



**DISCUSSION PAPER SERIES**

**039/26**

# **The Labor Market Impact of Occupation-Specific Technical Change: Inspecting the Mechanisms**

Fenella Carpena, Simon Galle

# The Labor Market Impact of Occupation-Specific Technical Change: Inspecting the Mechanisms

## Authors

Fenella Carpena, Simon Galle

## Reference

**JEL Codes:** J23, J24, E24, E25

**Keywords:** Technical change, labor reallocation, wage inequality, occupational heterogeneity, artificial intelligence, large language models

**Recommended Citation:** Fenella Carpena, Simon Galle (2026): The Labor Market Impact of Occupation-Specific Technical Change: Inspecting the Mechanisms. RFBerlin Discussion Paper No. 039/26

## Access

Papers can be downloaded free of charge from the RFBerlin website: <https://www.rfberlin.com/discussion-papers>

Discussion Papers of RFBerlin are indexed on RePEC: <https://ideas.repec.org/s/crm/wpaper.html>

## Disclaimer

*Opinions and views expressed in this paper are those of the author(s) and not those of RFBerlin. Research disseminated in this discussion paper series may include views on policy, but RFBerlin takes no institutional policy positions. RFBerlin is an independent research institute.*

*RFBerlin Discussion Papers often represent preliminary or incomplete work and have not been peer-reviewed. Citation and use of research disseminated in this series should take into account the provisional nature of the work. Discussion papers are shared to encourage feedback and foster academic discussion.*

*All materials were provided by the authors, who are responsible for proper attribution and rights clearance. While every effort has been made to ensure proper attribution and accuracy, should any issues arise regarding authorship, citation, or rights, please contact RFBerlin to request a correction.*

*These materials may not be used for the development or training of artificial intelligence systems.*

## Imprint

### RFBerlin

ROCKWOOL Foundation Berlin –  
Institute for the Economy  
and the Future of Work

Gormannstrasse 22, 10119 Berlin  
Tel: +49 (0) 151 143 444 67  
E-mail: [info@rfberlin.com](mailto:info@rfberlin.com)  
Web: [www.rfberlin.com](http://www.rfberlin.com)



# The Labor Market Impact of Occupation-Specific Technical Change: Inspecting the Mechanisms

Fenella Carpene

OsloMet

Simon Galle\*

BI Norwegian Business School

January 2026<sup>†</sup>

First version: October 2024

## Abstract

Motivated by the rise of artificial intelligence (AI), we set up a quantitative general-equilibrium model of the labor market impact of occupation-specific technical change. The highly tractable model crystallizes how three fundamental forces shape the impact of technical change on the occupational wage distribution: the input substitution elasticity, the final demand elasticity, and the labor supply (reallocation) elasticity. The difference between the former two elasticities determines whether machines and workers are gross complements, while the reallocation elasticity governs the magnitude of the distributional effects. We estimate the reallocation elasticities from group-by-occupation specialization changes, allowing for asymmetric reallocation and associated ripple effects on wages across occupations. After combining these estimates with externally disciplined demand and substitution parameters, as well as AI exposure measures, we shed light on the aggregate and distributional effects of occupation-specific advances in AI: wages in administrative services grow the least, ripple effects on less exposed occupations are substantial, AI modestly compresses the returns to education, and, on average, disproportionately benefits lower-income groups.

JEL: J23, J24, E24, E25

KEYWORDS: Technical change, labor reallocation, wage inequality, occupational heterogeneity, artificial intelligence, large language models

---

\*Corresponding author: [simon.galle@bi.no](mailto:simon.galle@bi.no).

<sup>†</sup>Previous versions of this paper were circulated under a shorter title, leaving out "occupation-specific." We are deeply grateful to Henrikke Gedde Rustad and Jenny Kolstad for excellent research assistance, and to Andrea Eisfeldt, Gregor Schubert, and Ben Miao Zhang for sharing their data. We thank Paweł Gola, Daniel Gross, Jochen Guntrner, Jonas Hamang, Svenn Jensen, Attila Lindner, Linnea Lorentzen, Katja Mann, Espen Moen, Peter Morrow, Plamen Nenov, Andrés Rodríguez-Clare, Ovidijus Stauskas, Alisa Tazhitdinova, Jose Pablo Vasquez, Jonathan Vogel, Yuan Zi, and seminar audiences at Aarhus, BI, CORA at Goethe University, Erasmus SE Rotterdam, Ghent, JKU Linz, Kristiania, OsloMet, and University of Oslo for insightful comments and valuable feedback. Galle acknowledges support from a Basic Research Funding grant by BI. All errors are our own.

## 1 Introduction

Technological change has long reshaped work through occupation-specific channels. Every wave of labor-saving innovation revives the now familiar tension between techno-optimists, who emphasize productivity gains, and techno-pessimists, who warn of job displacement and income loss. History suggests both forces often operate simultaneously. During the First Industrial Revolution, for instance, mechanisation displaced skilled artisans such as the Luddites, while others benefited from cheaper and more abundant textiles. A similar pattern emerged in the late twentieth century, when computerization reduced employment in routine occupations, including administration and bookkeeping. More recently, robotization has transformed manual production, and advances in generative artificial intelligence, in particular large language models (LLMs), have renewed job security concerns among programmers and other knowledge workers. These episodes underscore that the incidence of technological change is inherently occupational (see, e.g., [Burstein, Morales, and Vogel 2019](#); [Caunedo, Jaume, and Keller 2023](#); [Feigenbaum and Gross 2024](#)). Understanding this heterogeneity, in turn, is critical to characterize the incidence of gains and losses from new technologies, in the absence of redistribution.

This paper studies how occupation-specific technical change affects labor market outcomes in general equilibrium. The framework is simple and tractable, but designed to be portable across settings by focusing on a limited set of generally-applicable mechanisms. The model's key feature is an explicit and transparent joint adjustment of labor demand and labor supply to occupation-level technology shocks. The extent of worker displacement depends, of course, on the nature of the labor-demand response: specifically, whether substitution or scale effects dominate. Critically, the labor-demand response is just one side of the story, since workers reallocate in response to shifts in labor demand, dampening the impact of the shock for the stayers in their origin occupation and generating ripple effects on wages in other occupations. The interactions between labor demand and supply therefore jointly determine both the magnitude of displacement among affected workers and the impact on the overall distribution of wage changes. Since the model also tracks the productivity gains associated with technical change, it delivers a unified mapping from occupation-level shocks to aggregate outcomes and distributional incidence.

We parsimoniously integrate the above mechanisms in the framework through a concise set of elasticities. On the labor-*demand* side, we incorporate the fundamental interplay between input substitution and scale effects. Here, the within-occupation substitution elasticity ( $\sigma_o$ ) governs whether new technologies substitute for or complement

workers. This degree of substitutability is a critical aspect of any labor-market analysis of technical change. Moreover, our formulation accommodates the full range from almost perfect substitutability to complementarity. This flexibility matters, since some technologies, including LLMs, may complement labor in some occupations while substituting for it in others.

In turn, the final-demand elasticity  $\psi$  (in absolute value) determines how lower unit costs translate into expenditure on an occupation's output: a classic scale effect that can amplify or offset the within-occupation substitutability (Hicks, 1932; Robinson, 1933). Indeed, even if improved equipment is a substitute for workers, overall employment in an occupation might still increase if demand for the occupation's output picks up sufficiently.

On the labor-supply side, reallocation elasticities capture how easily workers move across occupations – a force critical for governing any unequal effects associated with worker displacement. Indeed if workers are equally talented in all occupations, labor displacement in one occupation is irrelevant from a distributional perspective. We microfound the reallocation elasticities in the dispersion of workers' comparative advantage: with low dispersion, workers reallocate easily and distributional effects are small; with stronger dispersion, reallocation is costly and effects are larger. Assuming Fréchet heterogeneity in abilities yields a constant reallocation elasticity (Lagakos and Waugh, 2013), spanning the extremes from perfectly elastic to perfectly inelastic occupational labor supply. Furthermore, to capture asymmetric mobility, we allow different elasticities within and across nests of occupations,  $\kappa$  and  $\mu$ , which generates differential wage spillovers (or ripple effects) when displaced workers disproportionately enter some occupations rather than others (Acemoglu and Restrepo, 2022; Lorentzen, 2024). Incorporating these asymmetries is therefore critical for capturing the uneven distributional implications of worker displacement.

