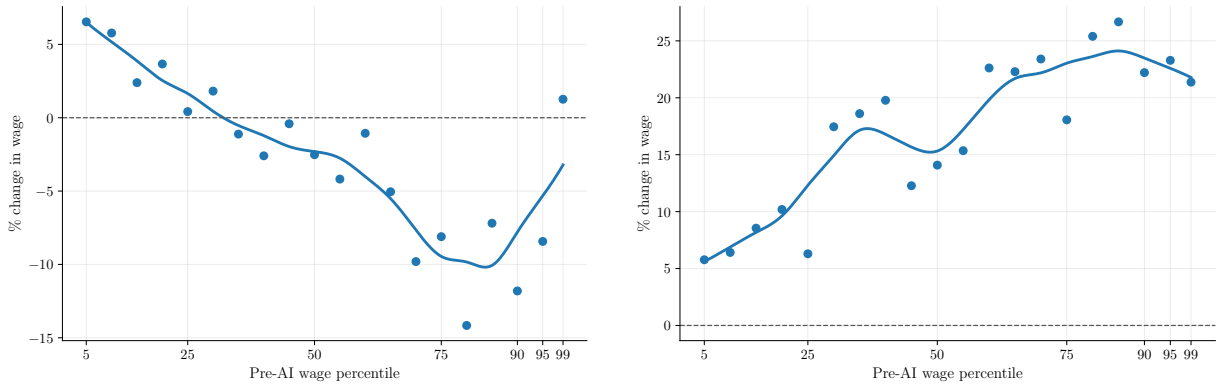

(a) Expertise Reductions Only

(b) Productivity Gains Only

Figure 8: Wage Changes in GE by Pre-AI Wage Percentile

Notes: Each panel plots the percentage change in occupational wages (general equilibrium) against the occupation's percentile in the baseline wage distribution. Points are employment-weighted averages within percentile bins of the pre-AI wage distribution. The two panels correspond to (i) expertise reductions only (scarcity channel), and (ii) productivity gains only (demand channel).

5.1.3 GE Responses and the Role of PSS

This subsection examines how much of the general-equilibrium adjustment driven by expertise reductions can be explained by the Potential Supply Shift (PSS). Recall that PSS is a deliberately simple, reduced-form measure: it captures how GenAI changes occupational entry barriers and the size of the eligible workforce, but it abstracts from the full structure of endogenous occupational choice, outside options, and general-equilibrium feedbacks embedded in the model. In this sense, PSS is a coarse summary statistic of the supply-side shock rather than a structural object.

Despite this simplicity, PSS is highly informative about equilibrium outcomes. Focusing on the *supply-only* counterfactual, Figures 9a and 9b plot general-equilibrium changes in occupational wages and employment against PSS. Occupations with larger expertise-driven expansions of effective labor supply experience systematically larger declines in wages and larger increases in employment. The relationship is strong and monotone, with PSS explaining a substantial share of cross-occupation variation in general-equilibrium responses ($R^2 \approx 0.54$ in both regressions). These patterns emerge even though the underlying equilibrium incorporates endogenous worker reallocation, wages, and outside op-

tions. That PSS—a deliberately simple, reduced-form measure—accounts for so much of this cross-occupation heterogeneity highlights its value as a summary statistic for GenAI's supply-side effects, one that can be computed directly from task-level data without solving for equilibrium.

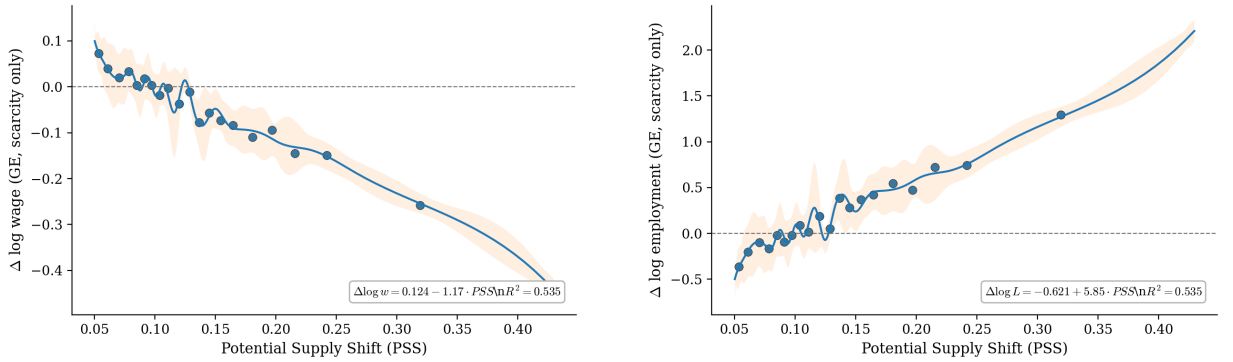

(a) GE Wage Changes vs. PSS (Supply Only) **(b)** GE Employ. Changes vs. PSS (Supply Only)

Figure 9: General-Equilibrium Supply Responses and Potential Supply Shift

Notes: Each panel plots binned scatterplots of general-equilibrium changes in occupational wages and employment against the Potential Supply Shift (PSS), focusing on the supply-only counterfactual. Changes are measured relative to the baseline equilibrium. Bins are employment-weighted. Occupations with zero PSS are excluded.

6 Dynamic Extension

6.1 Motivation

The static model in Sections 3–5 treats occupational entry barriers as one-dimensional expertise thresholds: any worker whose expertise exceeds the threshold can enter, regardless of their current occupation. This abstraction misses two features of real labor markets. First, switching costs in practice are pair-specific—moving from nursing to medicine is far easier than moving from nursing to law, even when the destination occupations require similar levels of overall expertise. Second, the static framework can only compare steady states before and after an AI shock; it is silent on how the equalizing effects of reduced entry barriers unfold over time. This section addresses both limitations with a dynamic model in which forward-looking workers face pairwise retraining distances that differ

across origin-destination pairs. The AI shock is modeled as a reduction in these pairwise retraining distances that varies across occupation pairs. The model then delivers a transition path that traces worker reallocation across occupations and the resulting implications for inequality over time.

6.2 Model

6.2.1 Dynamic Occupational Choice

Time is discrete and measured in years. Each worker is characterized by their current occupation $s \in \{1, \ldots, J\}$. Workers are infinitely lived and discount the future at rate $\beta \in (0, 1)$. A worker employed in occupation o earns the market wage w_o; there is no within-occupation ability heterogeneity, consistent with the static model.

Each period, a worker in occupation s chooses a destination o by solving the Bellman equation

$$V(s) \;=\; \tau \, \log \sum_{o=1}^{J} \exp\left[\frac{\log w_o \;-\; c_{so}\,\mathbf{1}_{o \neq s} \;+\; a_o \;+\; \beta\, V(o)}{\tau}\right], \tag{22}$$

where a_o is a time-invariant amenity, $\tau > 0$ is the dispersion of i.i.d. Type-I extreme value taste shocks, and c_{so} is the disutility of switching from occupation s to occupation o.

Heterogeneous switching costs. We decompose c_{so} into an origin-specific *stickiness* component and a destination-specific *barrier* component, both scaled by the retraining distance:

$$c_{so} \;=\; \left(\kappa_s^{\text{out}} + \kappa_o^{\text{in}}\right) d_{so}. \tag{23}$$

The origin parameter $\kappa_s^{\text{out}} \geq 0$ captures how costly it is to leave occupation s; it rationalizes the fact that some occupations (e.g. those with strong specific human capital or tenure attachment) see low outflow rates. The destination parameter $\kappa_o^{\text{in}} \geq 0$ captures how costly it is to enter o; it rationalizes the fact that some occupations (e.g. licensed professions) see low inflow rates from many origins. Allowing these two dimensions of heterogeneity lets the model match the full cross-occupation structure of observed mobility—who moves, from where, and to where—rather than only the aggregate mobility rate, and yields counterfactual predictions interpretable at the level of individual occupations.

The logit transition probability from s to o is

$$\pi(o \mid s) \;=\; \frac{\exp\big[v(o \mid s)/\tau\big]}{\sum_{o'=1}^{J} \exp\big[v(o' \mid s)/\tau\big]}, \qquad v(o \mid s) \;\equiv\; \log w_o - c_{so}\,\mathbf{1}_{o\neq s} + a_o + \beta\, V(o). \quad (24)$$

The stationary distribution $\mu(o) = \sum_s \pi(o \mid s)\,\mu(s)$ is the ergodic distribution of the resulting Markov chain. Labor supply in each occupation equals $L_o = \mu(o)$, and wages satisfy the same CES inverse demand as in the static model: $w_o = B_o \cdot L_o^{-1/\sigma}$.

6.2.2 AI Shock

We model the supply-side effects of AI as a permanent, unexpected change in pairwise retraining distances, $d_{so} \to d_{so}^{\text{AI}}$. Lower values of d_{so}^{AI}—as is the case for most occupation pairs—reflect the idea that access to GenAI makes it easier for workers in origin occupation s to acquire the skills required to enter destination o. By assumption, the switching-cost parameters $(\boldsymbol{\kappa}^{\text{out}}, \boldsymbol{\kappa}^{\text{in}})$ are structural and held fixed across pre- and post-AI equilibria, along with τ, productivity shifters $\{B_o\}$, and amenities $\{a_o\}$. The single channel through which the shock propagates is therefore the reduction in pairwise retraining distances.

This shock is the dynamic analogue of the supply-side mechanism in the static model. In the static framework, GenAI lowers occupation-specific expertise requirements and expands the set of workers qualified to perform a given occupation. In the dynamic setting, the same force operates through bilateral retraining frictions: GenAI reduces the time required to retrain across occupation pairs, making some transitions easier than others. The dynamic model therefore captures the same entry-barrier erosion mechanism as the static model, but in a richer environment in which accessibility depends on the worker's origin occupation rather than on a single common expertise threshold.

6.3 Data Inputs

6.3.1 Pre-AI Retraining Distance Matrix

The first input to the dynamic model is a directed retraining distance matrix d_{so} that measures, for each ordered pair of 3-digit SOC occupations (s, o), the time (in years) required

for a worker currently in occupation s to retrain for occupation o. We construct this matrix using LLM-based evaluations, prompting the LLM (using GPT-5 Nano) to assess the retraining time for each occupation pair based on the task bundles of the origin and destination occupations in the O*NET database.

The retraining matrix is asymmetric by construction: $d_{so} \neq d_{os}$ in general. Switching *into* a high-skill occupation (e.g., from nursing assistant to physician) involves a large retraining burden, while switching *out* involves a much smaller one, because the high-skill worker already possesses most of the skills needed in the lower-requirement occupation. The diagonal is zero ($d_{oo} = 0$). We aggregate from 6-digit SOC-level estimates to the 3-digit SOC level ($J = 94$ occupation groups) by averaging across detailed occupation pairs within each 3-digit group.

Validation against O*NET data. To validate the LLM-based retraining estimates, we construct an independent benchmark using O*NET occupational skill and knowledge data. For each occupation, O*NET reports proficiency levels on 35 skill dimensions and 33 knowledge dimensions (68 total). We compute a directed skill-shortfall measure for each occupation pair:

$$d_{so}^{\text{O*NET}} = \sum_{k=1}^{68} \max\bigl(0,\ v_o^k - v_s^k\bigr), \tag{25}$$

where v_o^k is the O*NET Level rating for occupation o on dimension k. This captures the additional proficiency a worker in s would need to acquire across all dimensions to meet the requirements of o. Like d_{so}, this measure is non-negative, zero on the diagonal, and asymmetric. We rescale $d_{so}^{\text{O*NET}}$ so that the mean off-diagonal entry equals 2 years.

Figure 10 plots the LLM-based retraining estimates against the O*NET skill-shortfall measure across all $J(J-1) = 8{,}742$ off-diagonal occupation pairs. The two measures are positively correlated, with Pearson $r = 0.49$ and Spearman $\rho = 0.65$. This correlation is notable given that the two measures are constructed from entirely different data: the LLM estimates reflect holistic assessments of retraining time, while the O*NET measure is a mechanical aggregation of skill-level gaps. The strong rank correlation confirms that the LLM-based retraining matrix captures the ordinal structure of occupational distances as reflected in skill requirements.

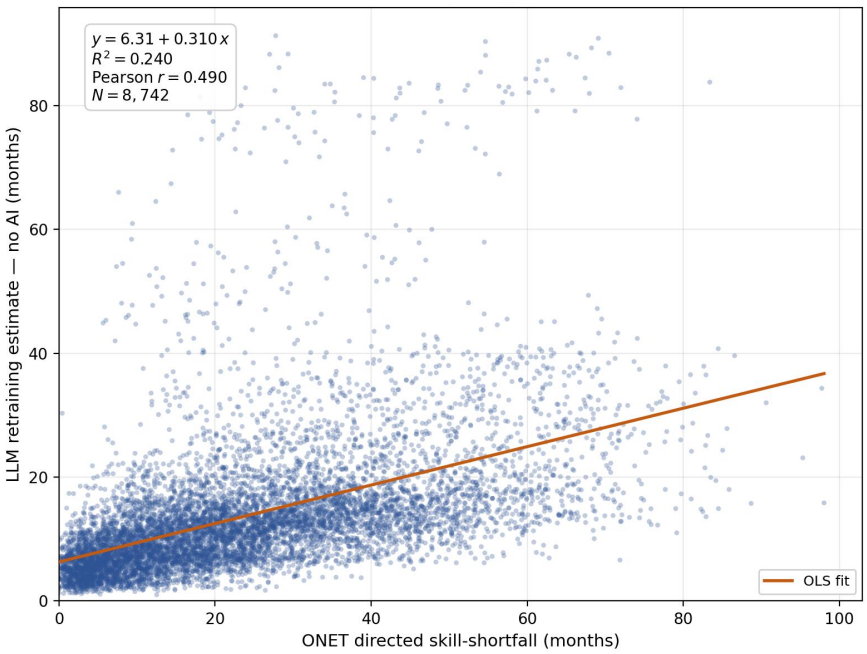

Figure 10: Validation: LLM Retraining Estimates vs. O*NET Skill Shortfall

Notes: Each point represents an ordered pair of 3-digit SOC occupations ($N = 8{,}742$ off-diagonal pairs). The x-axis is the O*NET directed skill-shortfall $d_{so}^{\text{O*NET}}$ (in months) and the y-axis is the LLM-based retraining estimate (in months). Orange lines are OLS fits.

6.3.2 Post-AI Retraining Distances

We construct the post-AI retraining distance matrix d_{so}^{AI} by re-estimating pairwise retraining times using the same LLM-based procedure as in Section 6.3.1, under the counterfactual that workers have access to a capable GenAI assistant. The employment-weighted mean of the off-diagonal retraining distance falls from 1.15 years pre-AI to 0.89 years post-AI, a reduction of approximately 23 percent.

6.3.3 Occupation-to-Occupation Flow Data

The third input is observed occupational mobility. We use the Occupational Transitions Public Data Set of Schubert et al. (2024), which records origin–destination flow shares for occupational switchers at the 6-digit SOC level based on a large longitudinal sample of U.S. workers. For each origin occupation, the file reports conditional transition shares $\widehat{\pi}(o \mid s, \text{switch})$ to each destination that sum to one, along with the total count of switchers

from that origin. We aggregate the 6-digit data to the model's SOC3 universe by summing switch counts and dropping within-SOC3 moves, which at our level of aggregation are "stay" events.

Two occupation-specific moments emerge from this aggregation and will be used to discipline the model: the origin-specific annual outflow rate m_s^{data} (the share of workers in occupation s who leave it over a year), and the share of switchers landing in each destination, $f_o^{\text{data}} = \sum_s \widehat{\pi}(o \mid s, \text{switch}) \cdot L_s^{\text{data}} / \sum_{s'} L_{s'}^{\text{data}}$, weighted by data employment at the origin.