This concise, tractable, and flexible combination of labor demand and supply allows crystallizing how capital-embodied technical change influences the distribution of relative wages. To inspect the central mechanisms, we first focus on a simplified case with a single reallocation elasticity  $\kappa$  and a fixed substitution elasticity  $\sigma$ . In that setup, we show that the log change in the relative wage of an exposed occupation vis-à-vis any other occupation is proportional to the log change in their relative unit costs, with elasticity

$$\frac{\sigma - \psi}{\sigma + \kappa - 1}.$$

The numerator is the net demand force, measuring gross substitutability between workers and machines. If input substitution dominates ( $\sigma > \psi$ ), a technology-induced cost de-

cline in  $\sigma$  reduces that occupation's relative wage; if scale expansion dominates ( $\psi > \sigma$ ), the relative wage will rise even when the cost share of labor falls ( $\sigma > 1$ ). In the denominator, the reallocation elasticity governs the magnitude of the distributional effects: these are strongest when  $\kappa \rightarrow 1$  (specific factors), while there is no wage inequality when  $\kappa \rightarrow \infty$  (workers are perfectly mobile); as in [Galle, Rodríguez-Clare, and Yi \(2023\)](#).

We use simulations to further illustrate the mechanisms of the model and which distributional patterns it can capture. The starting point is a productivity shock to machines in a specific occupation, examined for varying elasticities ( $\sigma, \psi, \kappa, \mu$ ). In addition to visualizing the above result, we illustrate the feedback loop between production cost, input prices, final goods prices, and expenditure shares. For instance, we show that the feedback loop between production, wages, and demand may lead to a *benign* labor-share decline: a higher final-demand elasticity  $\psi$  puts upward pressure on wages, which leads to input substitution away from labor. This potential for a benign labor-share decline appears underemphasized in the literature (see, e.g., [Karabarbounis 2024](#)). A second interesting pattern is that the expenditure share on an occupation, a *final-demand* object, grows as the *input*-substitution elasticity ( $\sigma$ ) rises, driven by the latter's downward pressure on an occupation's wage. The two previous patterns are both intuitive, but are only uncovered when wages are occupation-specific—a central feature of our model, incorporated in a tractable and transparent manner.

Finally, the same exercises overturn a common presumption that machine-augmenting shocks cannot produce falling real wages for important demographic subgroups. Indeed, our model does obtain wage stagnation or declines for plausible degrees of gross substitutability, and this pattern is further amplified by limited occupational mobility. This way, the simulations are able to capture the canonical patterns of wage polarization and clarify when exposure to technical change yields displacement versus expansion, or wage declines versus wage growth.

After illustrating the model, we estimate the reallocation elasticities with a transparent IV approach. In particular, we estimate the within-nest ( $\kappa$ ) and across-nest ( $\mu$ ) elasticities, which are critical for governing the ripple effects on wages, and we do so separately for young, middle-aged, and old workers. Here, nests are defined as sets of routine and non-routine intensive occupations, based on the dichotomy of [Autor and Dorn \(2013\)](#). Aligned with common intuitions, we find that reallocation is typically more elastic within than across nests, and that worker mobility tends to fall with age.

Because our framework is organized around a few interpretable elasticities, the machinery is portable to various settings, including historical waves of technical change, such as computerization ([Burstein et al., 2019](#)) or capital-embodied technical change more broadly ([Caunedo et al., 2023](#)). To illustrate the relevance of the model, we use

it to shed light on the GE impact of occupation-specific advances in artificial intelligence (AI), more specifically LLMs such as ChatGPT, Claude, or various others. Given the very recent rise of these models, an estimation of the model's parameters based on LLM-induced shocks is outside the scope of the current paper. Nevertheless, whenever possible, our calibration builds on findings from the emerging reduced-form literature on this topic, which finds displacement of workers in highly exposed occupations (see, e.g., [Brynjolfsson, Chandar, and Chen, 2025](#)), in particular in occupations where LLM use is more automation rather than augmentation oriented (measured as in [Handa, Tamkin, McCain, Huang, Durmus, Heck, Mueller, Hong, Ritchie, Belonax, et al., 2025](#)).

Our findings are intuitive and matter economically. First, all occupations experience real wage gains, driven by the broad applicability of LLMs across the majority of occupations ([Eloundou, Manning, Mishkin, and Rock, 2024](#)). Second, although everyone gains in absolute terms, scientific technicians and workers in administrative services gain the least. Both occupations combine high exposure to LLMs with fairly high gross substitutability of LLMs with human labor. Third, our estimated labor-supply elasticities imply that employment reallocation is a stronger margin of adjustment than inequality in wage changes. In other words, the ripple effects on wages from high to low-exposed occupations are substantial and compress the unequal effects of the introduction of LLMs.

To offer a more comprehensive analysis of the labor market impacts of technical change, we extend our model with frictional unemployment and intensive margin adjustment, following [Kim and Vogel \(2021\)](#). In this setup, we obtain an aggregate real income growth of 8.7%. Approximately 20% of this growth is accounted for by the increase in hours worked, while 30% is driven by a rise in the employment rate. Notably, income changes resulting from the rise of LLMs show a pro-poor distribution across demographic groups. Specifically, income changes exhibit a negative correlation of 27% with initial income levels. Moreover, AI has a modest negative effect on the education premium, with high-school dropouts experiencing the largest gains in real income. These findings indicate that LLMs have a labor market impact that strongly contrasts with previous episodes of “skill-biased” technical change.

**Literature** Our study is inspired by seminal papers on the distributional effects of technical change in a setting with labor reallocation (e.g., [Autor, Levy, and Murnane, 2003](#); [Acemoglu and Restrepo, 2022](#)). We generalize these papers by considering the full spectrum of substitutability between human and machine labor instead of focusing only on perfect substitutability. Moreover, our Roy-Fréchet setup tractably generalizes the classic Ricardo-Roy reallocation framework with step-wise occupational labor supply functions ([Acemoglu and Autor, 2011](#); [Acemoglu and Restrepo, 2022](#)) by allowing for strictly

upward-sloping labor supply to any discrete number of occupations.<sup>1</sup> <sup>2</sup> Finally, in contrast to models with two or three worker groups (e.g., low, middle, and high skill as in [Acemoglu and Autor, 2011](#)), we allow for any number of groups, allowing for a detailed demographic breakdown of occupational specialization.

This model of occupation-specific technical change with tractable labor reallocation across any number of occupations builds on the foundational work of [Burstein et al. \(2019\)](#), and [Caunedo et al. \(2023\)](#) – henceforth BMV and CJK respectively – as well as on the closely-related framework in [Galle and Lorentzen \(2024\)](#).<sup>3</sup> In contrast to BMV who use Cobb-Douglas, CJK has a general CES production function at the occupation level. This way, the CJK framework nests workhorse models such as [Katz and Murphy \(1992\)](#); [Krusell, Ohanian, Ríos-Rull, and Violante \(2000\)](#) and [Acemoglu and Autor \(2011\)](#).<sup>4</sup>

In this paper, we inspect the interplay between the central mechanisms in the BMV-CJK model. First, we pin down how the impact of technical change critically depends on the confluence of the input substitution, final demand, and labor supply elasticities. Second, we provide detailed quantitative comparative statics on the impact of these key elasticities. And third, we apply our framework to the rise of generative AI, more specifically LLMs, to shed light on the GE labor market impacts of this new and growing technology.