6.4 Calibration

Table 4 summarizes the calibration. The elasticity of substitution ($\sigma = 5$) and discount factor ($\beta = 0.95$) are set externally. The retraining distance matrix d_{so} and the post-AI matrix d_{so}^{AI} are constructed as described in Section 6.3. The pre-AI steady state is pinned down by two sequential inversions. First, amenities $\{a_o\}$ are recovered via a BLP-style contraction mapping that matches model-implied employment shares to data.[25] Second, given the model-implied labor supply, productivity shifters are backed out from $B_o = w_o^{\text{data}} \cdot L_o^{1/\sigma}$. Together, these inversions ensure that the pre-AI steady state exactly reproduces the observed wage and employment distributions. The taste-shock dispersion τ is estimated by the simulated method of moments targeting the cross-occupation variance of log wages. The $2J$ heterogeneous switching-cost parameters $(\boldsymbol{\kappa}^{\text{out}}, \boldsymbol{\kappa}^{\text{in}})$ are calibrated by a fixed-point iteration that targets two vectors of moments: the J origin-specific annual outflow rates $\{m_s^{\text{data}}\}$ and the J destination inflow shares among switchers $\{f_o^{\text{data}}\}$, giving a just-identified system. The calibrated model matches these flow moments essentially exactly: the correlation between model-implied and observed destination inflow shares is $+0.998$, and the correlation between model-implied and observed origin outflow rates is $+0.984$.

[25] Because the present value of a permanent flow amenity a_o is amplified by the factor $1/(1-\beta)$ through the Bellman recursion, the inversion applies a $(1-\beta)$ correction to translate between flow and continuation values.

Table 4: Dynamic Model: Calibrated Parameters

Parameter	Role	Value	Source / Target
β	Discount factor	0.95	Standard
σ	CES elasticity	5.0	Same as static model
τ	Choice dispersion	0.463	SMM: $\text{Var}(\log w_o)$
$\boldsymbol{\kappa}^{\text{out}}$	Origin stickiness (J)	calibrated	Origin outflow rates $\{m_s^{\text{data}}\}$
$\boldsymbol{\kappa}^{\text{in}}$	Destination barriers (J)	calibrated	Destination inflow shares $\{f_o^{\text{data}}\}$

Notes: The $2J$ occupation-specific cost parameters are jointly calibrated to the occupation-to-occupation flow moments of Schubert et al. (2024). At the calibrated parameters, the correlation between model and data inflow shares is $+0.998$ and between model and data outflow rates is $+0.984$.

6.5 Results

6.5.1 Long-Run Steady-State Comparison

Table 5 reports the long-run comparison of pre- and post-AI steady states. Between-occupation wage inequality declines across all measures: the variance of log wages falls from 0.208 to 0.190 (-8.7 percent), the $p90-p50$ gap contracts from 0.840 to 0.735 (-12.5 percent), and the $p90-p10$ gap narrows from 1.228 to 1.110 (-9.6 percent). The compression is particularly pronounced in the upper half of the distribution, consistent with the erosion of scarcity premia in high-wage occupations as reduced switching costs expand the pool of potential entrants. The magnitude of this decline is broadly consistent with the supply-only counterfactual in the static framework, where expertise reductions alone reduce the variance of log wages by approximately 7 percent (Table 3).

Table 5: Dynamic AI Counterfactual: Long-Run Steady-State Comparison

	$\text{Var}(\log w)$	$p90-p50$	$p90-p10$
Pre-AI steady state	0.208	0.840	1.228
Post-AI steady state	0.190	0.735	1.110
% change	-8.7%	-12.5%	-9.6%

Notes: Long-run (steady-state) comparison of the pre-AI and post-AI equilibria under the calibrated two-way switching-cost specification. $p90-p50$ and $p90-p10$ are employment-weighted log-wage percentile gaps computed using data employment shares, consistent with the static model's inequality statistics.

6.5.2 Transition Dynamics

A central advantage of the dynamic model over the static framework is that it tells us how the adjustment unfolds in time. Figure 11 plots the transition path of $\mathrm{Var}(\log w_o)$ following the shock.[26] The compression begins immediately—as soon as lower retraining distances make new transitions attractive—and proceeds rapidly. Half of the long-run decline in variance is realized within five years, eighty percent within nine years, and essentially all of it within roughly twenty-five years. The speed of adjustment is governed by the ratio of switching costs to idiosyncratic taste-shock dispersion, which is disciplined jointly by τ and by the heterogeneous κ vectors that match the observed cross-occupation flow matrix.

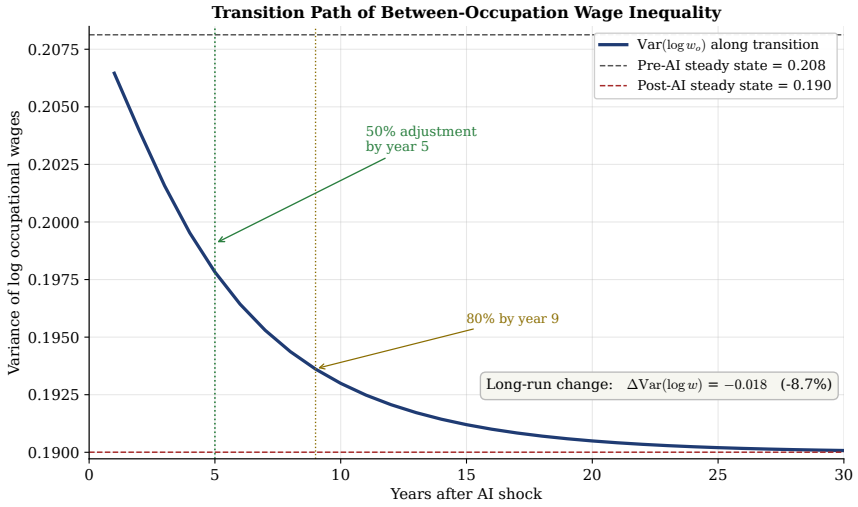

Figure 11: Transition Path of Between-Occupation Wage Inequality After the AI Shock

Notes: The figure plots the variance of log occupational wages (weighted by employment shares) along the transition path from the pre-AI to the post-AI steady state. The AI shock reduces pairwise retraining distances d_{so}, lowering occupational switching costs. Productivity shifters $\{B_o\}$, amenities $\{a_o\}$, and the switching-cost vectors $(\boldsymbol{\kappa}^{\text{out}}, \boldsymbol{\kappa}^{\text{in}})$ are held fixed.

[26] We solve the transition path between the pre- and post-AI steady states by backward induction on the value function given a guessed wage path, forward propagation of the distribution μ_t from μ_{pre}, and damped iteration on wages until the labor-market clearing condition $w_{o,t} = B_o L_{o,t}^{-1/\sigma}$ holds at every period.

6.5.3 Who Moves and Where

Because the model matches the joint distribution of observed flows, the post-AI counterfactual is interpretable at the level of individual occupations. Figure 12 reports the thirty occupations with the largest net employment inflows and the thirty with the largest net outflows under the long-run counterfactual. Three patterns emerge.

First, the expanding occupations are concentrated in high-wage cognitive occupations: Computer Occupations (+1.9M workers, +40 percent), Other Management (+1.7M, +50 percent), Lawyers (+1.1M, +136 percent), Financial Specialists (+1.0M, +34 percent), Marketing and Sales Managers (+0.7M, +64 percent), and Mathematical Science Occupations (+0.6M, +156 percent).[27] Second, wages fall in every expanding occupation, with declines ranging from roughly zero to twenty-four percent. This is the scarcity channel operating through reallocation: when AI reduces the retraining distance into a high-wage occupation from many origins, the scarcity premium the occupation commands erodes, even without any change in demand or productivity. Third, the contracting occupations are concentrated in low-wage service and manual jobs: Material Moving (−1.8M, −23 percent), Retail Sales (−1.6M), Food Service (−1.6M), Home Health Aides (−1.0M), Motor Vehicle Operators (−0.7M). In these occupations, wages generally rise modestly, by roughly 1 to 5 percent, because worker outflow tightens local supply.

Importantly, these results isolate the supply-side channel: they reflect the effects of GenAI on pairwise retraining distances, holding productivity, task sets, and demand conditions fixed. In practice, the net effect of GenAI on any occupation will also depend on demand-side forces—automation of existing tasks, productivity gains in retained tasks, and new task creation—that the dynamic model abstracts from.

Takeaways. The dynamic model yields three main conclusions. First, the direction of the supply-side channel identified in the static framework is robust: reductions in pairwise retraining distances compress between-occupation wage inequality in the long run, with $\mathrm{Var}(\log w_o)$ falling by roughly 9 percent. Second, the reallocation is highly concentrated: high-wage cognitive occupations absorb large inflows and experience meaningful

[27] These large percentage changes reflect the long-run steady-state comparison and should not be read as short-run forecasts; the transition dynamics above indicate that adjustments materialize gradually over years.

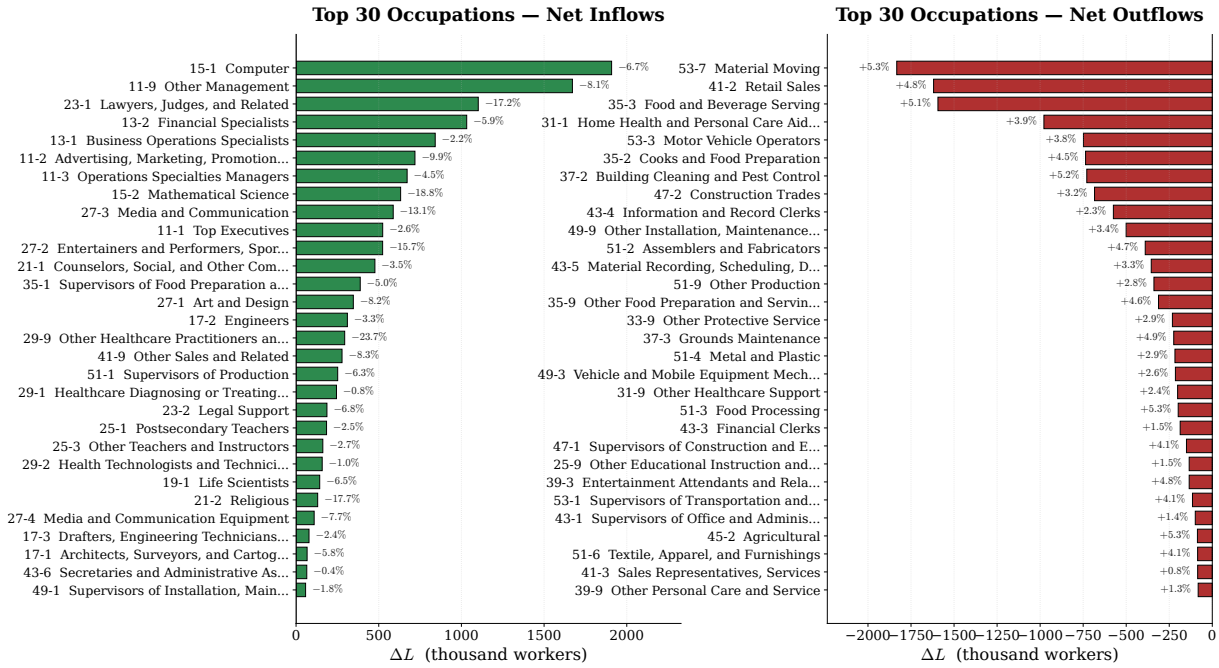

Figure 12: Occupation-Level Employment Reallocation Under the AI Counterfactual

Notes: Left panel: the thirty occupations with the largest net employment inflows, comparing pre- and post-AI long-run steady states. Right panel: the thirty occupations with the largest net outflows. Bars report ΔL in thousands of workers; the label next to each bar reports the percentage change in the occupation's predicted wage.

wage declines, while low-wage service occupations shed workers and see modest wage gains. Third, whereas the static framework can only compare steady states before and after the shock, the dynamic model traces the convergence path: half of the long-run compression materializes within five years and eighty percent within nine.

7 Conclusion

How will generative AI reshape wage inequality across occupations? This paper argues that an important part of the answer operates through a productivity–scarcity race. GenAI can raise wages by augmenting labor productivity, but it can also compress wage premia by lowering expertise-based entry barriers and expanding the set of workers who can perform high-expertise work. The distributional effects of GenAI depend on the relative strength of these channels and on where entry barriers are relaxed in the workforce

expertise distribution.

To capture the supply-side channel, we introduce and make publicly available a Potential Supply Shift (PSS) index, which maps task-level GenAI effects into changes in the share of the workforce qualified to perform each occupation. Using O*NET task data for 923 occupations, we combine LLM-based evaluations with a task-based framework that distinguishes two mechanisms through which GenAI alters expertise requirements. On the extensive margin, some tasks are automated and removed from labor's task bundle. On the intensive margin, tasks remain human-performed but become easier, reducing the expertise required to perform them. We validate this distinction using real-world AI usage data from the Anthropic Economic Index, which confirms that the two margins correspond to different patterns of human–AI interaction. Aggregating these effects yields a post-GenAI expertise threshold for each occupation. PSS then maps the implied change in expertise thresholds into the employment-weighted distribution of baseline expertise, capturing how many workers become newly eligible to work in each occupation.

We document three key facts. First, PSS increases with baseline expertise and wages across most of the distribution, but falls sharply in the extreme upper tail. This pattern is consistent with the thinness of the expertise distribution at very high thresholds, so comparable reductions in required expertise translate into only limited inflows of newly qualified workers. Second, PSS is positively associated with predicted productivity gains, implying that many upper-middle occupations combine sizable time savings with sizable reductions in entry barriers. Third, PSS varies widely even among occupations with nearly identical automation-exposure shares. This suggests that exposure measures, while informative about automatable task shares, are weak predictors of how GenAI reshapes occupational entry barriers and effective labor supply.

To interpret these patterns, we develop a general equilibrium model with discrete tasks, CES demand, and occupational choice with outside options. The model clarifies how GenAI affects wages through two opposing forces: productivity gains raise marginal revenue products and tend to widen the wage distribution, while expertise-driven supply expansions increase contestability and compress scarcity premia. Under additional distributional assumptions, we derive in Appendix A.4 closed-form expressions for the post-GenAI wage gradient and a sharp race condition: the wage structure steepens if and only if differential productivity gains, amplified by the elasticity of substitution, ex-

ceed the rate at which entry-barrier erosion expands the eligible pool across the expertise distribution. Quantitative counterfactuals show that expertise expansion alone is equalizing, while productivity gains alone increase dispersion. When both channels operate together, these forces partially offset. Under our baseline calibration, the net effect on cross-occupation wage inequality is modest, though the magnitude of each channel depends on assumptions about the strength and timing of GenAI's effects.

We complement the static analysis with a dynamic extension in which forward-looking workers face pairwise occupational switching costs, replacing the static model's one-dimensional expertise threshold with a richer occupation-to-occupation retraining cost matrix. Using LLM-based retraining estimates validated against O*NET skill gaps, we model the AI shock as a reduction in pairwise retraining distances. In the calibrated dynamic model, AI-driven reductions in retraining distances compress between-occupation wage inequality by approximately 9 percent, consistent with the equalizing direction of the scarcity channel in the static framework. Importantly, whereas the static model can only compare steady states before and after the shock, the dynamic model reveals that the equalizing effect begins immediately but materializes gradually: roughly 80 percent of the long-run inequality decline is realized within the first nine years, with the remaining adjustment occurring slowly as the economy approaches its new steady state.

Several limitations apply. The LLM-based counterfactual assessments—particularly how GenAI changes task expertise—are difficult to validate directly and should be interpreted cautiously. We hold the set of tasks and occupations fixed, abstracting from endogenous task creation (e.g., Acemoglu and Restrepo, 2019; Autor et al., 2024). More generally, realized impacts will depend on the pace and extent of adoption, which we do not model.

Overall, the results highlight that GenAI's distributional effects cannot be understood through productivity effects alone. If productivity gains dominate, GenAI is more likely to increase wage inequality by reinforcing returns in high-wage occupations. If entry-barrier erosion substantially expands access to high-expertise work, GenAI is more likely to reduce inequality by lowering occupational scarcity premia. The balance between these forces—the race between productivity and scarcity—will determine how GenAI reshapes the wage structure across occupations.

References

Acemoglu, D. (2025). The Simple Macroeconomics of AI. *Economic Policy*, 40(121):13–58.

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2022). Tasks, Automation, and the Rise in US Wage Inequality. *Econometrica*, 90(5):1973–2016.

Althoff, L. and Reichardt, H. (2026). Task-Specific Technical Change and Comparative Advantage. *CESifo Working Paper*.

Appel, H., Berger, B., Connors, J., Curran, T., and Ganguli, D. (2026). The Anthropic Economic Index. Research report, Anthropic. Version 4.

Artuç, E., Chaudhuri, S., and McLaren, J. (2010). Trade Shocks and Labor Adjustment: A Structural Empirical Approach. *American Economic Review*, 100(3):1008–1045.