We also further generalize the BMV-CJK model by allowing for ripple effects arising from asymmetric reallocation. Such ripple effects have been extensively documented in the local labor market literature ([Beaudry, Green, and Sand, 2012](#); [Fortin and Lemieux, 2015](#); [Tschopp, 2015](#)), and have recently been incorporated in labor-market models of technical change ([Acemoglu and Restrepo, 2022](#); [Ocampo Díaz, 2022](#)).<sup>5</sup> In the context of

---

<sup>1</sup>This way, we span the two extremes of perfectly inelastic and perfectly elastic labor supply, considered in [Caselli and Manning \(2019\)](#), and allow for a quantitative analysis of the intermediate cases.

<sup>2</sup>Recent related work employing a [Roy \(1951\)](#) model includes [Gola \(2021\)](#), who examines how the interplay of Roy-type selection across sectors and within-sector firm heterogeneity shapes the distributional effects of technical change, and [Beron and Magerman \(2024\)](#), who study the role of the input-output network in affecting income inequality after sector-specific productivity shocks.

<sup>3</sup>BMV and CJK assume CES preferences, while [Galle and Lorentzen \(2024\)](#) has Cobb-Douglas preferences across sectors and includes a gravity framework for world trade. This enables [Galle and Lorentzen \(2024\)](#) to compare the impact of trade and automation on US manufacturing. In a recent contribution, [Adachi \(2025\)](#) also examines the interplay of robotization and trade and estimates the input substitution elasticity between robots and labor. Relatedly, [Berlingieri, Boeri, Lashkari, and Vogel \(2024\)](#) employ a flexible framework to study how skill-bias in productivity at the firm level aggregates up to aggregate capital-skill complementarity. In the context of migration instead of technical change, [Burstein, Hanson, Tian, and Vogel \(2020\)](#) set up a model closely related to ours.

<sup>4</sup>Another closely related paper is [Humlum \(2022\)](#), who examines the labor market impacts of robotization. Compared to [Humlum \(2022\)](#), the model in our paper is more stylized, and we put heavy emphasis on understanding the interplay of the model's elasticities.

<sup>5</sup>Relatedly, minimum wage regulations have also been shown to generate substantial wage spillover effects ([Giupponi, Joyce, Lindner, Waters, Wernham, and Xu, 2024](#)), as has the decentralization of one occupation's wage determination ([Willén, 2021](#)).

displacement after an oil price drop, [Lorentzen \(2024\)](#) combines clean reduced-form evidence with novel structural methods to show that asymmetric ripple effects significantly influence the wage distribution in Norway.<sup>6</sup>

Finally, by dividing the population of workers into any finite number of demographic groups ([Hsieh, Hurst, Jones, and Klenow, 2019](#)), this type of model yields detailed insights on between-group inequality. Indeed, closely related to BMV, the model can yield insights on the college premium or the (non-causal) “returns to education,” the gender wage gap, the impact of technical change for older versus younger workers, or any combination of the above. Two contributions of our paper are to, first, further generalize this versatile model, and second, to shed light on the underlying mechanisms in the model.

## 2 Model

### 2.1 Setup

We write down a model with occupation-specific labor supply and demand. Labor demand is governed by a CES demand function across occupations, and a CES production function in each occupation, with machines and labor as inputs. Labor supply across occupations has a nested constant-elasticity setup, arising from a Roy-Fréchet foundation. The economy is perfectly competitive.

**Labor demand** The final good is produced by aggregating occupation outputs with a CES technology (as in BMV):

$$\tilde{Y} = \left[ \sum_o \nu_o^{\frac{1}{\psi}} Y_o^{\frac{\psi-1}{\psi}} \right]^{\frac{\psi}{\psi-1}},$$

where  $Y_o$  is the output produced by occupation  $o$ ,  $\nu_o$  is a demand shifter for this occupation, and  $\psi > 0$  is the elasticity of substitution across occupations’ output. The final good has a price  $P$ . It is used both for consumption and to produce machine input. We model machines as linearly produced from the final good, so the unit price of machine input is  $P$  and total machine use is  $M = \sum_o M_o$ .

---

<sup>6</sup>Specifically, [Lorentzen \(2024\)](#) draws on reduced-form microdata evidence and presents a log-normally distributed Roy model to generate rich reallocation patterns. In this paper, we focus on *constant* reallocation elasticities within and across nests, as we prioritize tractability in a less rich but more transparent framework. These constant reallocation elasticities arise from a Roy-type discrete choice model with extreme-value distributed heterogeneity ([Lagakos and Waugh, 2013](#); [Curuk and Vannoorenberghe, 2017](#)), where we introduce a nesting structure similar to [Kim and Vogel \(2021\)](#); [Zárate \(2022\)](#) and [Galle et al. \(2023\)](#) to capture asymmetric reallocation within versus across nests.

Each occupation produces an output

$$Y_o = \left[ \delta_o^{\frac{1}{\sigma_o}} (\gamma_o M_o)^{\frac{\sigma_o-1}{\sigma_o}} + (1 - \delta_o)^{\frac{1}{\sigma_o}} Z_o^{\frac{\sigma_o-1}{\sigma_o}} \right]^{\frac{\sigma_o}{\sigma_o-1}},$$

with inputs  $M_o$  as machines and  $Z_o$  the effective units of labor supplied to an occupation, and with  $\sigma_o$  the elasticity of substitution between machines and labor.<sup>7</sup> Below, we will focus on technical change as shocks to  $\gamma_o$ , which is occupation-specific productivity of machines.

Given this setup, the price of machines  $P$ , and the wage per effective unit of labor in an occupation  $w_o$ ,<sup>8</sup> the marginal cost for a unit of  $Y_o$  is

$$P_o = [\delta_o (P/\gamma_o)^{1-\sigma_o} + (1 - \delta_o) w_o^{1-\sigma_o}]^{\frac{1}{1-\sigma_o}}.$$

Cost minimization then implies that the cost share of labor in production of  $Y_o$  is

$$\omega_o \equiv \frac{w_o Z_o}{P_o Y_o} = (1 - \delta_o) \left( \frac{w_o}{P_o} \right)^{1-\sigma_o}. \quad (1)$$

Demand for occupation  $o$  is given by  $Y_o = \nu_o (P_o/P)^{-\psi} \tilde{Y}$ , with  $P \equiv \left[ \sum_o \nu_o P_o^{1-\psi} \right]^{\frac{1}{1-\psi}}$ . This implies that the expenditure share on goods from occupation  $o$  is:

$$\beta_o \equiv \frac{P_o Y_o}{P \tilde{Y}} = \nu_o \left( \frac{P_o}{P} \right)^{1-\psi}. \quad (2)$$

In turn, this implies that labor demand in occupation  $o$  is given by  $\omega_o \beta_o P \tilde{Y} / w_o$ .

**Labor supply** In our Roy model, workers are heterogeneous in their comparative advantage across  $O$  occupations. The set  $\mathcal{O} \equiv \{1, \dots, O\}$  of occupations is partitioned into  $M$  sets (or nests) with  $\mathcal{O} = \cup_m \mathcal{O}_m$ . Below, this nesting structure will be important for differential reallocation patterns within versus across nests. In addition, workers are split into demographic groups of workers (as in BMV and Hsieh et al. (2019)), where the definition of a group can be based on any pre-determined demographic variable – entailing that workers cannot switch group. This way, we can examine between-group inequality. In practice, we have  $G$  groups of workers, indexed by  $g$ ;  $L_g$  denotes the measure of workers in a group.

---

<sup>7</sup>As in Burstein et al. (2020), the equilibrium conditions we derive can also be obtained from a setup with perfect substitutability between machines and workers at the *task* level, where the substitutability at the *occupation* level arises from the dispersion in comparative advantage between the two inputs.

<sup>8</sup>Factor demand is:  $M_o = \frac{\delta_o \gamma_o^{\sigma_o-1} P^{-\sigma_o} Y_o}{[\delta_o (P/\gamma_o)^{1-\sigma_o} + (1 - \delta_o) w_o^{1-\sigma_o}]^{\frac{\sigma_o}{\sigma_o-1}}}$ ,  $Z_o = \frac{(1 - \delta_o) w_o^{-\sigma_o} Y_o}{[\delta_o (P/\gamma_o)^{1-\sigma_o} + (1 - \delta_o) w_o^{1-\sigma_o}]^{\frac{\sigma_o}{\sigma_o-1}}}$ .