Autor, D. (2024). Applying AI to Rebuild Middle Class Jobs. Technical report, National Bureau of Economic Research.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, 139(3):1399–1465.

Autor, D. and Kausik, B. (2026). Resolving the automation paradox: falling labor share, rising wages. *arXiv preprint arXiv:2601.06343*.

Autor, D., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Autor, D. and Thompson, N. (2025). Expertise. NBER Working Paper 33941, National Bureau of Economic Research.

Brynjolfsson, E., Chandar, B., and Chen, R. (2025). Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Working paper. Latest version available at https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica*, 87(3):741–835.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics*, 36(S1):S13–S70.

Cruces, G., Meijide, D. F., Galiani, S., Gálvez, R. H., and Lombardi, M. (2026). Does Generative AI Narrow Education-Based Productivity Gaps? Evidence from a Randomized Experiment. Technical report, National Bureau of Economic Research.

Danieli, O. (2024). Skill-Replacing Technology and Bottom-Half Inequality. Technical report, Working Paper.

de Souza, G. (2026). AI in the Office and the Factory: Evidence from Administrative Software Registry Data.

Dix-Carneiro, R. (2014). Trade Liberalization and Labor Market Dynamics. *Econometrica*, 82(3):825–885.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs Are GPTs: Labor Market Impact Potential of LLMs. *Science*, 384(6702):1306–1308.

Felten, E. W., Raj, M., and Seamans, R. (2023). Occupational Heterogeneity in Exposure to Generative AI. *Available at SSRN 4414065*.

Freeman, R. B. (1975). Overinvestment in College Training? *Journal of Human Resources*, 10(3):287–311.

Freund, L. B. and Mann, L. F. (2025). Job Transformation, Specialization, and the Labor Market Effects of AI. *CESifo Working Paper*.

Garicano, Luis and Li, Jin and Wu, Yanhui (2026). Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries. Working Paper.

Gmyrek, P., Berg, J., and Bescond, D. (2023). Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality. *ILO Working Paper*, 96.

Goldin, C. and Katz, L. F. (2008). *The Race between Education and Technology*. Harvard University Press, Cambridge, MA.

Goldman Sachs (2023). The Potentially Large Effects of Artificial Intelligence on Economic Growth. *Goldman Sachs*, 1(5):268–296.

Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial Intelligence and the Labor Market. Technical report, National Bureau of Economic Research.

Hosseini, S. M. and Lichtinger, G. (2025). Generative AI as Seniority-Biased Technological Change: Evidence From U.S. Resume and Job Posting Data. *Available at SSRN*.

ILO (2025). *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*. Number 140. ILO Working Paper.

IMF (2025). Global Impact of AI: Mind the Gap. Technical report, International Monetary Fund.

Johnston, A. and Makridis, C. (2025). The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis of AI Exposure. *Available at SSRN 5375017*.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963–1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Klein Teeselink, B. (2025). Generative AI and Labor Market Outcomes: Evidence from the United Kingdom. *Available at SSRN*.

Klein Teeselink, B. and Carey, D. (2026). AI, Automation, and Expertise. SSRN working paper.

Massenkoff, M. and McCrory, P. (2026). Labor Market Impacts of AI: A New Measure and Early Evidence.

OECD (2024). Miracle or Myth? Assessing the Macroeconomic Productivity Gains from Artificial Intelligence. Technical report, OECD Publishing.

O*NET (2023). O*NET database version 27.2. https://www.onetcenter.org/database.html. Accessed: 2026.

Schubert, G., Stansbury, A., and Taska, B. (2024). Employer concentration and outside options. *Working Paper*. Occupational Transitions Public Data Set.

The Budget Lab (2025). Evaluating the Impact of AI on the Labor Market: The Current State of Affairs. Accessed: 11-27-2025.

Tinbergen, J. (1974). Substitution of Graduate by Other Labour. *Kyklos*, 27(2):217–226.

Traiberman, S. (2019). Occupations and Import Competition: Evidence from Denmark. *American Economic Review*, 109(12):4260–4301.

A Supplemental Appendix

A.1 Task-Based Microfoundation: Hierarchical Feasibility as a Limiting Case

This appendix develops the hierarchical-feasibility ($\rho \to \infty$) special case of the power-mean entry barrier introduced in Section 3.1. Under hierarchical feasibility, the occupation's binding expertise requirement is determined by its single most demanding task.

Hierarchical feasibility. Assume that a worker with expertise e can perform task (o, t) if and only if $e \geq r_{ot}$. Production in occupation o requires completing *all* tasks assigned to labor, so the occupation's binding expertise requirement is

$$R_o^{\max} := \max_{t \in T_o} r_{ot}. \tag{26}$$

This corresponds to $\rho \to \infty$ in the general power-mean aggregator (1).

GenAI and changes in the binding requirement. Under hierarchical feasibility, GenAI affects the binding requirement through both margins described in Section 3.4.

On the *extensive margin*, GenAI automates a subset of tasks. Let $T_o^L := \{t \in T_o : a_{ot} = 0\}$ denote the tasks that remain assigned to labor. The post-automation binding requirement is

$$R_o^{\text{ext}} := \max_{t \in T_o^L} r_{ot}. \tag{27}$$

Automating the most demanding task can produce a discrete drop in the entry barrier; automating a less demanding task leaves the barrier unchanged.

On the *intensive margin*, tasks remain human-performed but access to GenAI lowers the expertise required. Let $\widetilde{r}_{ot} \leq r_{ot}$ denote the GenAI-assisted requirement for task t. Combining both margins:

$$\widetilde{R}_o^{\max} := \max_{t \in T_o^L} \widetilde{r}_{ot}. \tag{28}$$

Relation to the general power-mean aggregation. Equations (26)–(28) are the $\rho \to \infty$ limits of the general power-mean formulas (1) and (12) used in the main model. The hierarchical-feasibility case provides a clean theoretical benchmark in which a single task is binding, while the finite-ρ specification allows the full task distribution to matter. In the empirical implementation (Section 4.1), we set $\rho = 1$, which yields the weighted arithmetic mean. The general power mean nests both as special cases and allows intermediate specifications that place progressively more weight on high-expertise tasks as ρ increases above 1.

A.2 Sensitivity to the Aggregation Exponent ρ

The baseline analysis sets $\rho = 1$ in Equation (1), so that an occupation's entry barrier is the importance-weighted arithmetic mean of its task-level expertise requirements. This appendix examines the sensitivity of both descriptive and general-equilibrium results to the choice of ρ, considering values $\rho \in \{1, 2, 3, 5, 10\}$ that span from the arithmetic mean to a near-maximum aggregator.

A.2.1 Occupation-Level Expertise Rankings

Because the power mean is monotonically increasing in ρ, raising the aggregation exponent shifts the level of every occupation's entry barrier upward. The economically relevant question is whether the *ranking* of occupations changes materially. Figure A.1 reports pairwise Pearson and Spearman rank correlations of occupation-level expertise R_o across the five values of ρ.

All pairwise correlations exceed 0.88 (Pearson) and 0.94 (Spearman rank). Adjacent values—for example, $\rho = 1$ versus $\rho = 2$—correlate above 0.99. Even the most distant pair ($\rho = 1$ versus $\rho = 10$) exhibits a Pearson correlation of 0.88 and a Spearman rank correlation of 0.94. The near-preservation of rankings reflects the fact that most occupations have relatively homogeneous within-occupation task difficulty, so the particular exponent used to aggregate across tasks has limited scope to reorder occupations.

A.2.2 Potential Supply Shift (PSS)

Figure A.2 plots the Potential Supply Shift—the descriptive measure of how much the eligible workforce expands due to reduced entry barriers—against the baseline expertise percentile, for each value of ρ. The inverted-U pattern documented in the main text (Figure 6a) is present for all aggregation rules: PSS peaks in the middle of the expertise distribution and is small at both extremes.

As ρ increases, PSS compresses slightly. This is because higher ρ raises the level of both R_o and $\widetilde{R}_o$, pushing expertise requirements into a range where the employment-weighted CDF $F(\cdot)$ is thinner, thereby reducing the mass of newly eligible workers. The qualitative shape, however, is stable.

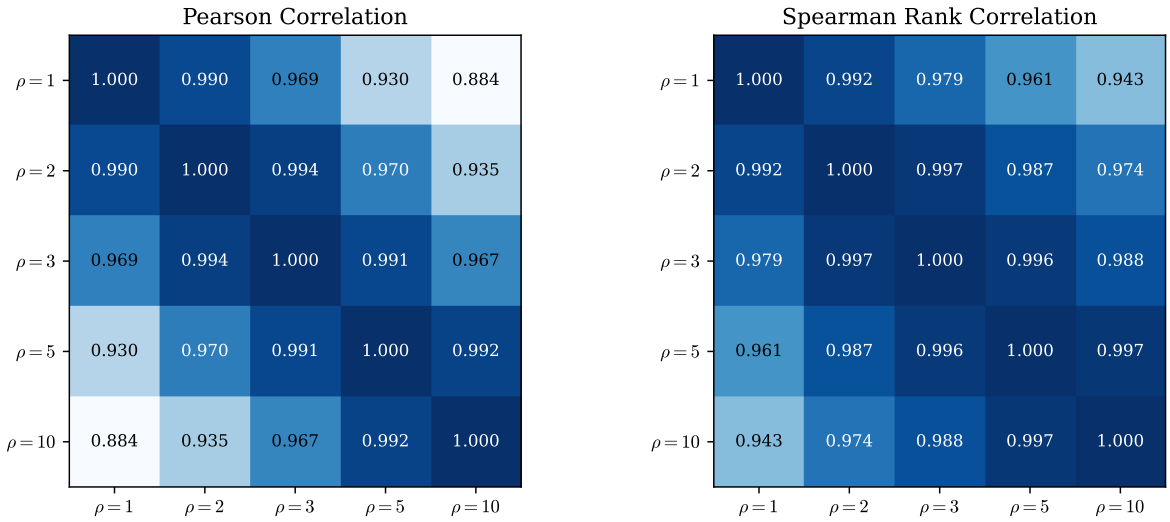

Figure A.1: Cross-ρ Correlations of Occupation-Level Expertise R_o

Notes: Each cell reports the pairwise correlation of occupation-level expertise requirements R_o computed under two different values of ρ in Equation (1). Left panel: Pearson correlation. Right panel: Spearman rank correlation. The sample consists of 923 occupations with non-missing task-level expertise and importance weights.

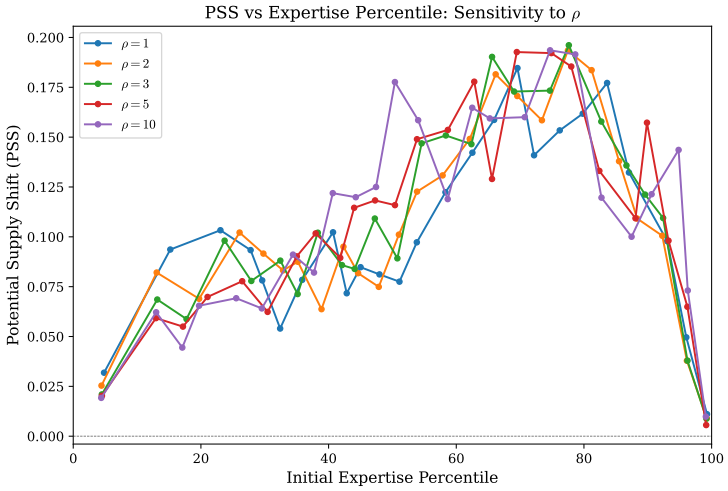

Figure A.2: PSS vs. Expertise Percentile Under Different ρ

Notes: Each line plots employment-weighted bin averages of the Potential Supply Shift (PSS) against occupation-level expertise percentiles, for a given value of ρ in the power-mean aggregator (1). PSS measures the change in the share of workers whose expertise exceeds the occupation's entry barrier. 20 equally spaced bins of the expertise percentile distribution are used.

A.2.3 General-Equilibrium Wage Effects

Figure A.3 presents the distributional pattern of supply-channel wage changes across the pre-AI wage distribution (analogous to Figure 8a in the main text). The U-shaped pattern is robust: low-wage occupations gain as barriers fall and workers gain access to higher-paying jobs; middle-wage occupations experience the largest wage declines as contestability increases; and high-wage occupations are relatively insulated. The curves are qualitatively identical across all five aggregation rules, with the magnitude of wage changes attenuating as ρ increases.

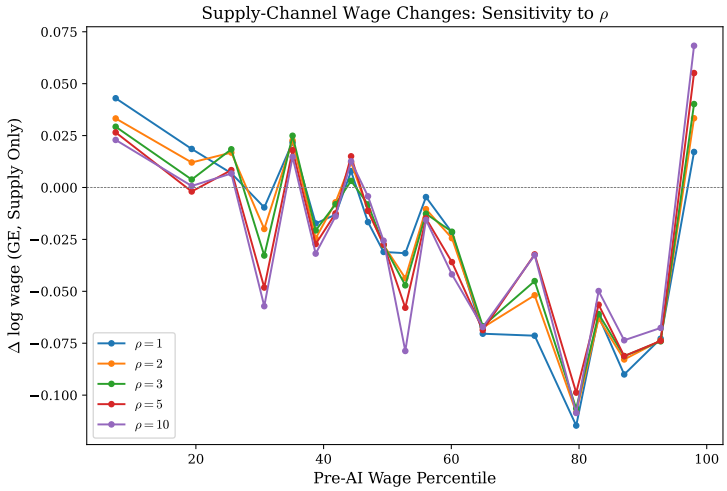

Figure A.3: Supply-Channel Wage Changes by Pre-AI Wage Percentile Under Different ρ

Notes: Each line plots employment-weighted bin averages of the percentage change in occupational wages (general equilibrium, supply channel only) against the occupation's percentile in the baseline wage distribution, for a given value of ρ. This figure is the multi-ρ analog of Figure 8a. 20 equally spaced bins of the pre-AI wage percentile distribution are used.

Table A.1 reports the main inequality statistics from the general-equilibrium counterfactuals under each aggregation rule. The supply channel (expertise reductions only) reduces the variance of log wages by 4.0–6.6%, with larger compressions at lower ρ. The productivity channel raises inequality by 26–31%, and the combined effect is a net increase of 15–21%. Across all values of ρ, the productivity channel dominates the supply channel.

Table A.1: GE Inequality Effects by Power-Mean Exponent ρ

	$\rho = 1$	$\rho = 2$	$\rho = 3$	$\rho = 5$	$\rho = 10$
Calibrated parameters					
τ	0.717	0.805	0.861	0.929	0.980
Baseline Var(log w)	0.201	0.192	0.188	0.184	0.183
% change in Var(log w)					
Supply only	-6.6	-6.1	-5.5	-4.7	-4.0
Productivity only	$+25.7$	$+27.9$	$+28.8$	$+29.8$	$+30.7$
Both channels	$+14.7$	$+16.9$	$+18.2$	$+19.4$	$+20.6$

Notes: Each column reports general-equilibrium results under a different power-mean aggregation exponent ρ for the task-to-occupation expertise mapping. τ is the logit taste parameter calibrated to match observed employment shares. "Supply only" isolates the wage effects of reduced entry barriers. "Productivity only" isolates the wage effects of higher occupational productivity. "Both channels" combines both mechanisms. All percentage changes are relative to the pre-AI baseline.

A.3 Mapping Potential Productivity Gains into Model Productivity Shocks

As described in Section 3, the model's demand-side GenAI shock is a labor-augmenting productivity increase, $A_o^{AI} = A_o \exp(\pi_o)$, where π_o is the log change in effective productivity in occupation o (Section 3.5). We interpret PPG_o as the fraction of total work time that can be saved within occupation o when GenAI is used for augmentation, holding the occupation's task bundle fixed.

Let T_o denote the baseline labor time required to produce one unit of occupational output. If GenAI reduces the time requirement by a fraction PPG_o, then

$$T_o^{AI} = (1 - PPG_o)\, T_o.$$

Because labor productivity is inversely related to time per unit, the implied multiplicative productivity change is

$$\frac{A_o^{AI}}{A_o} = \frac{1/T_o^{AI}}{1/T_o} = \frac{1}{1 - PPG_o}.$$

Equating this with $A_o^{AI} = A_o \exp(\pi_o)$ yields

$$\pi_o = \log\left(\frac{A_o^{AI}}{A_o}\right) = -\log(1 - PPG_o). \tag{29}$$

For example, if $PPG_o = 0.20$, then GenAI reduces time per unit by 20 percent and raises effective labor productivity by $1/(1 - 0.20) = 1.25$, so $\pi_o = \log(1.25)$.

A.4 Analytical Characterization of the Expertise Premium

This appendix characterizes how GenAI affects the *expertise premium*—the elasticity of wages with respect to occupational expertise requirements—under distributional assumptions that deliver closed-form results. The expertise premium, $\xi := d\log w / d\log r$, measures how steeply wages rise with the expertise threshold an occupation demands. It is related to, but distinct from, overall wage inequality: changes in ξ shift the slope of the wage–expertise relationship, whereas wage inequality also depends on the distribution of workers across expertise levels and general-equilibrium reallocation. The assumptions required for tractability (two occupations, or a Pareto expertise distribution with power-law demand) are more restrictive than the general model in Section 3.5, which places no parametric restrictions and which we solve numerically in Section 5. The analytical models are therefore complementary to the quantitative framework: they sacrifice generality for transparency and make explicit which primitives determine whether GenAI steepens or flattens the expertise premium. Full derivations and proofs are in Appendix A.6.

A.4.1 Two-Occupation Model

Consider two occupations, Low (L) and High (H), with expertise requirements $R_L = 0 < R_H$. Let $m_L := F(R_H)$ be the mass of workers qualified for L only and $m_H := 1 - F(R_H)$ the mass qualified for both. Define the expertise premium $x := w_H/w_L$. GenAI affects the economy through the productivity channel, $\Delta\log(B_H/B_L) \approx \frac{\sigma-1}{\sigma}(PPG_H - PPG_L)$, and the scarcity channel, $PSS_H = dm_H > 0$.

Proposition 1 (The productivity–scarcity race)**.** *Under simultaneous PPG and PSS shocks, the change in the expertise premium is:*

$$d\log x^* = \mathcal{D}\left[\underbrace{\frac{\sigma-1}{\sigma}\left(PPG_H - PPG_L\right)}_{\text{productivity: widens premium}} - \underbrace{\frac{1}{\sigma}\cdot\frac{\rho}{1+\rho}\left(\frac{1}{L_H}+\frac{1}{L_L}\right)\cdot PSS_H}_{\text{scarcity: compresses premium}}\right], \tag{30}$$

where $\rho := (w_H/w_L)^{1/\tau}$ and $\mathcal{D} := 1/(1+\mathcal{E}/\sigma) \in (0,1)$ is a GE dampening factor.

The premium widens when the PPG differential is large relative to the scarcity term, and compresses when PSS is large. The two-occupation model delivers *level* predictions:

PSS and PPG enter as occupation-specific magnitudes. It also yields a wage-level decomposition (Appendix A.6): the scarcity channel *raises* w_L (workers exit L, reducing its supply) and *lowers* w_H (new entrants erode the scarcity premium), so both effects compress the premium.

A.4.2 Continuous Assignment Model and the Wage Gradient

To characterize the cross-occupation wage structure, consider a continuum of occupations indexed by baseline expertise requirement r, with Pareto-distributed expertise ($F(e) = 1 - (\underline{r}/e)^{\theta}, \theta > 1$), power-law demand ($B(r) = b\, r^{\beta}$), and CES final demand with elasticity $\sigma > 1$. Under perfect sorting ($\tau = 0$), each worker chooses the highest-wage feasible occupation, and the model yields a tractable power-law characterization of wages, $w(r) \propto r^{\xi}$. The central object of interest is the expertise premium:

$$\xi := \frac{d \log w}{d \log r}, \tag{31}$$

the elasticity of wages with respect to occupational expertise. A higher ξ means that wages rise more steeply with expertise, so occupations that require more expertise command proportionally higher wages.[28] In the baseline economy,

$$\xi = \beta + \frac{\theta + 1}{\sigma}. \tag{32}$$

We next characterize how the expertise premium changes when technology shifts both labor demand and labor supply. The key distributional objects are two gradients that capture how PPG (the demand shift) and PSS (the supply shift) vary systematically across the expertise distribution:

$$\mathcal{G}_{PPG} := \frac{d\, PPG(r)}{d \log r}, \qquad \mathcal{G}_{PSS} := \frac{d \log\big(S^{AI}(r)/S(r)\big)}{d \log r}, \tag{33}$$

where $S(r) = 1 - F(r)$ is the baseline eligible pool (as in Equation (8)) and $S^{AI}(r) = 1 - F(\widetilde{R}(r))$ is the post-GenAI eligible pool. $\mathcal{G}_{PPG}$ measures the rate at which productivity

[28] Note that ξ is distinct from cross-occupation wage inequality, such as $\mathrm{Var}(\log w)$, which depends both on the wage gradient and on the distribution of workers across occupations.

gains rise with expertise; $\mathcal{G}_{PSS}$ measures the rate at which the *proportional* eligible-pool expansion grows with expertise.

Proposition 2 (Post-GenAI Expertise Premium)**.** *Under Pareto expertise, power-law demand, log-linear productivity gains, and perfect sorting ($\tau = 0$), the post-GenAI economy admits the following analytical characterization of the expertise premium:*

$$\xi^{AI} = \underbrace{\beta + \frac{\sigma - 1}{\sigma}\,\mathcal{G}_{PPG}}_{\textit{demand gradient}} + \underbrace{\frac{\theta + 1 - \mathcal{G}_{PSS}}{\sigma}}_{\textit{scarcity gradient}}, \tag{34}$$

and the change in the expertise premium is

$$\xi^{AI} - \xi = \frac{1}{\sigma}\left[(\sigma - 1)\,\mathcal{G}_{PPG} - \mathcal{G}_{PSS}\right]. \tag{35}$$

Equation (35) yields a sharp *race condition*: the expertise premium steepens—meaning that GenAI increases the return to occupational expertise—if and only if

$$(\sigma - 1)\,\mathcal{G}_{PPG} > \mathcal{G}_{PSS}. \tag{36}$$

Productivity dominates when $\mathcal{G}_{PPG}$ is large, reflecting strongly skill-biased productivity gains, and when σ is large, so that substitutability amplifies the demand response. Scarcity dominates when $\mathcal{G}_{PSS}$ is large, which occurs when barrier compression is strong and the expertise tail is thin, so that even modest reductions in expertise requirements generate large proportional expansions of the eligible pool at high expertise levels. The race condition depends only on observable cross-occupation gradients and the CES elasticity, making these objects sufficient statistics for characterizing the effect of GenAI on the expertise premium.

A.5 Empirical Assessment of the Analytical Race Condition

We complement the quantitative counterfactuals in Section 5.1.2 with a direct empirical assessment of the race condition derived in Appendix A.4. That characterization implies that the expertise premium steepens if and only if $(\sigma - 1)\,\mathcal{G}_{PPG} > \mathcal{G}_{PSS}$ (Equation (36)), where $\mathcal{G}_{PPG}$ and $\mathcal{G}_{PSS}$ denote, respectively, the gradients of PPG and of the log proportional expansion in the eligible pool with respect to log expertise (Equation (33)). Because these objects are defined over observables, the race condition can be evaluated directly in the data. In this sense, $\mathcal{G}_{PPG}$ and $\mathcal{G}_{PSS}$ are sufficient statistics for the change in the expertise premium.

We estimate $\mathcal{G}_{PPG}$ as the slope from an employment-weighted regression of occupation-level PPG_o on $\log R_o$, where R_o is the baseline expertise requirement (months of training without GenAI). We estimate $\mathcal{G}_{PSS}$ as the slope from an employment-weighted regression of $\log\left(S_o^{AI}/S_o\right)$ on $\log R_o$, where $S_o = 1 - F(R_o)$ and $S_o^{AI} = 1 - F(\widetilde{R}_o)$ are the baseline and post-GenAI eligible pools defined in Section 4.1.[29]

Table A.2 reports the estimates. In the baseline specification (employment-weighted, $\sigma = 5$), $\hat{\mathcal{G}}_{PPG} = 0.036$ and $\hat{\mathcal{G}}_{PSS} = 0.133$, both precisely estimated. The race condition yields $(\sigma - 1)\hat{\mathcal{G}}_{PPG} = 0.143 > 0.133 = \hat{\mathcal{G}}_{PSS}$: the productivity force slightly exceeds the scarcity force, implying that GenAI *steepens* the expertise premium. The implied change is $\Delta\xi = +0.002$, an increase of less than 1 percent relative to the baseline empirical expertise premium ($\hat{\xi} \approx 0.31$, estimated by regressing log wages on log expertise). To gauge the magnitude: across the $p90$–$p10$ range of baseline expertise, $\Delta\xi = +0.002$ widens the log wage gap by about 0.007, which is negligible relative to the baseline $p90$–$p10$ log wage gap of approximately 1.15. The two forces are therefore close in magnitude, with the productivity channel slightly dominating on average. The result is robust to using unweighted OLS ($\Delta\xi = +0.003$).

The closed-form race condition points toward slight steepening, consistent in direction with the quantitative counterfactuals in Section 5.1.2, which also show that wage inequality rises modestly on net, though with a larger magnitude. This gap between the expertise-premium prediction and the full-model inequality prediction reflects precisely

[29] Because GenAI raises the effective expertise requirement for a small number of occupations (about 7 percent of occupations, representing 6 percent of employment), $S_o^{AI} < S_o$ in those cases. These observations are included in all regressions; excluding them strengthens the result.

Table A.2: Empirical Assessment of the productivity–scarcity Race Condition

Specification	$\hat{\mathcal{G}}_{PPG}$	$\hat{\mathcal{G}}_{PSS}$	$(\sigma-1)\hat{\mathcal{G}}_{PPG}$	$\hat{\mathcal{G}}_{PSS}$	$\Delta\xi$
Employment-weighted	0.036	0.133	0.143	0.133	+0.002
	(0.002)	(0.006)			
Unweighted	0.039	0.143	0.156	0.143	+0.003
	(0.003)	(0.009)			

Notes: $\hat{\mathcal{G}}_{PPG}$ is the slope from regressing PPG_o on $\log R_o$; $\hat{\mathcal{G}}_{PSS}$ is the slope from regressing $\log(S_o^{AI}/S_o)$ on $\log R_o$. Standard errors in parentheses. $\Delta\xi = [(\sigma-1)\hat{\mathcal{G}}_{PPG} - \hat{\mathcal{G}}_{PSS}]/\sigma$. $\sigma = 5$. $N = 885$ occupations.

the distinction between the two objects: the expertise premium captures the slope of the wage–expertise relationship, whereas overall inequality also reflects worker reallocation and general-equilibrium adjustments across the full occupation distribution.

A.6 Analytical Appendix: The productivity–scarcity Race

This appendix provides the full derivations underlying the analytical results in Appendix A.4. Section A.6.1 develops the two-occupation model. Section A.6.2 develops the continuous assignment model with Pareto-distributed expertise. Under perfect sorting ($\tau = 0$), the model admits an exact, globally valid, closed-form wage gradient expressed directly in terms of the cross-occupation gradients of PPG and PSS. A robustness result then establishes that this gradient is unchanged when sorting is imperfect ($\tau > 0$).

A.6.1 Two-Occupation Model

Setup. There are two occupations, Low (L) and High (H), with expertise requirements $R_L = 0 < R_H$. Workers draw expertise e from a continuous distribution $F(\cdot)$ with density f. Define $m_L := F(R_H)$ (mass qualified for L only) and $m_H := 1 - F(R_H)$ (mass qualified for both), with $m_L + m_H = 1$. Demand shifters are $B_L, B_H > 0$, the CES elasticity is $\sigma > 1$, and logit choice dispersion is $\tau > 0$.

Workers with $e < R_H$ can only work in L. Workers with $e \geq R_H$ choose between L and H via multinomial logit. Defining $\rho := (w_H/w_L)^{1/\tau}$, labor supply is:

$$L_H = m_H \cdot \frac{\rho}{1+\rho}, \tag{37}$$

$$L_L = m_L + m_H \cdot \frac{1}{1+\rho}. \tag{38}$$

Labor demand is $w_o = B_o \, L_o^{-1/\sigma}$ for $o \in \{L, H\}$.

Equilibrium characterization. The equilibrium skill premium $x^* := w_H/w_L$ is the unique solution to:

$$x = \frac{B_H}{B_L} \cdot \left(\frac{m_H \, \rho}{m_L(1+\rho) + m_H} \right)^{-1/\sigma}, \qquad \rho = x^{1/\tau}. \tag{39}$$

Proof. Taking the ratio of the two demand equations and substituting (37)–(38) yields (39). Denoting the right-hand side $\Psi(x)$: as $x \to 0$, $\Psi(x) \to \infty$; as $x \to \infty$, $\Psi(x) \to (B_H/B_L)(m_H/m_L)^{-1/\sigma} < \infty$. Since Ψ is continuous and strictly decreasing while x is the identity, there is a unique crossing. □

Under perfect sorting ($\tau \to 0$), all high-type workers sort into H:

$$L_L = m_L, \quad L_H = m_H, \quad w_o = B_o\, m_o^{-1/\sigma}, \quad x^* = \frac{B_H}{B_L}\left(\frac{m_H}{m_L}\right)^{-1/\sigma}. \tag{40}$$

Comparative statics with PPG and PSS. GenAI affects the two-occupation economy through PPG and PSS. The productivity channel operates through the demand-shifter ratio:

$$\Delta \log \frac{B_H}{B_L} \;\approx\; \frac{\sigma - 1}{\sigma}\left(PPG_H - PPG_L\right). \tag{41}$$

The scarcity channel operates through the expansion of m_H:

$$PSS_H := F(R_H) - F(\widetilde{R}_H) = dm_H > 0. \tag{42}$$

Define the GE dampening factor $\mathcal{D} := 1/(1 + \mathcal{E}/\sigma) \in (0, 1)$, where $\mathcal{E} := d\log(L_H/L_L)/d\log x|_m > 0$.

Proof of Proposition 1. Totally differentiate $\log x = \log(B_H/B_L) - (1/\sigma)\log(L_H/L_L)$. At fixed x, increasing m_H by PSS_H changes the employment ratio by $\partial \log(L_H/L_L)/\partial m_H|_x = [\rho/(1+\rho)](1/L_H + 1/L_L)$. Combining with the endogenous response $d\log(L_H/L_L) = \mathcal{E}\, d\log x$ and solving yields Equation (30). □

Wage-level decomposition under perfect sorting. When $\tau \to 0$:

$$\Delta \log w_L = \underbrace{\frac{\sigma - 1}{\sigma} PPG_L}_{\text{productivity}} + \underbrace{\frac{1}{\sigma}\log\frac{m_L}{m_L - PSS_H}}_{\text{scarcity gain: workers exit } L}, \tag{43}$$

$$\Delta \log w_H = \underbrace{\frac{\sigma - 1}{\sigma} PPG_H}_{\text{productivity}} - \underbrace{\frac{1}{\sigma}\log\frac{m_H + PSS_H}{m_H}}_{\text{scarcity loss: workers enter } H}. \tag{44}$$

Proof. From (40), $w_o^{AI} = B_o^{AI}(m_o^{AI})^{-1/\sigma}$. Using $\Delta \log B_o \approx (\sigma - 1)/\sigma \cdot PPG_o$, $m_L^{AI} = m_L - PSS_H$, and $m_H^{AI} = m_H + PSS_H$. □

The scarcity channel *raises* w_L (workers exit L, reducing its supply) and *lowers* w_H

(new entrants increase supply and erode the scarcity premium). Both effects compress the premium. The productivity channel raises both wages, but by potentially different amounts.

A.6.2 Continuous Assignment Model with Pareto Expertise

Environment. Occupations form a continuum indexed by baseline expertise requirement $r \in [\underline{r}, \infty)$, $\underline{r} > 0$. Demand is $B(r) = b\, r^{\beta}$, $b > 0$, $\beta \geq 0$. Expertise is Pareto: $1 - F(e) = (\underline{r}/e)^{\theta}$, $f(e) = \theta \underline{r}^{\theta}\, e^{-(\theta+1)}$, $e \geq \underline{r}$, $\theta > 1$. CES final demand implies $w(r) = B(r)\, \ell(r)^{-1/\sigma}$, $\sigma > 1$.