Labor supply follows a Roy–Fréchet structure with nests. Conceptually, workers first choose a nest (e.g., routine vs. non-routine) based on nest-level comparative advantage, and then choose an occupation within that nest based on occupation-level comparative advantage. This structure yields two distinct reallocation elasticities: within-nest mobility ( $\kappa$ ) and across-nest mobility ( $\mu$ ).

Indeed, workers within each group differ in their productivity across occupations. To allow for flexibility in reallocation patterns, each worker's productivity is determined in two steps (following [Galle et al. \(2023\)](#), Section 8).<sup>9</sup> In step 1, workers learn about their nest-specific productivity  $z_{\mathcal{O}_m}$  and decide in which nest to work. Then, in step 2, they learn about their occupation-specific productivity  $z_o$ . Total productivity of a worker in an occupation  $o \in \mathcal{O}_m$  is therefore  $z_o z_{\mathcal{O}_m}$ . Here, all productivity draws  $z_o$  in nest  $\mathcal{O}_m$  are drawn independently from a nest-specific Fréchet distribution with shape parameter  $\kappa_{gm}$  and scale parameters  $\tilde{A}_{go}$ . In turn, the nest-specific productivities are also drawn independently from a Fréchet distribution, but now with shape parameter  $\mu_g$  and scale parameters  $A_{gm}$ .

Workers' earnings are given by  $w_o z_o z_{\mathcal{O}_m}$ : the product of their productivity with the occupation-specific wage  $w_o$ . Workers sort into the occupation where they obtain the highest earnings.

We work backwards and start in step 2. Conditional on sorting in nest  $\mathcal{O}_m$ , the share of workers in group  $g$  that work in occupation  $o \in \mathcal{O}_m$  is

$$\pi_{go|\mathcal{O}_m} = \frac{\tilde{A}_{go} w_o^{\kappa_{gm}}}{\sum_{n \in \mathcal{O}_m} \tilde{A}_{gn} w_n^{\kappa_{gm}}}. \quad (3)$$

This expression clarifies that the Fréchet scale parameters govern the cross-group pattern of comparative advantage, while the shape parameter ( $\kappa_{gm}$ ) becomes the within-nest reallocation elasticity.

Define a group's nest-specific wage index as

$$\tilde{\Phi}_{gm} = \left( \sum_{n \in \mathcal{O}_m} \tilde{A}_{gn} w_n^{\kappa_{gm}} \right)^{1/\kappa_{gm}}. \quad (4)$$

Standard properties of the Fréchet imply that the resulting supply of effective labor units to an occupation is

---

<sup>9</sup>Isomorphic results can be achieved using productivity draws from a nested Fréchet distribution (see e.g., [Zárate \(2022\)](#)). However, our two-step framework allows the cross-nest reallocation to be more elastic than within-nest reallocation, which is not allowed when assuming a nested Fréchet structure, but which is sometimes required by the data (see below).

$$Z_{go} = \bar{z}_{\mathcal{O}_m} \tilde{\zeta}_{gm} \frac{\tilde{\Phi}_{gm}}{w_o} \pi_{go|\mathcal{O}_m} L_g, \quad (5)$$

where  $\bar{z}_{\mathcal{O}_m}$  is the average nest-specific productivity of workers sorting into nest  $m$ , and  $\tilde{\zeta}_{mg} \equiv \Gamma(1 - 1/\kappa_{gm})$ . As a result, for those workers sorting into nest  $m$ , average earnings per worker are constant across occupations, namely

$$\frac{w_o Z_{go}}{\pi_{go|\mathcal{O}_m} L_g} = \bar{z}_{\mathcal{O}_m} \tilde{\zeta}_{gm} \tilde{\Phi}_{gm}.$$

This result also clarifies that the earnings shares across occupations equal employment shares.

In step 1, workers sort across nests based on their expected earnings in each nest:  $z_{\mathcal{O}_m} \tilde{\zeta}_{gm} \tilde{\Phi}_{gm}$ . This results in an analogous share expression as above:

$$\pi_{g\mathcal{O}_m} = \frac{A_{gm} \left( \tilde{\zeta}_{gm} \tilde{\Phi}_{gm} \right)^{\mu_g}}{\sum_{m'} A_{gm'} \left( \tilde{\zeta}_{gm'} \tilde{\Phi}_{gm'} \right)^{\mu_g}}, \quad (6)$$

where  $\pi_{g\mathcal{O}_m}$  is the share of workers in nest  $m$  and the shape parameter  $\mu_g$  has become the cross-nest reallocation elasticity. In turn, the combination of (3) and (6) implies that the unconditional employment share in an occupation is

$$\pi_{go} = \pi_{go|\mathcal{O}_m} \pi_{g\mathcal{O}_m}. \quad (7)$$

Analogous to the within-nest wage index, it is useful to define the cross-nest wage index as

$$\Phi_g \equiv \left( \sum_m A_{gm} \tilde{\zeta}_g^{\mu_g} \tilde{\Phi}_{gm}^{\mu_g} \right)^{1/\mu_g}. \quad (8)$$

Still following an analogous logic as in (5), a group's total earnings is now  $I_g \equiv \sum_o w_o Z_{go} = \zeta_g \Phi_g L_g$ , with  $\zeta_g \equiv \Gamma(1 - 1/\mu_g)$ . Defining a group's average income as  $i_g \equiv I_g / L_g = \zeta_g \Phi_g$ , we see that the wage index  $\Phi_g$  summarizes all the endogenous variation in average income across groups.

**Labor Market Equilibrium** The equilibrium between labor demand and labor supply in each occupation is given by  $\omega_o \beta_o P \tilde{Y} = w_o Z_o$ , with  $Z_o \equiv \sum_g Z_{go}$ . But since  $w_o Z_{go} = \pi_{go} I_g$  and  $P \tilde{Y} = \sum_o \sum_g \pi_{go} I_g / \omega_o$ , we can write this equilibrium as:

$$\omega_o \beta_o \sum_n \sum_g \frac{\pi_{gn} I_g}{\omega_n} = \sum_g \pi_{go} I_g. \quad (9)$$

**Counterfactuals** We are interested in the counterfactual equilibrium after the occurrence of technology shocks, which we solve for using Jones's exact hat algebra, where  $\hat{x} \equiv x'/x$ . We will primarily focus on capital-embodied technology shocks ( $\hat{\gamma}_o$ ), but also allow for labor-eliminating shocks ( $\hat{\delta}_o$ ), or labor supply shocks ( $\hat{A}_{go}$ ) – all specific to an occupation, or nest when considering  $\hat{A}_{gm}$ . The counterfactual labor market equilibrium is given by

$$\omega_o \beta_o \hat{\omega}_o \hat{\beta}_o \sum_n \sum_g \frac{\pi_{gn} \hat{\pi}_{gn} I_g \hat{I}_g}{\omega_n \hat{\omega}_n} = \sum_g \pi_{go} \hat{\pi}_{go} I_g \hat{I}_g. \quad (10)$$

This is a system of  $O$  equations that allow us to solve for  $O$  unknowns: the wage changes  $\hat{w}_o$ . Setting the final good price as the numeraire, all the hat variables are then a function of the data, the productivity shock  $\hat{\gamma}_o$  and the wage changes. We document the expressions for all the hat variables in Appendix Section A.1. For future purposes, note that the shift in occupation-specific labor demand is captured by  $\hat{\omega}_o \hat{\beta}_o$ , since the summation term on the left hand side is a sum across all occupations.

## 2.2 Technical change and the wage distribution

Our focus is on the impact of capital-embodied technical change ( $\hat{\gamma}_o$ ), which exogenously shifts labor demand. To crystallize the impact of these shifts in labor demand on the wage distribution, we consider a simplified setting with a single group, a single nest ( $G = 1$ ), allowing us to drop the subscript  $g$ , and a single value for the input substitution elasticity  $\sigma_o = \sigma$ .

**Proposition 1.** *Assume there is a single group and a single nest, with fixed reallocation elasticity  $\kappa$  and a fixed input substitution elasticity  $\sigma$ . If there are only capital-embodied technology shocks ( $\hat{\gamma}_o$ ), then log relative wage changes satisfy*

$$\ln \left( \frac{\hat{w}_o}{\hat{w}_n} \right) = \left( \frac{\sigma - \psi}{\kappa + \sigma - 1} \right) \ln \left( \frac{\hat{P}_o}{\hat{P}_n} \right). \quad (11)$$

The short, three-step proof appears in Appendix Section A.2.<sup>10</sup>

---

<sup>10</sup>In their analysis of the labor market impact of immigration, Burstein et al. (2020) show that a similar elasticity governs the relative wage responses to immigrant inflows in a locality. Their model is similar to ours, featuring upward-sloping labor supply and a comparable final-demand structure. However, they use immigrant and native workers as inputs in occupational production, while our model uses workers and

**Intuition** Proposition 1 follows from the counterfactual labor market equilibrium in Equation (10). First, on the labor demand side, we obtain

$$\hat{\omega}_o \hat{\beta}_o = \hat{w}_o^{1-\sigma} \hat{P}_o^{\sigma-\psi}, \quad (12)$$

which follows from combining the expressions for the labor share (1) and the expenditure share (2), assuming  $\hat{\delta}_o = \hat{\nu}_o = 1$ . Second, on the labor supply side, we have  $\hat{\pi}_o = \hat{w}_o^\kappa / \hat{I}^\kappa$ . Combining these labor demand and labor supply expressions in (10), while abstracting from aggregate terms, delivers

$$\hat{w}_o^{1-\sigma-\kappa} \propto \hat{P}_o^{\sigma-\psi}, \quad (13)$$

which yields Equation (11) after taking logs and comparing two occupations.<sup>11</sup>

In this expression,  $\hat{P}_o$  incorporates the technology shock. This shock affects labor demand through two mechanisms (see Equation (12)). First, higher  $\gamma_o$  lowers  $P_o$ , the overall cost of production, thereby affecting labor's cost share—an effect governed by the elasticity of *input substitution*  $\sigma_o$ . Second, lower  $P_o$  affects demand for the occupation's output—a *scale effect* governed by the final demand elasticity  $\psi$ .

In combination,  $\sigma - \psi$  determines whether machines and labor are gross substitutes, where the input substitution effect dominates the scale effect ( $\sigma > \psi$ ), or gross complements ( $\sigma < \psi$ ). Indeed, when machines and labor are gross substitutes, the dominant input substitution effect ensures that the decline in  $\hat{P}_o$  associated with a positive productivity shock puts downward pressure on labor demand, leading to a decline in relative wages (and vice versa in case of gross complements).

Notice that Proposition 1 implies that when the input substitution and the scale effect exactly offset each other (i.e., when  $\sigma = \psi$ ), there is no change in relative wages. In this case, the term incorporating the technology shock, namely  $\hat{P}_o^{\sigma-\psi}$  in Equation (12), cancels out. Hence, there is no exogenous shock to labor demand and the labor market remains in the initial equilibrium. The remaining common shift in real wages then arises from the change in purchasing power.

The denominator in (11),  $(\kappa + \sigma - 1)$ , regulates how strongly wages respond to demand shocks, for a given degree of gross substitutability. It arises from the left hand side of (13), which has the endogenous response in equilibrium wages. Here, a higher reallocation elasticity  $\kappa$  entails more elastic occupational labor supply and therefore smaller

---

machines.

<sup>11</sup>Why do the non-hat occupation-specific terms, namely  $\pi_o$  and  $\omega_o \beta_o$ , disappear from the Proposition while being present in Equation (10)? This follows from the initial equilibrium (9), where with a single group we have that  $\frac{\pi_o}{\omega_o \beta_o} = \sum_n \frac{\pi_n}{\omega_n}$ . This aggregate term cancels out when taking the ratio of wage changes relative to another occupation.

equilibrium wage changes. Finally,  $(\sigma - 1)$  mitigates the labor demand shock: wages respond to this shock, leading to input substitution that partially offsets the original shift in labor demand.

**Contribution** The insight in Proposition 1 lies in establishing a concise quantitative relationship between equilibrium wages and marginal costs of production across occupations, highlighting the roles of both labor supply and labor demand.<sup>12</sup> In particular, the proposition integrates the impact of the gross substitutability between workers and machines, on the demand side, and of the reallocation elasticity, on the supply side, in one expression.

The concept that wages are determined by the equilibrium between labor demand and labor supply is at the core of labor economics. It is therefore hardly new. In the context of capital-embodied technical change, the more specific insight that the interplay between the input substitution elasticity and the final demand elasticity shapes wage outcomes, dates back to at least [Hicks \(1932\)](#) and [Robinson \(1933\)](#).<sup>13</sup> On the labor supply side then, [Galle et al. \(2023\)](#), among others, emphasize the quantitative role of the reallocation elasticity in determining wage inequality. Still, the contribution of Proposition 1 lies in combining, crystallizing and quantifying these existing insights from labor demand and labor supply.

The above reasons establish why Proposition 1 is useful. However, since it describes a relationship between endogenous variables, namely wages and costs, it only indirectly speaks to the impact of exogenous technology shocks. The following proposition addresses this limitation.

**Proposition 2.** *Under the same assumptions as Proposition 1, there is a strictly decreasing relationship between occupations' relative costs and their relative machine productivity:*

$$\frac{\hat{P}_o}{\hat{P}_n} = F\left(\frac{\hat{\gamma}_o}{\hat{\gamma}_n}\right), \quad F'(\cdot) < 0.$$

---

<sup>12</sup>[Burstein et al. \(2020\)](#) derive a related result in the context of immigration. While their framework features native and immigrant labor as inputs, rather than machines and workers, it shares an analogous labor supply and demand structure.

<sup>13</sup>Indeed, CJK, referring to both studies, provide the expression for the cross-price elasticity of labor demand: how occupational labor demand responds to changes in the price of capital:

$$-\frac{d \ln Z_o}{d \ln \gamma_o} = \frac{\kappa(\psi - \sigma)(1 - \omega_o)}{\psi + \kappa + (\sigma - \psi)(1 - \omega_o)} \quad (14)$$

While this expression is related, it applies to the cross-price elasticity of *labor demand*, and does not speak directly about our variables of interest, namely equilibrium *relative wages*. Indeed, the expression for the CJK elasticity holds locally, while Proposition 1 is a global result based on exact hat-algebra.

Relatedly, relative equilibrium wages are decreasing in relative machine productivity when the input substitution effect dominates the scale effect, and vice versa:

$$\text{sign} \left( \frac{d(\hat{w}_o/\hat{w}_n)}{d(\hat{\gamma}_o/\hat{\gamma}_n)} \right) = \text{sign} \left( \frac{\psi - \sigma}{\kappa + \sigma - 1} \right).$$

Hence,  $(\psi - \sigma)$  determines the gross complementarity of workers and machines.

The first part of Proposition 2 establishes the monotonic relationship between relative technology shocks and relative costs—the right hand side variable in Proposition 1. The second part then leverages this relationship to demonstrate the monotonic relationship between relative wages and relative technology shocks.

Appendix A.3 has the proof. There, we first establish that, for a single occupation, an increase in machine productivity ( $\hat{\gamma}_o$ ) lowers marginal cost of production ( $\hat{P}_o$ ). This relationship is highly intuitive: since the marginal cost is a CES cost index of the cost of machine input and the cost of labor input, the only way for marginal cost to decline when machine productivity increases, would be when an occupation's wage increases substantially. Importantly though, Proposition 1 has established a monotonic relationship between an occupation's wage and its marginal cost. The proof builds on this link to demonstrate that an occupation's marginal cost has to fall when its machine productivity increases. Once this declining relationship is established for a single occupation, it is straightforward to establish the relationship between the cross-occupation ratios  $\hat{P}_o/\hat{P}_n$  and  $\hat{\gamma}_o/\hat{\gamma}_n$ , while Proposition 1 has established the relationship between relative wages and relative costs ( $\hat{P}_o/\hat{P}_n$ ).