Exact baseline equilibrium under perfect sorting ($\tau = 0$). Under perfect sorting, each worker with expertise e chooses the highest-wage feasible occupation. With baseline requirements $R(r) = r$ and wages increasing in r (verified below), the assignment is $r(e) = e$: each worker works in the occupation that exactly matches their expertise. Employment density equals expertise density:

$$\ell(r) = f(r) = \theta\, \underline{r}^{\theta}\, r^{-(\theta+1)}, \qquad r \geq \underline{r}, \tag{45}$$

which integrates to one. Substituting into the wage equation:

$$w(r) = b r^{\beta} \left(\theta \underline{r}^{\theta} r^{-(\theta+1)}\right)^{-1/\sigma} = b(\theta \underline{r}^{\theta})^{-1/\sigma}\, r^{\beta + (\theta+1)/\sigma}.$$

Since $\beta + (\theta+1)/\sigma > 0$, wages are increasing in r, confirming the sorting assumption. The baseline wage gradient is:

$$\xi := \frac{d \log w}{d \log r} = \beta + \frac{\theta + 1}{\sigma}. \tag{46}$$

This expression is exact and holds globally on $[\underline{r}, \infty)$, with no approximation involved.

GenAI shocks: two observable gradients. *PPG gradient.* GenAI raises productivity: $A(r)^{AI} = A(r) \exp(\pi(r))$, where $\pi(r) \approx PPG(r)$ for moderate gains. Under log-linearity,

$PPG(r) = PPG_0 + \mathcal{G}_{PPG} \cdot \log(r/\underline{r})$, where

$$\mathcal{G}_{PPG} := \frac{d\,PPG(r)}{d\log r} \tag{47}$$

is the rate at which productivity gains rise with expertise. Through $B \propto A^{1-1/\sigma}$, the demand exponent becomes $\beta + (\sigma-1)\mathcal{G}_{PPG}/\sigma$.

PSS gradient. GenAI reduces the entry requirement from $R(r) = r$ to $\widetilde{R}(r) \leq r$. Under anchored power-law compression $\widetilde{R}(r) = \underline{r}^{1-\delta} r^{\delta}$ (which satisfies $\widetilde{R}(\underline{r}) = \underline{r}$, $0 < \delta \leq 1$):

$$S(r) = \left(\frac{\underline{r}}{r}\right)^{\theta}, \qquad S^{AI}(r) = \left(\frac{\underline{r}}{r}\right)^{\delta\theta}, \qquad \frac{S^{AI}(r)}{S(r)} = \left(\frac{r}{\underline{r}}\right)^{(1-\delta)\theta}, \tag{48}$$

with constant log-derivative:

$$\mathcal{G}_{PSS} := \frac{d\log\big(S^{AI}(r)/S(r)\big)}{d\log r} = (1-\delta)\theta. \tag{49}$$

$\mathcal{G}_{PSS}$ is positive when $\delta < 1$ (GenAI compresses barriers), meaning the supply expansion is proportionally larger for higher-expertise occupations. It is determined by two primitives: barrier compression $(1-\delta)$ and tail thinness (θ).

Proof of Proposition 2 (exact post-GenAI wage gradient under $\tau = 0$). Under $\tau = 0$ and the post-AI barrier $\widetilde{R}(r) = \underline{r}^{1-\delta} r^{\delta}$, worker e accesses occupations up to $r^*(e) = \underline{r}^{(\delta-1)/\delta} e^{1/\delta}$. Since wages are increasing in r (verified below), the assignment is $r^*(e)$. By change of variables ($e = \widetilde{R}(r) = \underline{r}^{1-\delta} r^{\delta}$, $de = \underline{r}^{1-\delta}\delta r^{\delta-1} dr$):

$$\ell^{AI}(r) = f(\widetilde{R}(r)) \cdot |\widetilde{R}'(r)| = \frac{\theta \underline{r}^{\theta}}{(\underline{r}^{1-\delta} r^{\delta})^{\theta+1}} \cdot \underline{r}^{1-\delta}\delta r^{\delta-1} = \theta\delta\, \underline{r}^{\delta\theta}\, r^{-(\delta\theta+1)}.$$

This integrates to one: $\theta\delta\underline{r}^{\delta\theta} \int_{\underline{r}}^{\infty} r^{-(\delta\theta+1)} dr = 1$. The post-AI demand shifter is $B^{AI}(r) = be^{(\sigma-1)PPG_0/\sigma} r^{\beta+(\sigma-1)\mathcal{G}_{PPG}/\sigma}$. From inverse demand:

$$w^{AI}(r) = B^{AI}(r)\, \ell^{AI}(r)^{-1/\sigma} \propto r^{\beta+(\sigma-1)\mathcal{G}_{PPG}/\sigma+(\delta\theta+1)/\sigma}.$$

Using $\mathcal{G}_{PSS} = (1-\delta)\theta$: $(\delta\theta+1)/\sigma = (\theta+1-\mathcal{G}_{PSS})/\sigma$. The gradient is therefore $\xi^{AI} = \beta + (\sigma-1)\mathcal{G}_{PPG}/\sigma + (\theta+1-\mathcal{G}_{PSS})/\sigma$, and the change is $\xi^{AI} - \xi = [(\sigma-1)\mathcal{G}_{PPG} - \mathcal{G}_{PSS}]/\sigma$. No approximation is involved at any step. □

Robustness: the wage gradient is independent of τ. The perfect-sorting assumption delivers exact closed forms because employment density equals the expertise density evaluated at the assignment function. A natural question is whether the wage gradient changes when sorting is imperfect ($\tau > 0$), so that workers sometimes choose feasible occupations other than the highest-wage option.

With $\tau > 0$ and logit occupational choice (utility $u = \log w(r) + \varepsilon(r)$, $\varepsilon \sim$ Type-I EV with scale τ), conjecturing $w(r) = C(\tau) r^{\xi}$ and defining $\zeta := \xi/\tau + 1$, the inclusive value is $\Phi(e) = C^{1/\tau}(e^{\zeta} - \underline{r}^{\zeta})/\zeta$. For interior r ($r/\underline{r}$ bounded away from one), $\Phi(e) \approx C^{1/\tau} e^{\zeta}/\zeta$, and aggregating over eligible workers:

$$\ell(r) \propto r^{\xi/\tau} \int_r^{\infty} e^{-(\theta+1+\zeta)}\, de \propto r^{\xi/\tau} \cdot r^{-(\theta+\zeta)} = r^{-(\theta+1)},$$

since $\xi/\tau - \theta - \zeta = -(\theta+1)$. From demand: $\xi = \beta + (\theta+1)/\sigma$, confirming the gradient is independent of τ. The constant $C(\tau)$ depends on τ, but ξ does not. The boundary correction is of order $(\underline{r}/r)^{\zeta}$ with $\zeta > 1$, vanishing rapidly for all but the lowest-expertise occupations. The post-GenAI extension follows identically, yielding ξ^{AI} as in Proposition 2.

This τ-independence reflects the self-similarity of the Pareto distribution: with power-law wages and Pareto expertise, the logit choice probabilities maintain the same power-law structure regardless of τ. Changing τ rescales wage and employment levels but preserves the gradient.

Reading the formula. The post-GenAI wage gradient decomposes into a demand gradient and a scarcity gradient. The demand gradient $\beta + (\sigma-1)\mathcal{G}_{PPG}/\sigma$ captures how fast labor demand rises with expertise; GenAI steepens this when $\mathcal{G}_{PPG} > 0$. The scarcity gradient $(\theta+1-\mathcal{G}_{PSS})/\sigma$ captures how fast wages rise due to progressively scarcer labor; GenAI flattens this by $\mathcal{G}_{PSS}/\sigma$ because proportionally larger eligible-pool expansions at higher expertise levels erode the scarcity advantage.

Decomposition table.

	Demand gradient	Scarcity gradient	Wage gradient ξ
Baseline	β	$\frac{\theta+1}{\sigma}$	$\beta + \frac{\theta+1}{\sigma}$
PPG only	$\beta + \frac{\sigma-1}{\sigma}\mathcal{G}_{PPG}$	$\frac{\theta+1}{\sigma}$	$\xi + \frac{\sigma-1}{\sigma}\mathcal{G}_{PPG}$
PSS only	β	$\frac{\theta+1-\mathcal{G}_{PSS}}{\sigma}$	$\xi - \frac{\mathcal{G}_{PSS}}{\sigma}$
PPG + PSS	$\beta + \frac{\sigma-1}{\sigma}\mathcal{G}_{PPG}$	$\frac{\theta+1-\mathcal{G}_{PSS}}{\sigma}$	$\xi + \frac{(\sigma-1)\mathcal{G}_{PPG} - \mathcal{G}_{PSS}}{\sigma}$

PPG affects only the demand gradient; PSS affects only the scarcity gradient. The two forces enter additively with opposite signs.

The race condition. The wage gradient steepens if and only if $(\sigma-1)\,\mathcal{G}_{PPG} > \mathcal{G}_{PSS}$. Productivity dominates when productivity gains are strongly skill-biased ($\mathcal{G}_{PPG}$ large), substitutability is high (σ large), or GenAI does not differentially erode barriers at higher expertise levels ($\mathcal{G}_{PSS}$ small). Scarcity dominates when barrier compression is strong and the expertise tail is thin ($\mathcal{G}_{PSS} = (1-\delta)\theta$ large), substitutability is low (σ small), or productivity gains are not concentrated in high-expertise work.

PSS shape in the continuum. Under Pareto expertise and anchored compression:

$$PSS(r) = \left(\frac{r}{\underline{r}}\right)^{\delta\theta} - \left(\frac{r}{\underline{r}}\right)^{\theta}. \tag{50}$$

The *level* of PSS is hump-shaped in r: zero at $r = \underline{r}$, rising through the middle of the distribution, and declining in the upper tail as both S^{AI} and S vanish. This matches the empirical pattern in Figure 6a. The *proportional expansion* $S^{AI}(r)/S(r)$ is, by contrast, monotonically increasing. It is the proportional expansion whose gradient, $\mathcal{G}_{PSS}$, governs the wage gradient.

Gradient vs. cross-occupation inequality. The wage gradient is a *local* object: the log-elasticity of w with respect to r, constant across all $r \geq \underline{r}$. Under perfect sorting, it is exact

and globally valid. It flattens unambiguously when $\mathcal{G}_{PSS} > 0$.

Cross-occupation wage *inequality* ($\text{Var}(\log w)$) depends on both the gradient and the employment-weighted distribution of workers across occupations. Under barrier compression with sorting in the unbounded continuum, the latter changes: workers reach further up the ladder, stretching the occupation distribution. This stretching can offset the flatter gradient in an unbounded setting. In the quantitative model with a *fixed finite set of occupations*, this stretching does not arise: occupations are fixed, and barrier compression redistributes workers among existing occupations. The finite model therefore delivers the expected result that barrier compression is equalizing, consistent with the gradient. The continuum gradient captures the *force* that drives equalization in the finite model; the finite occupation structure is needed for inequality measurement.

Connection between the two models.

Force	**Two-occupation**	**Continuum**
Productivity	$\frac{\sigma-1}{\sigma}(PPG_H - PPG_L)$	$\frac{\sigma-1}{\sigma}\mathcal{G}_{PPG}$
Scarcity	$\frac{\rho}{(1+\rho)\sigma}\left(\frac{1}{L_H}+\frac{1}{L_L}\right)PSS_H$	$\frac{\mathcal{G}_{PSS}}{\sigma}$
Race condition	PPG gap vs. PSS term	$(\sigma-1)\,\mathcal{G}_{PPG} \gtrless \mathcal{G}_{PSS}$
GE dampening	$\mathcal{D} = \frac{1}{1+\mathcal{E}/\sigma}$	Built into gradient

The two formulations capture the same economic forces. The two-occupation model delivers *level* predictions with PSS and PPG entering as occupation-specific magnitudes. The continuum model characterizes the *gradient* of wages across the expertise distribution, with $\mathcal{G}_{PPG}$ and $\mathcal{G}_{PSS}$ measuring how PPG and the proportional PSS expansion vary systematically with expertise.

Mapping to empirical objects.

Parameter	Empirical counterpart
$\mathcal{G}_{PPG}$	Slope of PPG_o against $\log R_o$ (or expertise percentile) across occupations
$\mathcal{G}_{PSS}$	Slope of $\log(S_o^{AI}/S_o)$ against $\log R_o$ across occupations
β	Cross-occupation demand gradient (from baseline wages and employment)
θ	Pareto shape of expertise distribution
σ	CES elasticity of substitution across occupations
τ	Logit dispersion (does not affect the wage gradient)

A.7 Automation as a Demand-Side Channel: An Extension with Automation Services

This appendix extends the benchmark model by introducing a classical automation/displacement margin on the demand side. The extension nests the benchmark model when automation services are prohibitively expensive or when the technology offers no scope for substitution away from labor. The goal is conceptual: to show how a task-based measure of automation *feasibility* enters the occupational labor-demand schedule, and to clarify which additional objects must be specified to quantify its equilibrium incidence. A key distinction is that our task-based exposure measures discipline *technological scope* for substitution (captured by I_{ot}), not the equilibrium level of adoption, which depends on relative prices and substitution elasticities.

A.7.1 Task Structure and Automation Services

Each occupation o consists of tasks $t \in T_o = \{1, \ldots, J_o\}$ with O*NET weights $w_{ot} > 0$ and normalized weights

$$\omega_{ot} := \frac{w_{ot}}{\sum_{s \in T_o} w_{os}}, \qquad \sum_{t \in T_o} \omega_{ot} = 1.$$

Let $\ell_{ot} \geq 0$ denote labor time allocated to task (o, t) and let $k_{ot} \geq 0$ denote *automation services* (e.g., AI services or AI-enabled capital services) used in that task. We interpret k_{ot} as an input measured in efficiency units and priced at a rental rate r, which is an *effective* unit cost that subsumes compute/model costs as well as complementary integration, supervision, and organizational adjustment. Labor is augmented by GenAI through π_o as in (11).

We model each task as a CES combination of effective labor and automation services:

$$y_{ot} = \left((1 - I_{ot})^{\frac{1}{\eta}} \left(e^{\pi_o} \ell_{ot} \right)^{\frac{\eta-1}{\eta}} + I_{ot}^{\frac{1}{\eta}} k_{ot}^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \qquad I_{ot} \in [0, 1],\ \eta > 1. \tag{51}$$

The parameter I_{ot} captures the scope for automation in task (o, t). When $I_{ot} = 0$, the task is strictly labor-only; when $I_{ot} > 0$, the producer can substitute between labor and automation services depending on relative prices.

Task outputs are aggregated into occupational output via a CES aggregator:

$$Y_o = A_o \left(\sum_{t \in T_o} \omega_{ot}^{\frac{1}{\eta}} y_{ot}^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}. \tag{52}$$

Define total labor and automation services in occupation o:

$$L_o := \sum_{t \in T_o} \ell_{ot}, \qquad K_o := \sum_{t \in T_o} k_{ot},$$

and define the task-weighted automation capability (occupation exposure) as

$$\mathcal{I}_o := \sum_{t \in T_o} \omega_{ot} I_{ot} \in [0, 1]. \tag{53}$$

A.7.2 Occupation-Level Reduced Form and Unit Cost

Under cost-minimizing allocation across tasks for given totals (L_o, K_o), the nested CES structure (51)–(52) reduces to an occupation-level CES between effective labor and automation services:

$$Y_o = A_o \Big((1 - \mathcal{I}_o)^{\frac{1}{\eta}} \left(e^{\pi_o} L_o\right)^{\frac{\eta-1}{\eta}} + (\mathcal{I}_o)^{\frac{1}{\eta}} \left(K_o\right)^{\frac{\eta-1}{\eta}} \Big)^{\frac{\eta}{\eta-1}}. \tag{54}$$

Producers rent automation services at a fixed rental rate $r > 0$ (perfectly elastic supply). Given (w_o, r), cost minimization implies the standard CES input ratio in efficiency units:

$$\frac{K_o}{e^{\pi_o} L_o} = \frac{\mathcal{I}_o}{1 - \mathcal{I}_o} \left(\frac{w_o}{r\, e^{\pi_o}} \right)^{\eta}. \tag{55}$$

The corresponding unit cost (and competitive output price) is

$$p_o = c_o(w_o, r; \pi_o, \mathcal{I}_o) = \frac{1}{A_o} \left[(1 - \mathcal{I}_o) \left(\frac{w_o}{e^{\pi_o}} \right)^{1-\eta} + \mathcal{I}_o\, r^{1-\eta} \right]^{\frac{1}{1-\eta}}. \tag{56}$$

A.7.3 Labor Demand with Automation Services

Final demand remains CES across occupations as in (3). Under perfect competition, $p_o = c_o(\cdot)$ and CES demand implies $Y_o = \vartheta_o p_o^{-\sigma} Y$. Combining demand with (56) and the implied factor shares yields an occupational labor-demand schedule:

$$L_o^D(w_o; \widetilde{\Theta}_o) = D_o \, e^{(\sigma-1)\pi_o} \, w_o^{-\sigma} \, (1 - \mathcal{I}_o) \left[(1 - \mathcal{I}_o) + \mathcal{I}_o \left(\frac{r \, e^{\pi_o}}{w_o} \right)^{1-\eta} \right]^{\frac{\sigma-\eta}{\eta-1}}, \tag{57}$$

where $D_o := Y \, \vartheta_o \, A_o^{\sigma-1}$ is an occupation-specific shifter and $\widetilde{\Theta}_o = (r, \pi_o, \mathcal{I}_o, \eta, \sigma)$.