### 3 Model Illustration

We now further explain and illustrate the functioning of the model using a simple baseline case. In this scenario, we focus solely on the distribution of wage changes and therefore consider only a single group to abstract from between-group inequality. For simplicity, we limit the analysis to three occupations, labeled 1, 2, and 3. The first two occupations are placed in one nest, while the third occupation is in a separate nest. The first two occupations each have an employment and expenditure share of 10%. A capital-embodied technology shock is applied to production in occupation 1.

Our illustrative exercise conducts comparative statics by altering one parameter value at a time while keeping the others constant. Unless otherwise indicated, the fixed parameter values are  $\psi = 1.34$ ,  $\mu = 1.4$ ,  $\kappa = 2.9$ , and  $\sigma = 1.95$  for all occupations, which is why the subscript on  $\sigma_o$  is dropped. For now, these values merely serve an illustrative pur-

pose, but they are based on the calibrated values from Section 7.1.<sup>14</sup>

### 3.1 Role of the within- and cross-nest reallocation elasticities $\kappa$ and $\mu$

It is clear that the differences in wage changes are largest when the reallocation elasticities are lowest, i.e. at their lower limit of unity. Indeed, in Figure 1, panel (a), we see that within-nest wage change differences are largest when  $\kappa \rightarrow 1$ , with *negative* wage changes for the directly affected occupation,<sup>15</sup> and that cross-nest wage change differences are largest when  $\mu \rightarrow 1$  (panel b). As the within-nest reallocation elasticity grows, the ripple effects on wage changes within the nest grow. Indeed, for  $\kappa \rightarrow \infty$ , there is perfect convergence of within-nest wage changes due to perfectly elastic reallocation within the nest. An analogous pattern again holds across nests; conditional on  $\kappa \rightarrow \infty$ , there is also perfect wage convergence across all occupations as  $\mu \rightarrow \infty$ .

The Roy-Fréchet model thereby tractably spans the two extremes of perfectly immobile versus perfectly mobile labor across occupations, as well as the associated impact on the wage distribution. When labor supply becomes more elastic, the cross-occupation differences in wage changes disappear, precisely because the ripple effects across occupations are maximized. At the same time, increased discrepancies between the within ( $\kappa$ ) and the cross-nest ( $\mu$ ) reallocation elasticities sharpen the difference in which occupations are affected by the ripple effects. It is therefore critical to obtain estimates on these two elasticities, which is our focus in the empirical section.


---

<sup>14</sup>Specifically,  $\psi = 1.34$  is exactly the calibrated value, while  $\mu = 1.4$  and  $\kappa = 2.9$  are the estimated values for old workers. Finally  $\sigma = 1.95$  is the mean value of the calibrated  $\sigma_o$  across occupations.

<sup>15</sup>The existing literature argues that models with factor-augmenting technological change struggle to generate falling real wages (see e.g. [Acemoglu and Restrepo \(2022\)](#)). Figure 1, panel (a), provides a first counterexample of how our model, with factor-augmenting, capital-embodied technological change, easily generates falling real wages for the directly exposed workers.

Figure 1: Role of the reallocation elasticities

(a) Within-nest wage convergence with  $\kappa$       (b) Cross-nest wage convergence with  $\mu$



Notes: These figures are generated for a machine productivity shock  $\hat{\gamma}_o = 1.1$  in the first occupation (solid blue line), while the other two occupations are not shocked. Occupation 2 (red dashed line) is co-nested with 1, while occupation 3 is in the other nest (orange long-short dashed line). In panel (a),  $\mu = 1.4$  – the default value -, while in panel (b),  $\kappa = 150$ , in order to focus on the cross-nest wage change differences.

### 3.2 Role of input substitution elasticity ( $\sigma$ )

Changes in the input substitution elasticity ( $\sigma$ ) have both standard and less standard effects on the labor market impact of technical change. First, almost by definition, *input substitution* away from labor toward machines increases with  $\sigma$  after a positive productivity shock to machines (see Figure 2, panel a). Indeed, given that  $\hat{\omega}_o = (\hat{w}_o/\hat{P}_o)^{1-\sigma}$ , the cost share of labor falls with  $\sigma$  when the production cost in an occupation ( $\hat{P}_o$ ) declines. Moreover, we notice that  $\hat{w}_o/\hat{P}_o$  actually falls with  $\sigma$  (see Appendix Figure B.1, panel a), entailing that the endogenous wage adjustment dampens the increase in the substitution effect.

Second, there is *demand substitution* toward the shocked occupation (Figure 2, panel b), arising from the decline in the occupation's price (Figure B.1, panel b). Interestingly, an increase in the production-side parameter  $\sigma$  further strengthens this demand-side effect. Since the productivity effect is constant with  $\sigma$ , the output price continues to fall due to the decline in the occupation's wage (Figure 2, panel d), reducing the production cost. Quantitatively, the increase in the demand substitution effect is modest, which is due to the low elasticity of demand ( $\psi = 1.34$ ). Importantly though, the increased demand substitution arises from relaxing restrictive assumptions that are common in the

automation literature.<sup>16</sup>

How does the combination of the above input substitution and demand substitution effects shape overall changes in labor demand? The answer comes from Proposition 1, illustrated in panel (d). When  $\sigma = \psi$ , the two substitution effects exactly cancel out in their impact on relative labor demand. In that case, all occupations experience the same real wage increase, arising from the expansion of the PPF. When  $\sigma < \psi$ , the shocked occupation has an increase in relative labor demand compared to the other occupations (panel c), and a higher relative wage (panel d). Naturally, when  $\sigma > \psi$ , the situation reverses and input substitution dominates demand substitution. In the current case, we then actually obtain negative wage changes for the directly affected occupation once  $\sigma$  reaches 2.<sup>17</sup>

**Ripple effects** As is clear from our discussion on the role of the reallocation elasticities, the nesting structure introduces important ripple effects of productivity shocks in one occupation on wages in other occupations. Indeed, the wage change of the shocked occupation “spills over” to its co-nested occupation, as is clear from panel (d), where the wage change for the co-nested occupation lies in between the wage changes of the shocked occupation and the other occupation.

To understand the underlying mechanisms, first focus on the case where  $\sigma < \psi$  and the relative wage in this co-nested occupation is below the shocked occupation but higher than in other occupations. Why is this? Well, when  $\sigma < \psi$ , the increased demand for workers in the shocked occupation also pushes up wages in the co-nested occupation, as workers in that latter occupation are most prone to moving to the shocked occupation (see Appendix Figure B.1, panel c). When  $\sigma > \psi$ , the pattern is reversed, with a decline in relative demand for workers in the shocked occupation, which leads to increased labor supply to the co-nested occupation. As a result, the wage decline in the shocked occupation ripples over to the co-nested occupation in terms of a reduced relative wage.


**Broader GE effects** For the non-shocked, non-co-nested occupation, the real wage change increases with  $\sigma$ . This is due to the shocked occupation’s price falling more and more

---

<sup>16</sup>First, since the mechanism is driven by an increased occupation-specific wage, it does not occur in any model with homogeneous labor (e.g. Acemoglu and Restrepo, 2018). Second, we document how the demand substitution varies with  $\sigma$ , which is impossible to do in a model featuring only perfect substitutability between machines and the automation-exposed labor type (e.g., Autor et al., 2003).

<sup>17</sup>Whether and for which  $\sigma$  the model generates negative real wage changes for a positive productivity shock depends on how large the expenditure share is on an occupation. In the current setup, we have low expenditure shares (=10%) for the directly affected occupation, which results in a modest aggregate productivity increase. In contrast, when all expenditure shares are set at 1/3, we only obtain negative real wage changes for high  $\sigma$  and small productivity shocks (see Appendix Figures B.3, B.4, and B.5).