Equation (57) makes clear why the automation channel is difficult to quantify with occupation-level exposure data alone. First, the effect of $\mathcal{I}_o$ depends on the relative price re^{π_o}/w_o, but r is an *effective* deployment cost rather than an observed market price. Second, the magnitude of substitution depends on the within-occupation elasticity η. Third, even holding (r, η) fixed, the sign and size of the general-equilibrium effect depend on σ, which governs the strength of output expansion when unit costs fall. As a result, the same exposure profile $\{\mathcal{I}_o\}$ can imply sharply different wage and inequality effects across plausible calibrations.

A.7.4 Interpretation and Relation to the Benchmark Model

The benchmark model corresponds to a limit case of this extension. If $\mathcal{I}_o = 0$ for all occupations (no scope for automation services), or if r is sufficiently large that the cost-minimizing mix uses essentially no automation services, then (57) collapses to the familiar CES demand schedule $L_o^D = D_o e^{(\sigma-1)\pi_o} w_o^{-\sigma}$, and the only channels shaping wages are productivity and scarcity.

In the main text, we therefore focus on the productivity–scarcity race using the benchmark model, and we treat automation/displacement as an additional demand-side channel whose quantitative incidence is explored only through sensitivity analysis (Appendix A.8).

A.8 Sensitivity to Calibration Choices

This appendix documents that the distributional implications of the demand-side automation channel in Appendix A.7 are sensitive to parameters that are not well measured in available data. This motivates our decision to treat automation as an extension rather than a centerpiece of the quantitative analysis. That said, the sign of the automation channel may be more constrained than the parameter sensitivity below suggests: Autor and Kausik (2026) show that when the prevailing labor share exceeds its wage-maximizing level—a condition they estimate holds across twelve industrialized economies—automation that further reduces labor share raises average wages, though it may still increase inequality across sectors.

A.8.1 Which Additional Parameters Does Automation Require?

Relative to the benchmark productivity–scarcity model, the automation extension requires at least three additional objects:

1. **The effective price of automation services, r.** This is the unit rental rate of AI/automation services in efficiency units comparable to labor costs. Importantly, r should be interpreted as an *effective* cost of deploying automation services in production (including integration, complementary engineering, monitoring/supervision, compliance, and liability), not merely the posted price of a model API. This effective cost is not directly observed at the occupation level and may vary across settings.

2. **The within-occupation elasticity of substitution, η.** This governs how sharply firms substitute from labor to automation services when relative prices change. The K/L ratio scales as $(w_o/(re^{\pi_o}))^{\eta}$ in (55), so even moderate uncertainty about η can imply large differences in implied substitution.

3. **The mapping from task exposure to technology parameters.** Empirically, we observe task-level exposure classifications and construct occupation-level exposure shares. Translating these into the structural objects (I_{ot}) or $\mathcal{I}_o$ requires a modeling choice (binary vs. continuous, thresholds, and how exposure relates to substitution possibilities), and exposure alone does not identify equilibrium adoption without (r, η).

The final-demand elasticity σ also matters for automation incidence because it governs the scale (output expansion) response to lower unit costs. While σ is more commonly parameterized in applied work, plausible values still vary across contexts.

A.8.2 Analytical Sensitivity: Scale Effects vs. Substitution Effects

A key reason automation is hard to sign is that an increase in automation capability $\mathcal{I}_o$ generates two opposing forces in (57):

1. *Substitution effect:* holding output fixed, more scope for automation shifts the cost-minimizing input mix away from labor.

2. *Scale effect:* if automation reduces unit cost, output demand for occupation o expands under CES final demand, raising the demand for all inputs.

Which force dominates depends on (σ, η) and on relative prices through re^{π_o}/w_o.

To illustrate, define

$$u_o := \left(\frac{r\,e^{\pi_o}}{w_o}\right)^{1-\eta} = \left(\frac{w_o}{r\,e^{\pi_o}}\right)^{\eta-1}, \qquad B_o := (1-\mathcal{I}_o) + \mathcal{I}_o u_o.$$

Holding wages fixed, the impact of automation capability on labor demand implied by (57) is:

$$\left.\frac{\partial \log L_o^D}{\partial \mathcal{I}_o}\right|_w = -\frac{1}{1-\mathcal{I}_o} + \frac{\sigma-\eta}{\eta-1}\cdot\frac{u_o-1}{B_o}. \tag{58}$$

The first term in (58) is a mechanical substitution force: increasing $\mathcal{I}_o$ lowers labor's weight in the CES aggregator. The second term is a scale effect that operates through unit costs and final-demand expansion; its magnitude and sign are governed by $\frac{\sigma-\eta}{\eta-1}$ and by the relative-price index u_o. When $\sigma < \eta$, substitution tends to dominate and labor demand falls as automation feasibility rises. When $\sigma > \eta$, scale effects can offset or reverse substitution, so labor demand (and wages) may rise even as automation becomes more feasible.

A.8.3 Distributional Sensitivity Across Occupations

Even with constant r, automation incentives vary systematically with baseline wages because the relative price re^{π_o}/w_o varies across occupations. Holding (r, η) fixed, higher-wage occupations face stronger incentives to substitute toward automation services (lower re^{π_o}/w_o), which can compress upper-tail wages if substitution dominates. But when σ is large, cost reductions can generate substantial output expansion, potentially raising labor demand and wages in the same high-wage occupations. Because w_o is endogenous in general equilibrium, these relative-price and scale mechanisms interact, making the distributional incidence of automation especially sensitive to (r, η, σ).

A.8.4 Empirical Sensitivity Exercises

In the quantitative appendix to the paper, we implement a parameter-sweep exercise based on the automation extension (Appendix A.7). We vary (r, η, σ) over a grid of plausible values and recompute the automation-only counterfactual (and the joint counterfactual with productivity and scarcity).

Table A.3: Automation Sensitivity Grid: Parameter Choices

Parameter	Low	Baseline	High
Automation-service rental rate r (relative to median wage)	0.25	1.00	2.00
Within-occupation substitution elasticity η	2	10	20
Final-demand elasticity across occupations σ	2	5	8

Notes: Grid used in the sensitivity exercises for the automation-services extension. The purpose is to illustrate that automation incidence can vary substantially with (r, η, σ).

Table A.4: Wage Inequality Under Automation-Only Counterfactual: Sensitivity

Calibration	$\mathrm{Var}(\log w)$	$p10$	$p50$	$p90$	$p50 - p10$	$p90 - p50$	$p90 - p10$
Low r, High η	0.1735	9.9791	10.1250	10.8491	0.1459	0.7241	0.8701
Baseline (r, η, σ)	0.1255	10.4775	10.8432	11.3561	0.3657	0.5129	0.8786
High r, Low η	0.2141	10.3991	10.8096	11.4928	0.4105	0.6832	1.0937

Notes: Table reports employment-weighted inequality moments for log annual wages under the automation-only counterfactual, for three illustrative calibrations. Percentiles are of the employment-weighted wage distribution.

Across plausible calibrations, the automation channel can generate small or large ef-

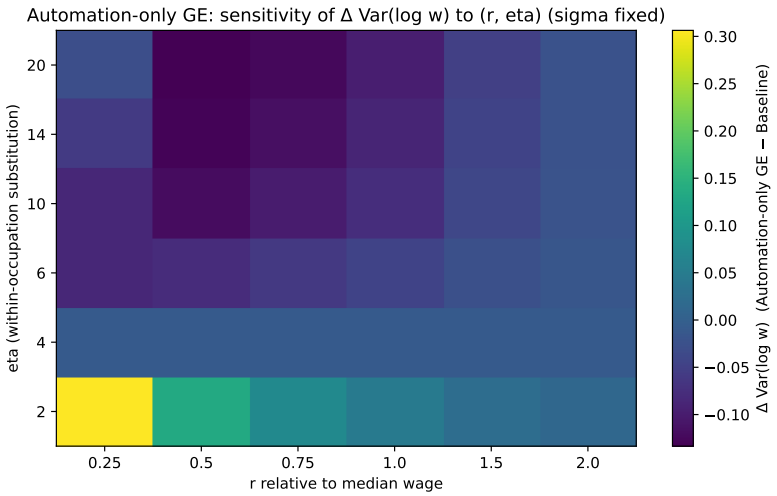

Figure A.4: Sensitivity of Aggregate Inequality Effects to (r, η)

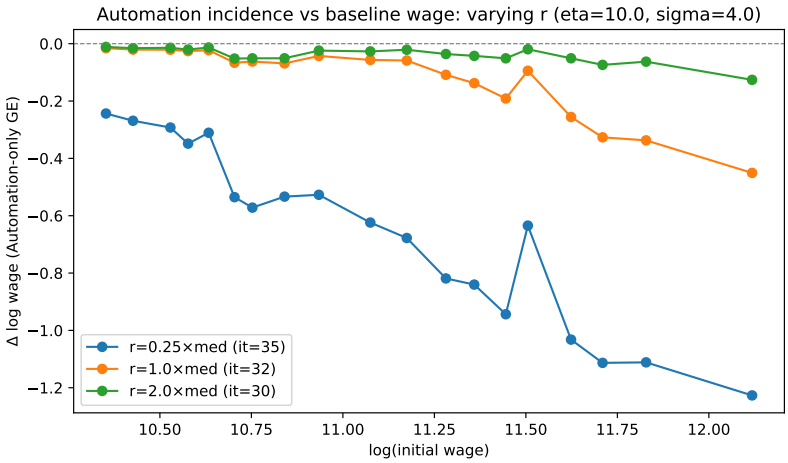

Figure A.5: Sensitivity of the Cross-Occupation Incidence of Automation

fects, and it can either compress or widen the wage distribution. This sensitivity reflects genuine economic ambiguity (scale vs. substitution) amplified by limited empirical information on the effective cost of automation services and substitution elasticities. For this reason, we treat automation as an extension and emphasize the productivity and scarcity channels—where our task-based measures provide more direct empirical discipline—in the benchmark quantitative analysis.

A.9 LLM Prompt—Automation Exposure Measure

"""
You are a "Dual-Expert" evaluator with two distinct, high-level competencies:

1. **Domain Expert:** You deeply understand the occupation listed below, including tacit knowledge, informal practices, interpersonal nuance, real-world context, and the hidden complexity that non-experts overlook.

2. **Gen AI Expert:** You understand the real capabilities and limitations of modern LLMs, agentic workflows, RAG systems, tool use (Python/APIs), and multimodal models. You can distinguish between "demo-level" and "production-reliable" AI behavior.

Your job is to integrate BOTH perspectives.

Evaluate the automation exposure of a specific occupational task: **Current December 2025 LLM abilities

—

AUTOMATION RUBRIC (T0–T4 — USE EXACT DEFINITIONS)

T0 — No Automation System cannot perform any meaningful component of the task. Typically highly physical, deeply emotional, or restricted by legal requirements.

T1 — Low Automation System can perform 0–50% of the task. Core work relies on physical action, real-world perception, in-person nuance, or tacit knowledge AI cannot substitute.

T2 — Moderate Automation (Hybrid) System can perform 50–80% of the task at high quality. Human involvement remains essential for: - physical actions, - in-person communication, - or tacit, context-dependent judgment.

T3 — High Automation (Human-in-the-Loop) System can perform 80–100% at high quality, BUT human oversight is required because of: - liability/safety, - stakeholder trust expectations, - rare catastrophic failure modes.

T4 — Full Automation (Autonomous) System performs 100% of the task with high quality. Human oversight is *not routinely needed* and humans are *not liable* for errors. The entire workflow is digitally executable end-to-end.

—

CAPABILITY ASSUMPTIONS
LLM(Current, Year 2025)**
CAN:
- Multi-step planning and reasoning
- Advanced multimodal analysis (vision/audio/docs/charts)
- Retrieval-augmented generation with large knowledge bases
- API/Python/Excel/SQL tool use
- Code generation, debugging, and software automation
- Structured workflow execution and agentic autonomy
- Up-to-date factual retrieval via internet tools
- Data analysis, visualization, and cleaning
- Idea generation, brainstorming, and research assistance
- High-quality writing, summarization, and translation
- Professional communication and long-context coherence
- Document extraction, comparison, and synthesis
- Business analytics, forecasting, and competitive intelligence
- UX/UI design, creative writing, images, audio, and video generation
- App/website creation and end-to-end automation
- Long-term memory and personalization (where enabled)
- Self-correction, critique, and verification loops
- Persistent agentic behaviors across sessions
- Scientific reasoning: literature synthesis, experimental design

CANNOT: - Perform any physical actions
- Conduct tactile or in-person inspection
- Guarantee correctness in high-stakes situations
- Substitute for human judgment where regulation requires a human
- Provide genuine emotional presence or tacit human understanding

OUTPUT FORMAT (MANDATORY JSON)
Return exactly:

{ ″T″: ″T0 — T1 — T2 — T3 — T4″ }
″″″

A.10 Comparing Automation Exposure Measures

This appendix compares our updated occupation-level automation exposure measure to the original measure developed by Eloundou et al. (2024). Figure A.6 plots, for each occupation, the share of tasks classified as automatable under our GPT-5.2-based measure against the corresponding share implied by the original GPT-4-based measure. Each point represents a single O*NET occupation. The figure shows a strong positive relationship between the two measures.

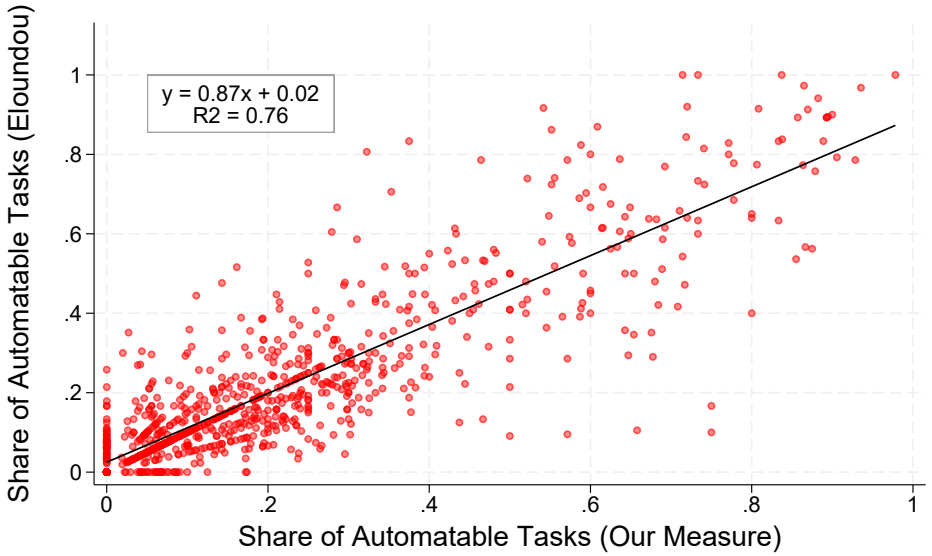

Figure A.6: Comparing Our Automation Exposure Measure to Eloundou et al. (2024)

Notes: Each point represents an O*NET occupation. The vertical axis reports the share of tasks classified as automatable using the automation measure of Eloundou et al. (2024), while the horizontal axis reports the corresponding share under our updated GPT-5.2-based measure.