Figure 2: Role of  $\sigma$



Notes: These figures are generated for a machine productivity shock  $\hat{\gamma}_o = 1.1$  in the first occupation (solid blue line), while the other two occupations are not shocked. Occupation 2 (red dashed line) is co-nested with 1, while occupation 3 is in the other nest (orange long-short dashed line).

Panel (a) shows the change in the cost share of labor  $\hat{\omega}_o = \left(\hat{w}_o/\hat{P}_o\right)^{1-\sigma}$ , while panel (b) displays the change in the expenditure share on an occupation  $\hat{\beta}_o = \left(\hat{P}_o\right)^{1-\psi}$ . Next, panel (c) shows the change in labor demand as a share of total expenditure ( $\hat{\omega}_o \hat{\beta}_o$ ), while panel (d) depicts the real wage changes.

and its expenditure share increasing, which raises real wages in the other occupations. Since real wages increase with  $\sigma$ , or equivalently, the final good becomes cheaper, production substitutes away from labor in all occupations. We therefore observe a perfectly benign decline in the labor share in the non-shocked occupations, arising from increased real wages (see Figure 2, panel a). This effect is larger when the expenditure share on the shocked occupation is larger (see Appendix Figure B.3). It is therefore important to account for how technical change in one occupation may affect factor prices in other occupations when examining the root causes of declines in the labor share as in Grossman and Oberfield (2022).

**Results when  $\sigma < 1$**  So far, we have focused on an environment where  $\sigma > 1$ . However, we also perform the above comparative statics for the case with  $\sigma < 1$ , and obtain analogous insights (see Appendix Figures B.6 and B.7).

### 3.3 Role of demand substitution elasticity ( $\psi$ )


By definition, an increased demand elasticity ( $\psi$ ) leads to stronger expenditure switches toward the good that experiences the positive productivity shock (Figure 3, panel b). In turn, this increase in demand for output from the occupation then leads to upward pressure on the occupation price (see Appendix Figure B.2, panel b), dampening the increase in the expenditure share.

Since we are assuming that  $\sigma > 1$ , there is always a drop in the level of the labor share, for any value of  $\psi$  (Figure 3, panel a). In addition, this *supply-side* substitution effect deepens when the *demand* elasticity increases. Indeed, the larger is  $\psi$ , the stronger is the demand effect (recall panel b), which puts increasingly upward pressure on the occupation's wage (panel d), thereby further reducing the labor share. While the decline in the level of the labor share reflects a displacement effect from technical change, the further decline with  $\psi$  arises from a fully benign demand-side mechanism.

Examining real wage changes in further detail (panel d), we find that for the shocked occupation, wage changes are low, even negative, when  $\psi$  is close to unity, but become larger and larger when  $\psi$  grows. As before, this is a reflection of the relative size of the labor substitution effect (governed by  $\sigma$ ) and the demand effect (governed by  $\psi$ ) on net labor demand changes (panel c). The higher  $\psi$ , the more important the demand effect becomes relative to the substitution effect.

As was the case for  $\sigma$ , the wage changes from the shocked occupation continue to ripple over to the co-nested occupation (panel d). Indeed, when  $\psi < \sigma$ , the relative wage change is lower for this occupation than for the other non-shocked occupation, and this

Figure 3: Role of  $\psi$



Notes: These figures are generated for a machine productivity shock  $\hat{\gamma}_o = 1.1$  in the first occupation (solid blue line), while the other two occupations are not shocked. Occupation 2 (red dashed line) is co-nested with 1, while occupation 3 is in the other nest (orange long-short dashed line). Panel (a) shows the change in the cost share of labor  $\hat{\omega}_o = (\hat{w}_o/\hat{P}_o)^{1-\sigma}$ , while panel (b) displays the change in the expenditure share on an occupation  $\hat{\beta}_o = (\hat{P}_o)^{1-\psi}$ . Next, panel (c) shows the change in labor demand as a share of total expenditure ( $\hat{\omega}_o \hat{\beta}_o$ ), while panel (d) depicts the real wage changes.

turns around for  $\psi > \sigma$ . As a result of these relative wage changes, relative labor demand is higher for this occupation when  $\psi < \sigma$ , and vice versa (panel c).

## 4 Taking stock and planning ahead

Before moving to the application, we take stock. This paper asks how occupation-specific technical change affects the labor market. To answer this question, we need to under-

stand how technical change affects the interplay of labor demand and labor supply. This paper presents a parsimonious model that organizes this interplay around four elasticities, and the resulting elasticity parameter space spans a wide range of labor market outcomes. Indeed, the setup generalizes canonical skill-biased and polarization models, and streamlines task-based formulations (Acemoglu and Restrepo, 2022). Propositions 1 and 2, characterize how the model’s key mechanisms determine equilibrium wage changes and ripple effects, while the simulations illustrate how final demand, input substitution, and reallocation interact in general equilibrium.

This class of models has been used to study computerization (BMV) or broader capital-embodied technical change (CJK). Our version generalizes BMV on the labor-demand side (allowing richer input substitutability as in CJK and Galle and Lorentzen (2024)) and both BMV and CJK on the labor-supply side (allowing asymmetric reallocation and ripple effects). Next, we apply the framework to LLM adoption, mapping occupation-level exposure into aggregate and distributional impacts across demographic groups and quantifying effects on returns to education, gender earnings gaps, and age-earnings profiles. The result is a single, transparent machinery that links occupation-level AI exposure to macro and distributional outcomes without sacrificing tractability.

## 5 Data

Our empirical analysis uses data from IPUMS USA, focusing on private sector employees aged 24-60, excluding the primary sector. We define demographic groups by gender, five education levels, three age bins, and the nine regional Census divisions, yielding 270 groups.<sup>18</sup>

Following Caunedo et al. (2023), we classify employment into nine 1-digit level occupations (see Table 1). This aggregation reduces sampling noise in group-specific employment shares, while retaining sizable heterogeneity in LLM exposure (see below). We then partition occupations into routine and non-routine nests following Autor and Dorn (2013), a natural nesting structure given the empirical evidence that reallocation and associated wage convergence across these nests is far from perfectly elastic.

**Estimation Data** Our estimation analysis uses the 2000 Census Sample and the 2007 3-year American Community Survey (ACS) from IPUMS. This period is selected because of the presence of substantial shifts in labor demand driven by exogenous trade or tech-

---

<sup>18</sup>The levels of attained education are: (i) lower than high school, (ii) high school, (iii) some college, (iv) college degree, (v) post-graduate degree. The age groups are defined as follows: below 33 years, between 33 and 46 years, and above 46 years. Each age group represents approximately one-third of the sample.

Table 1: Summary statistics on the occupations

	Occupation	Nest	$\beta_o$	mean $\pi_{go}$	SD( $\pi_{go}$ )	$AI_o$
1	Management	NR	0.30	0.23	0.13	0.39
2	Professionals	NR	0.21	0.18	0.17	0.33
3	Sc. technicians	NR	0.07	0.06	0.04	0.42
4	Sales	R	0.11	0.11	0.04	0.30
5	Admin services	R	0.08	0.10	0.08	0.54
6	Low-skill services	NR	0.06	0.11	0.11	0.11
7	Mechanics and transport	NR	0.13	0.15	0.19	0.06
8	Precision prod.	R	0.02	0.02	0.02	0.20
9	Machine operators	R	0.03	0.05	0.05	0.12

Notes: The table lists the nine occupations in our data. The third column explains the nest the occupation belongs to, namely non-routine (NR) or routine (R). The fourth lists the expenditure share on the occupation ( $\beta_o$ ), the fifth the average employment share ( $\pi_{go}$ ), the sixth the standard deviation across groups of the  $\pi_{go}$ , and the final column has the relative exposure to generative artificial intelligence as measured in [Eisfeldt et al. \(2025\)](#).

nology shocks ([Autor and Dorn, 2013](#); [Autor, Dorn, and Hanson, 2013](#); [Fort, Pierce, and Schott, 2018](#); [Acemoglu and Restrepo, 2020](#); [Galle and Lorentzen, 2024](#)). This focus ensures that labor *demand* shocks are the primary source of labor reallocation in this period, which is crucial for reliably estimating the labor *supply* elasticities.