A.11 PSS Dispersion Among Occupations with Similar Exposure Shares

This appendix complements the main-text discussion in Section 4.4. To make the dispersion of PSS at fixed exposure concrete, Figure A.7 displays PSS for all occupations with exposure between 29 and 31 percent (approximately the 75th percentile of the exposure distribution). Despite their nearly identical aggregate exposure shares, these occupations exhibit large differences in PSS, reinforcing the conclusion that exposure alone is an incomplete statistic for characterizing GenAI-induced changes in occupational entry barriers and effective labor supply.

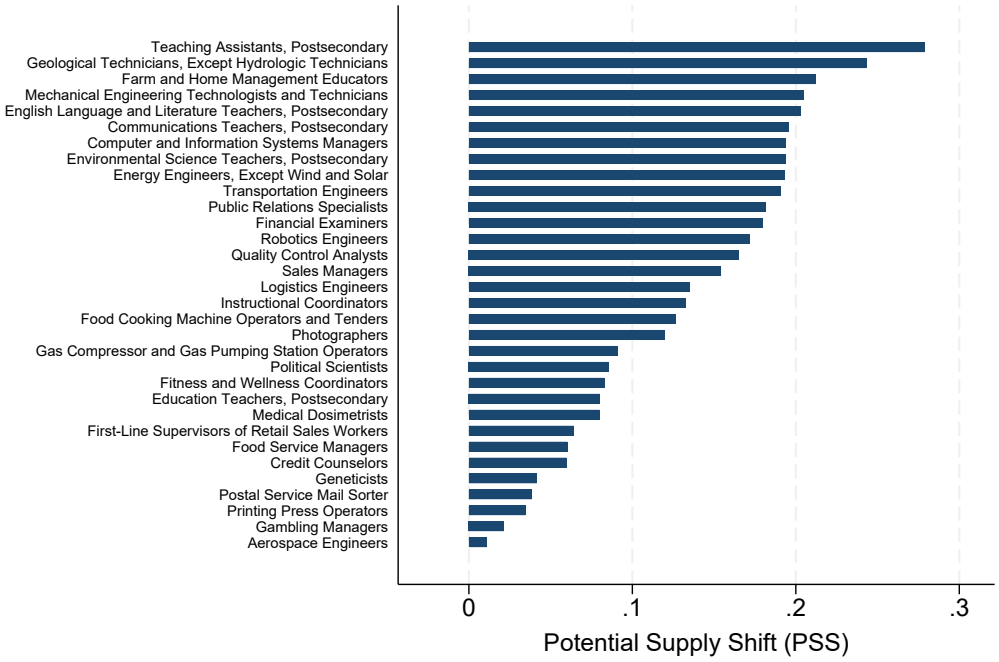

Figure A.7: PSS of Occupations with ∼30% Automatable Tasks

Notes: The figure displays PSS for all occupations with 29–31 percent exposed tasks (approximately the 75th percentile of the exposure distribution), illustrating that occupations with nearly identical exposure shares can exhibit large differences in PSS.

A.12 LLM Prompt—Task-Level Expertise Measure

"""
You are rating the expertise required to perform individual job tasks from O*NET.
By "expertise" we mean specialized knowledge, training, or skill that: - acts as a barrier to entry (not everyone can do it), and - typically commands a wage premium.
Generic physical or social tasks that almost any adult could do with minimal training are "no or low expertise".

Give each task a 1–5 expertise score:
1 = No or minimal expertise - Generic or very basic tasks - Can be learned quickly with little training.
2 = Low expertise - Some short training required - Limited occupation-specific knowledge.
3 = Moderate expertise - Solid occupation-specific knowledge - Often some credential, apprenticeship, or substantial on-the-job learning.
4 = High expertise - Advanced specialized knowledge - Significant training, degree, or certification.
5 = Very high expertise - Deep specialized expertise - Often advanced professional or graduate-level training.

Output STRICT JSON ONLY:
"expertise_score": 1-5 integer,
"expertise_label": "no_or_minimal" — "low" — "moderate" — "high" — "very_high",
"confidence": 0.0-1.0

Do NOT include explanations. Do NOT include text before or after the JSON. Output ONLY valid JSON. No markdown. No comments. No words before or after.

"""

A.13 LLM Prompt—Required Training Time

""" You are rating the expertise required to perform individual job tasks from O*NET.

For each task, estimate the minimum amount of specific education and training time (in months) required for a typical adult to learn to perform this task at a professional level.

Guidelines: - Consider specialized knowledge, training, or skill needed. - Focus on training/education time required to reach professional competence, not years of generic work experience. - Respond with a single number in months (e.g., 0.5, 6, 18, 36, 60).

Output STRICT JSON ONLY:

```
{
        "training_months": number,
        "confidence": 0.0-1.0
}
```

Do NOT include explanations. Do NOT include text before or after the JSON. Output ONLY valid JSON. No markdown. No comments. No words before or after. """

A.14 Expertise Distributions Under GenAI Scenarios

This appendix plots the distributions of four occupation-level expertise measures, each computed as a task-weighted average (with weights of 1 for core tasks and 0.5 for supplemental tasks) under different GenAI scenarios. The *baseline* measure X_o^{noAI} averages task-level expertise ratings without any GenAI effects. The *extensive-margin* measure $X_o^{extensive}$ averages expertise over only the tasks that are *not* classified as automatable (i.e., tasks with $a_{ot} = 0$), thereby capturing the effect of removing automated tasks from the occupation's task bundle. The *intensive-margin* measure $X_o^{intensive}$ averages the GenAI-assisted expertise ratings x_{ot}^{withAI} over the full set of tasks, capturing the reduction in required expertise when workers perform all tasks with GenAI assistance. Finally, the *combined* measure $X_o^{combined}$ applies both margins simultaneously: it averages GenAI-assisted expertise over only the non-automated tasks, as defined in equation (19).

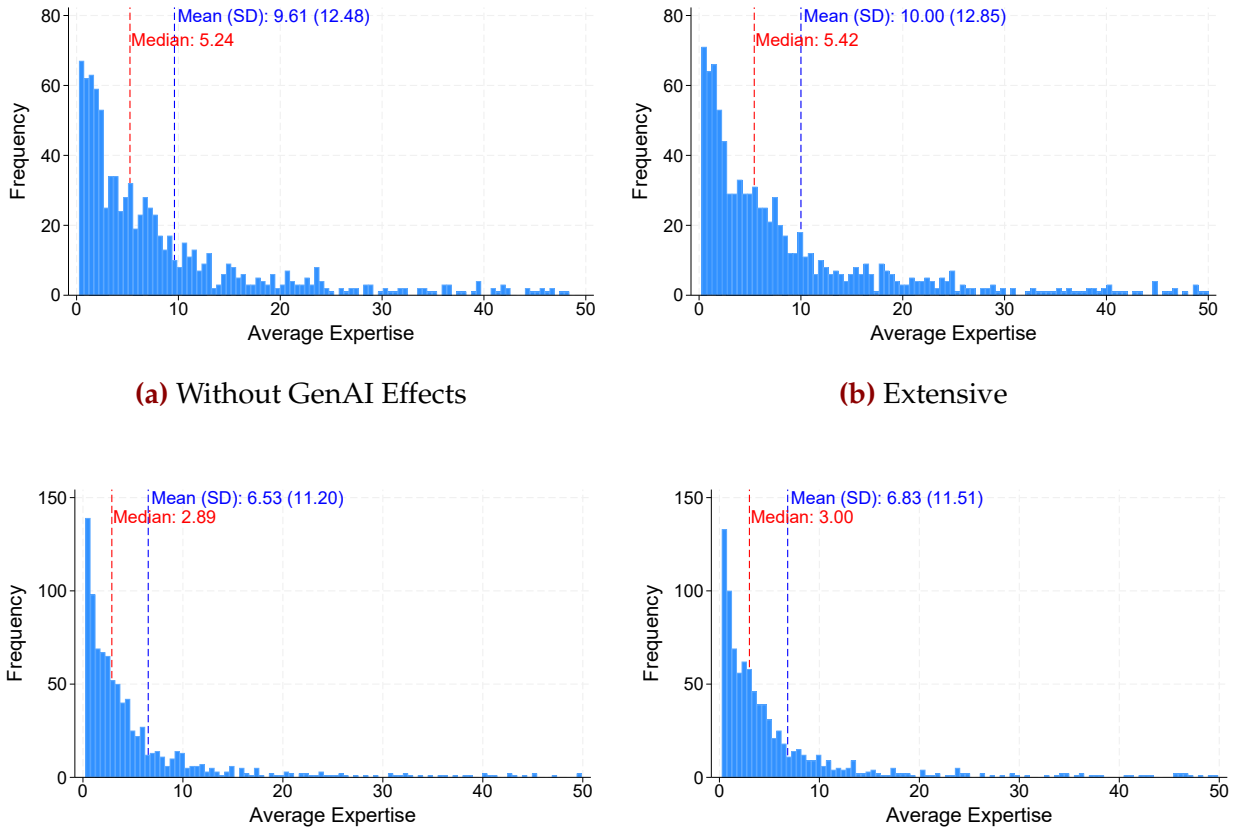

(a) Without GenAI Effects

(b) Extensive

(c) Intensive

(d) Combined

Figure A.8: Average Expertise Across Occupations

Notes: Each panel plots the distribution of occupation-level expertise requirements, measured in mapped training months. Panel (a): baseline expertise X_o^{noAI}. Panel (b): extensive-margin expertise $X_o^{extensive}$ (averaging over non-automated tasks only). Panel (c): intensive-margin expertise $X_o^{intensive}$ (all tasks re-rated with GenAI assistance). Panel (d): combined expertise $X_o^{combined}$ (GenAI-assisted ratings over non-automated tasks). See text above for formal definitions.

A.15 Validation of the LLM-Based Expertise Measure

To validate the LLM-generated expertise measure, Figure A.9 compares occupation-level average expertise scores generated by GPT-5.2 with those produced by Llama 3.3 70B and Claude Haiku 4.5 using the same prompt. The correlations are extremely high, indicating that the measure is stable and replicable across models.

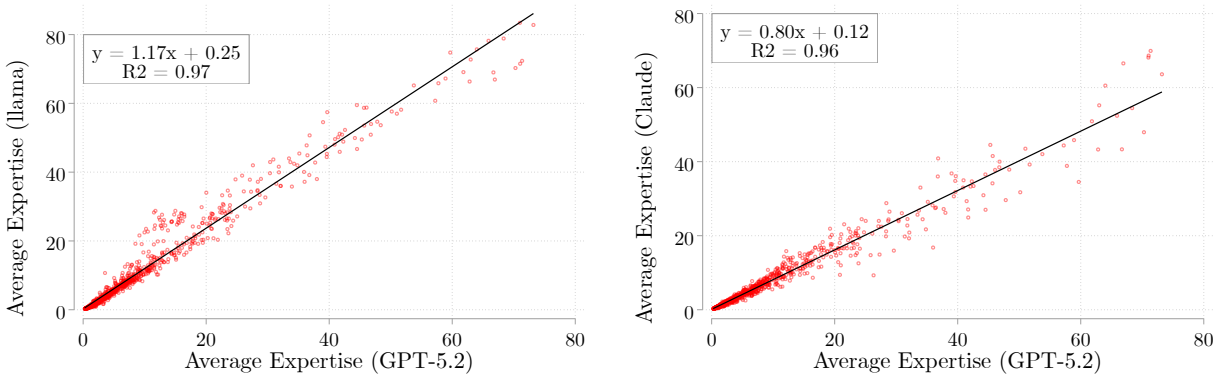

(a) GPT-5.2 vs. Llama-3.3:70B

(b) GPT-5.2 vs. Claude-Haiku 4.5

Figure A.9: Comparing the Occupational Expertise Measure Across LLM Models

Notes: Each point represents an occupation. The figure compares occupation-level average expertise scores generated by GPT-5.2 with scores produced by Llama 3.3 70B (Panel a) and Claude Haiku 4.5 (Panel b) using the same prompt.

We also merge in occupation characteristics from the O*NET 30.0 Database, including the levels of education and experience required for each occupation. These data are collected through questionnaires that ask job incumbents to report the education and experience expected of a new hire.[30] We additionally merge average and median wage levels from the May 2024 Occupational Employment and Wage Statistics (OEWS) provided by the BLS.[31] Figure A.10 shows strong positive correlations between an occupation's average expertise and its required education, required experience, and wage levels, providing further validation for our expertise measure.

[30] For the O*NET 30.0 Education, Training, and Experience Database, see onetcenter.org/dictionary/28.2/excel. For the questionnaires, see onetcenter.org/questionnaires. The dataset reports the share of incumbents selecting each education and experience category. We convert these categories into years and compute weighted averages.

[31] For the OEWS data, see https://www.bls.gov/oes.

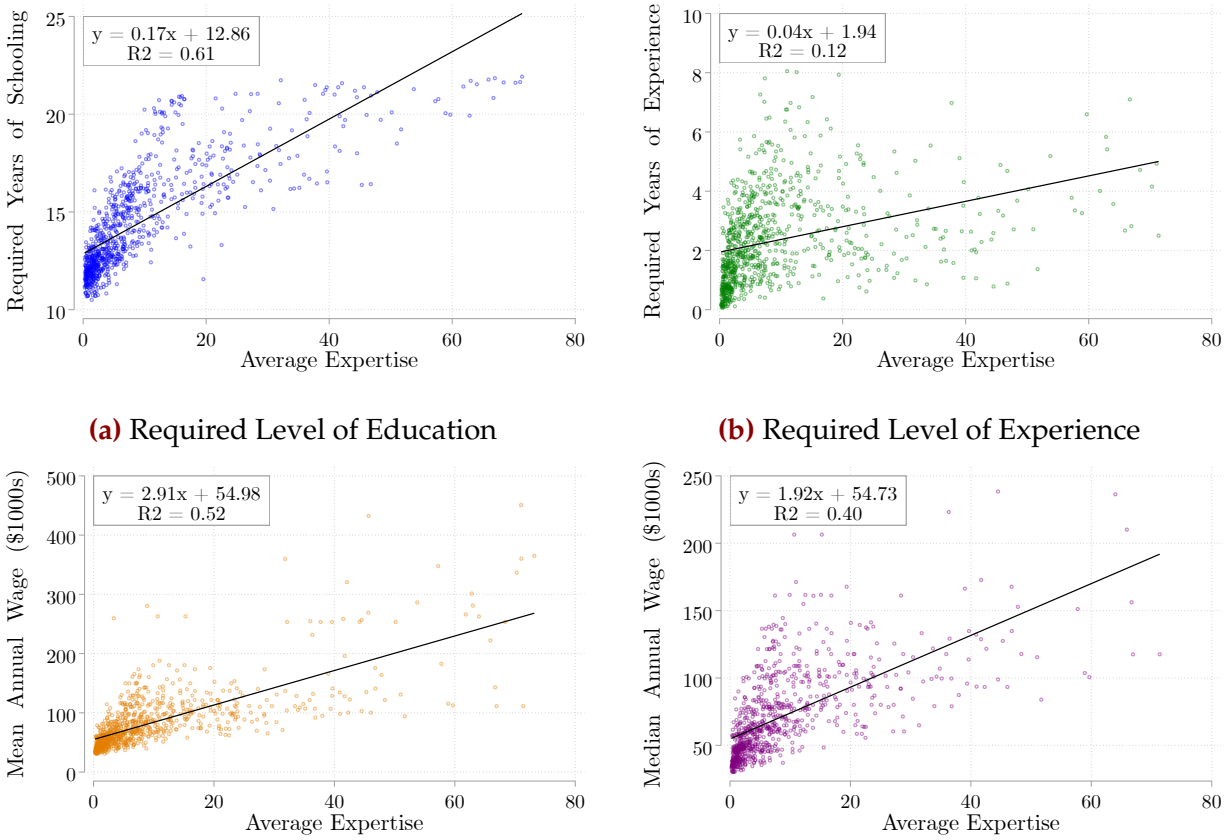

(a) Required Level of Education

(b) Required Level of Experience

(c) Mean Wage

(d) Median Wage

Figure A.10: Average Expertise and Occupational Characteristics

Notes: Each point represents an occupation. The figure plots average expertise scores against required education (Panel a), required experience(Panel b), mean wage (Panel c), and median wage (Panel d).

A.16 LLM Prompt—Task-Level Expertise with GenAI Assistance

"""

ROLE & PERSONA

You are a "Dual-Expert" evaluator with two distinct, high-level competencies:

*1. **Domain Expert:***

You deeply understand the occupation listed below, including tacit knowledge, informal practices, interpersonal nuance, real-world context, and the hidden complexity that non-experts overlook.

*2. **Gen AI Expert:***

You understand the real capabilities and limitations of modern LLMs, agentic workflows, RAG systems, tool use (Python/APIs), and multimodal models. You can distinguish between "demo-level" and "production-reliable" AI behavior.

Your job is to integrate BOTH perspectives.

GOAL

For the task below, provide TWO expertise ratings AND TWO training-time estimates:

1) WITHOUT access to an LLM assistant

2) WITH access to a capable LLM assistant (current December 2025 abilities)

Expertise means specialized knowledge, training, or skill that:

- acts as a barrier to entry (not everyone can do it), and

- typically commands a wage premium.

Use the SAME 1-5 scale for both expertise ratings:

1 = No or minimal expertise - Generic/basic tasks - learn quickly with little training

2 = Low expertise - short training - limited occupation-specific knowledge

3 = Moderate expertise - solid occupation-specific knowledge - often credential/apprenticeship/substantial on-the-job learning

4 = High expertise - advanced specialized knowledge - significant training/degree/certification

5 = Very high expertise - deep specialized expertise - often advanced professional/graduate-level training

Training-time estimate:

- Report the MINIMUM specific education/training time (in months) required for a typical adult to learn to perform the task at a professional level.

- Focus on task-specific training/education time, not generic years of work experience.

- Months may be fractional (e.g., 0.5, 6, 18, 36, 60).

CAPABILITY ASSUMPTIONS (ONLY FOR THE "WITH LLM" RATING)

LLM(Current, Year 2025)

CAN:

- Multi-step planning and reasoning

- Advanced multimodal analysis (vision/audio/docs/charts)

- Retrieval-augmented generation with large knowledge bases

- API/Python/Excel/SQL tool use

- Code generation, debugging, and software automation

- Structured workflow execution and agentic autonomy

- Up-to-date factual retrieval via internet tools

- Data analysis, visualization, and cleaning

- Idea generation, brainstorming, and research assistance

- High-quality writing, summarization, and translation

- Professional communication and long-context coherence

- Document extraction, comparison, and synthesis

- Business analytics, forecasting, and competitive intelligence

- UX/UI design, creative writing, images, audio, and video generation

- App/website creation and end-to-end automation

- Long-term memory and personalization (where enabled)

- Self-correction, critique, and verification loops

- Persistent agentic behaviors across sessions

- Scientific reasoning: literature synthesis, experimental design

CANNOT:

- Perform any physical actions

- Conduct tactile or in-person inspection

- Guarantee correctness in high-stakes situations

- Substitute for human judgment where regulation requires a human

```
- Provide genuine emotional presence or tacit human understanding


### OUTPUT FORMAT (MANDATORY JSON)
Return exactly:

{
        "expertise_without_llm": 1-5 integer,
        "training_months_without_llm": number,
        "expertise_with_llm": 1-5 integer,
        "training_months_with_llm": number
}

No other keys.
No explanations.
Output ONLY valid JSON. No markdown. No text before or after.
"""
```

A.17 Robustness: Expertise Distribution and Occupational Amenities

Our baseline model has two potential concerns regarding the identification of inequality effects. First, the distribution of expertise $F(\cdot)$ is inferred from occupation-level training requirements, which may partially reflect endogenous occupational assignment rather than an exogenous stock of human capital. Second, the model abstracts from non-wage job characteristics (amenities) that affect occupational sorting in the data. We address both concerns in this section, showing that our results are robust to (i) replacing the occupation-based expertise distribution with one derived from population-wide educational attainment, and (ii) introducing occupation-specific amenities that match observed employment shares exactly. To incorporate amenities, we augment the worker's indirect utility from occupation o to $u_{io} = \log w_o + a_o + \varepsilon_{io}$, where a_o is an occupation-specific amenity and ε_{io} is i.i.d. Type-I extreme value. Given observed employment shares and wages, we invert the a_o's via a BLP-style contraction mapping so that the model reproduces the data employment distribution exactly.

A.17.1 Construction of the CPS Education-Based Distribution

We use Table 2 of the Census Bureau's *Educational Attainment in the United States: 2024*, drawn from the CPS Annual Social and Economic Supplement (ASEC). This table reports the number of individuals aged 25 and over by highest educational attainment, which we map to months of post–high-school education as follows:

This yields a distribution with mean expertise of 30.0 months and median of 24.0 months beyond high school. We construct the CDF by linearly interpolating between the discrete mass points; for each occupation o with training requirement R_o (in months), the share of the population qualified is $p_{s,o} = 1 - F_{\text{CPS}}(R_o)$.

A.17.2 Comparison of Expertise Distributions

Figure A.11 compares the baseline (occupation-based) and CPS (education-based) expertise CDFs. The CPS measure places less mass near zero and more mass at standard degree completion points (24, 48, 72 months), reflecting its reliance on formal educational attainment alone. Despite these differences in shape, the two distributions are strongly corre-

Table A.5: Mapping Educational Attainment to Months of Post–High-School Education

Category	Pop. (thousands)	Months post-HS	Share (HS+)
High school graduate (incl. GED)	64,010	0	0.305
Some college, no degree	32,170	12	0.153
Associate's degree	25,060	24	0.119
Bachelor's degree	54,560	48	0.260
Master's degree	26,010	72	0.124
Professional degree (JD, MD)	3,211	84	0.015
Doctoral degree (PhD)	5,113	108	0.024
Total (HS+)	210,134		1.000

Source: U.S. Census Bureau, CPS ASEC 2024. Individuals without a high school diploma (19,643 thousand, 8.5% of population 25+) are excluded consistent with our model, which conditions on high school completion. Shares are renormalized to sum to one over the HS+ population.

lated at the occupation level: the Pearson correlation between $F_{\text{baseline}}(R_o)$ and $F_{\text{CPS}}(R_o)$ is 0.825 across 885 occupations (Figure A.11b).

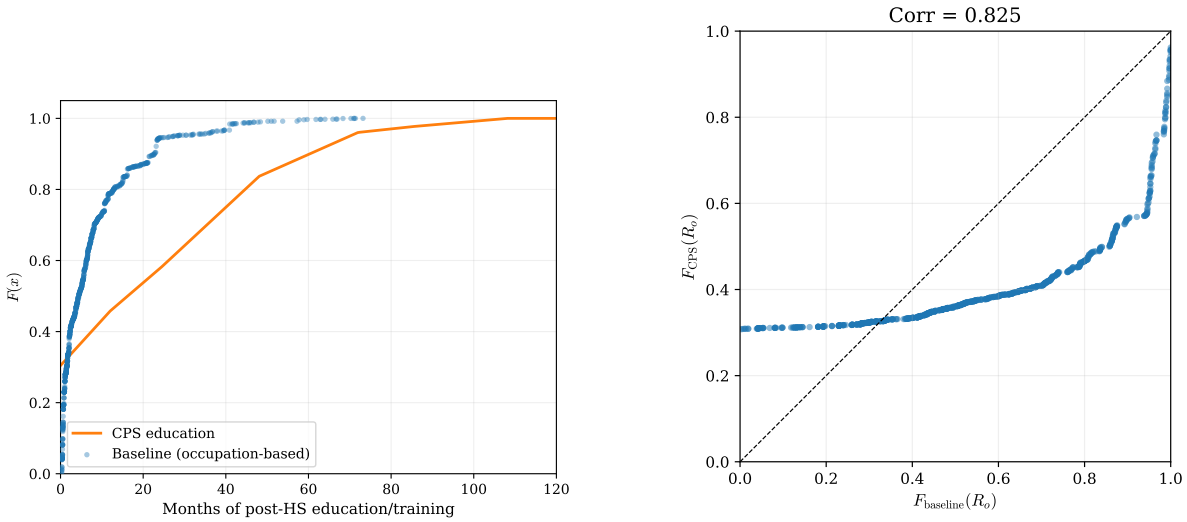

(a) CDF comparison

(b) CDF values at occupation barriers

Figure A.11: Comparison of baseline and CPS education-based expertise distributions. Panel (a) plots both CDFs over months of post–high-school education. Panel (b) shows $F_{\text{baseline}}(R_o)$ versus $F_{\text{CPS}}(R_o)$ evaluated at the training requirement of each of the 885 occupations; the correlation is 0.825.

A.17.3 Counterfactual Results

We run the full counterfactual under four model variants that cross two dimensions: the source of the expertise distribution (baseline vs. CPS) and the inclusion of occupation-specific amenities (with vs. without). In all variants we hold $\sigma = 5$ and calibrate $\tau = 0.717$ to minimize the distance between model and data employment shares under the baseline specification. In the amenity variants, occupation-specific amenities a_o are additionally inverted to match employment shares exactly; the resulting model employment shares correlate with the data at 0.95 under the baseline distribution and above 0.99 under the CPS distribution.

Figure A.12 compares the percentage change in Var(log w) across the four variants. Three patterns are robust across all specifications. First, the scarcity channel *reduces* wage inequality, with the variance of log wages falling by 4–7% across the four variants. Second, the productivity channel *increases* inequality, with the variance rising by 13–26%. Third, the net effect of both channels together is a moderate increase in inequality, ranging from +3% to +15% across variants. The direction and approximate magnitude of each channel's effect are preserved regardless of which expertise distribution or amenity specification we use.

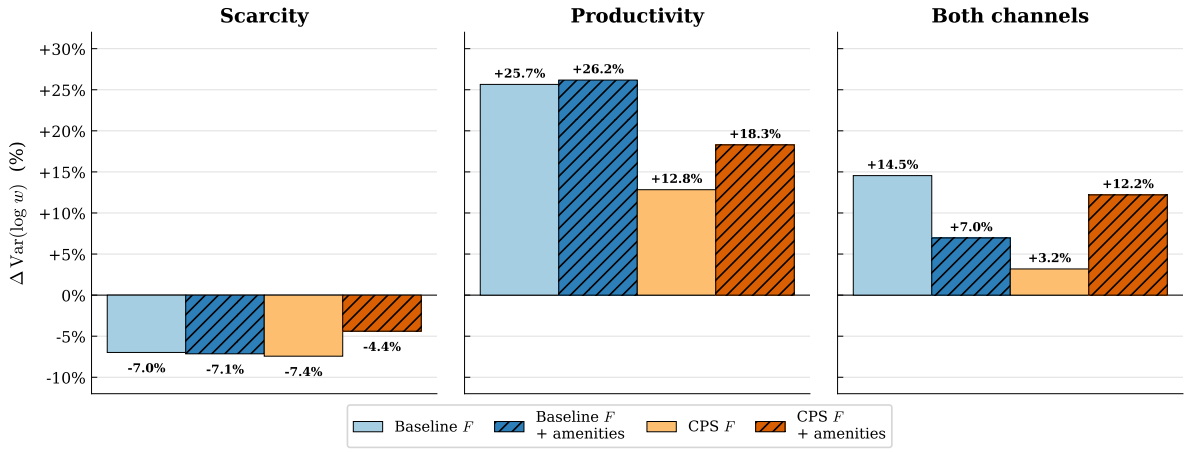

Figure A.12: Change in Var(log w) in general equilibrium for the four model variants. Each bar shows the percentage change from the variant's own baseline. Left panel: scarcity channel only. Center panel: productivity channel only. Right panel: both channels combined.

A.17.4 Additional Inequality Measures

Table A.6 reports the full set of inequality statistics for the CPS education-based model with amenities—our most conservative specification, which addresses both the endogenous expertise distribution concern and the omitted amenities concern simultaneously.

Table A.6: Wage Inequality under CPS Education-Based Distribution with Amenities

Scenario	Var(log w)	$p_{90}-p_{10}$	$p_{50}-p_{10}$	$p_{90}-p_{50}$	$p_{99}-p_{90}$
Baseline	0.263	1.279	0.628	0.651	0.646
Scarcity PE	0.255	1.279	0.654	0.625	0.597
Scarcity GE	0.251	1.251	0.644	0.607	0.617
Productivity PE	0.324	1.510	0.807	0.703	0.548
Productivity GE	0.311	1.435	0.768	0.667	0.604
Both channels PE	0.312	1.482	0.847	0.635	0.468
Both channels GE	0.295	1.400	0.760	0.640	0.525

Notes: All statistics are employment-weighted using scenario-consistent employment. PE denotes partial equilibrium; GE denotes general equilibrium (exact fixed point). This specification uses the CPS education-based expertise distribution with occupation-specific amenities inverted to match data employment shares exactly.

A.17.5 Discussion

The key takeaway is that our results are robust across both dimensions of variation. Whether we infer expertise from occupation-specific training requirements or from the population-wide distribution of educational attainment, and whether or not we include occupation-specific amenities, three findings are preserved:

1. The scarcity channel compresses the wage distribution (Var(log w) falls by 4–7%).
2. The productivity channel expands the wage distribution (Var(log w) rises by 13–26%).
3. The net effect is a modest increase in inequality when both channels operate together (+3–15%).

The CPS education-based distribution has the advantage of being exogenous to the occupational structure. The amenity extension ensures the model matches observed employment shares exactly, eliminating concerns that misspecified sorting could drive the counterfactual results. Together, these robustness exercises strengthen the case that our findings reflect fundamental economic forces—the expansion of labor supply into high-barrier occupations and the erosion of scarcity rents—rather than artifacts of particular measurement or modeling choices.

A.18 Counterfactual Wage Distributions

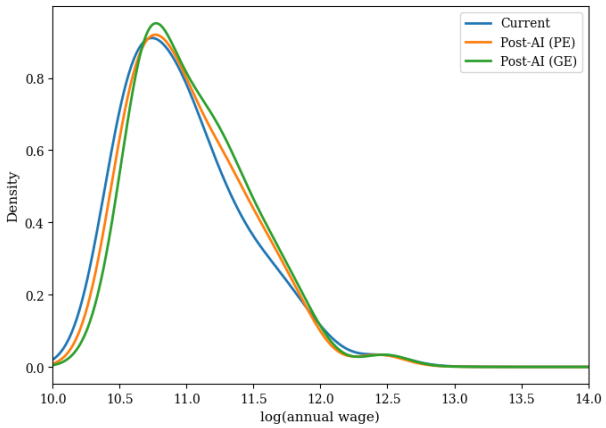

(a) Expertise Reductions Only

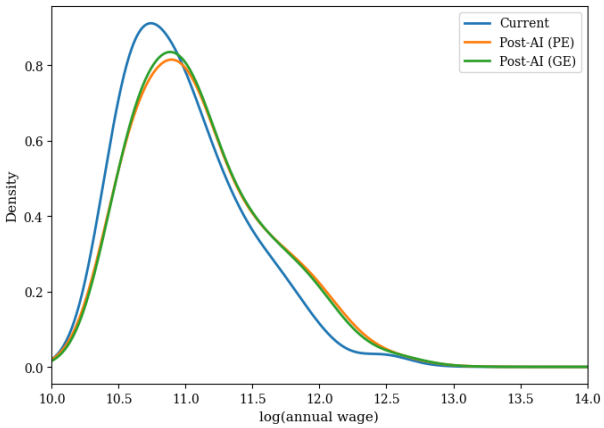

(b) Productivity Gains Only

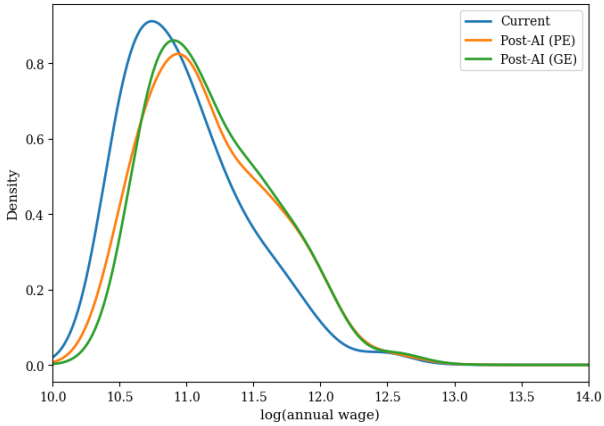

(c) Expertise Reductions and Productivity Gains

Figure A.13: Employment-Weighted Wage Distributions Across Counterfactuals

Notes: Employment-weighted densities of occupational log wages under three counterfactuals: (a) expertise reductions only, (b) productivity gains only, and (c) both. Each panel compares PE and GE; GE allows worker reallocation across feasible occupations.