In the estimation, we measure  $I_g$  as the average hourly earnings in a group so that changes in measured income reflect wage adjustments rather than shifts in hours worked or participation.<sup>19</sup> Employment shares ( $\pi_{go|\mathcal{O}_m}$ ,  $\pi_{g\mathcal{O}_m}$ ) are measured as the share of workers in that occupation or nest.

**Simulation data** For the counterfactual analysis on the impact of LLMs, we need data on the baseline values in Equation (10):  $\omega_o$ ,  $\beta_o$ ,  $\pi_{go}$ , and  $I_g$ . To this end, we use the 2022 5-year ACS. There we observe total labor income for each demographic group in each occupation:  $I_{go}$ . An implication of the Roy-Fréchet setup is that employment shares equal earnings shares. In the simulations, we therefore measure employment shares as  $\pi_{go} = I_{go}/I_g$ , since this choice ensures that groups' labor income and employment shares are consistent with the income data. In turn, we measure revenue on an occupation as  $R_o = I_o/\omega_o$ . The expenditure share on that occupation is then measured as  $\beta_o = R_o/\sum_o R_o$ . Finally, we follow BMV and CJK by setting the cost share of labor at  $\omega_o = 0.76$ , based on estimates in [Burstein, Cravino, and Vogel \(2013\)](#). While we do not have a precise measure of  $\omega_o$ , the sensitivity of our counterfactual results to values of  $\omega_o$  is negligible. Since our calibration targets labor productivity changes, the inferred tech-

<sup>19</sup>The latter are not the focus of the baseline model but are captured by the extended model (Section 8).

nology shocks automatically adjust with the value of  $\omega_o$ , ensuring predicted wages and employment are virtually unchanged. Section 7.2 provides further detail.

## 6 Estimation of occupational labor supply

### 6.1 Estimation strategy

#### 6.1.1 Estimation specification

To estimate the within- and cross-nest reallocation elasticities, we start by deriving the estimation specification. First, note from Equations (3), (6), (7), (8) and using that  $\hat{\pi}_{go|\mathcal{O}_m} = \frac{\hat{\pi}_{go}}{\hat{\pi}_{g\mathcal{O}_m}}$ , we obtain:<sup>20</sup>

$$\hat{\Phi}_g = \hat{w}_o \hat{\pi}_{go|\mathcal{O}_m}^{-\frac{1}{\kappa_{gm}}} \hat{\pi}_{g\mathcal{O}_m}^{-\frac{1}{\mu_g}} \hat{A}_{go}^{\frac{1}{\kappa_{gm}}} \hat{A}_{gm}^{\frac{1}{\mu_g}}.$$

Taking logs, imposing a common  $\kappa_{gm} = \kappa$  and  $\mu_g = \mu$ , and using that  $\hat{\Phi}_g = \hat{I}_g$ , we obtain our primary estimation equation:

$$\ln \hat{I}_g = \alpha_o + \beta_1 \ln \hat{\pi}_{go|\mathcal{O}_m} + \beta_2 \ln \hat{\pi}_{g\mathcal{O}_m} + \varepsilon_{go}, \quad (15)$$

where  $\alpha_o \equiv \ln \hat{w}_o$  and  $\varepsilon_{go} \equiv \ln \left( \hat{A}_{go}^{\frac{1}{\kappa_{gm}}} \hat{A}_{gm}^{\frac{1}{\mu_g}} \right)$ . This estimation specification links changes in group income to changes in occupational specialization. The term  $\ln \hat{\pi}_{go|\mathcal{O}_m}$  captures *within-nest* decline in specialization for group  $g$  (holding wages fixed), while  $\ln \hat{\pi}_{g\mathcal{O}_m}$  captures *across-nest* specialization declines. Intuitively, if a particular group exhibits higher occupational expansion than the average group – still holding wages fixed, then this group faces more negative exposure to national-level wage shifts. The group is therefore forced to expand in occupations where it initially had low shares, lowering its occupational specialization. The coefficient estimands map directly to the reallocation elasticities:  $\beta_1 \equiv -1/\kappa$  and  $\beta_2 \equiv -1/\mu$ , measuring how earnings respond to changes in occupational specialization.

Specification (15) holds for all occupations within a given group. But how can all

---

<sup>20</sup>As intermediate step, we obtained from (3), (6), (7), and (8) that

$$\hat{\pi}_{go} = \hat{A}_{go} \hat{A}_{gm} \hat{w}_o^{\kappa_{gm}} \left( \sum_{n \in \mathcal{O}_m} \pi_{gn|\mathcal{O}_m} \hat{A}_{gn} \hat{w}_n^{\kappa_{gm}} \right)^{\frac{\mu_g}{\kappa_{gm}} - 1} \hat{\Phi}_g^{-\mu_g}.$$

Given equations (3) and (7), we can rewrite this expression as

$$\hat{\pi}_{go} = \hat{A}_{go} \hat{A}_{gm} \hat{w}_o^{\kappa_{gm}} \left( \frac{\hat{A}_{go} \hat{w}_o^{\kappa_{gm}}}{\hat{\pi}_{go}} \hat{\pi}_{g\mathcal{O}_m} \right)^{\frac{\mu_g}{\kappa_{gm}} - 1} \hat{\Phi}_g^{-\mu_g}.$$

occupations in a group simultaneously experience higher percentage growth than the average group? The explanation is that groups with above-average occupational expansion experience growth primarily in occupations where they initially had fewer workers, and declines in occupations where they initially had more workers. Indeed, a modest absolute movement of workers from a larger to a smaller occupation results in a small percentage decrease in the large occupation but a large percentage increase in the smaller one. Hence, groups growing their smaller occupations will have higher percentage growth (positive or negative) in all occupations. In economic terms, these groups expand in occupations where they have a comparative *disadvantage* and shrink in the ones where they have a comparative advantage (Galle et al., 2023).

The error term  $\varepsilon_{go}$  consists of occupation-by-group productivity shocks. To the extent that these shocks are correlated with demographic characteristics, we control for them with demographic controls. Specifically, as controls we include fixed effects for a group's defining demographic characteristics: their education level, their gender, their age bin, and their geographic region (Census division). The remaining error term in our estimation then consists of unobservable productivity shocks.

Since the model allows for heterogeneity in the reallocation elasticities across groups, we can estimate our specification separately for different demographic subgroups. We will focus on age groups, as the data shows the most meaningful heterogeneity along this dimension.

### 6.1.2 Shift-share instruments

The regressors in our estimation specifications positively correlate with the error term, since groups' occupational productivity shocks ( $\hat{A}_{go}\hat{A}_{gm}$ ) affect reallocation into that occupation (see Equations (3) and (6)). Hence, an OLS estimation of our specifications will exhibit upward bias in the  $\beta$  coefficient estimates. Since the implied reallocation elasticities are the negative inverse of the coefficient, OLS estimates of these implied elasticities also have an upward bias.

We instrument specialization changes using a shift-share strategy with predetermined group *shares* interacted with national occupation (or nest) reallocation shocks as *shifts*. This strategy builds on our Roy-Fréchet expressions for specialization changes (Equation (18) in Appendix A.1), where the denominator of that ratio is a wage index that summarizes earnings potential for a group. Intuitively, when national labor demand shifts in favor of some occupations, groups that are initially more concentrated in those occupations experience an improvement in their earnings potential and an increase in their occupational specialization.

Concretely, to implement our IV approach, we first construct national reallocation

shocks. Start from  $r_o \equiv w_o Z_o / \sum_n w_n Z_n$  as the national income share of an occupation, such that  $\hat{r}_o$  is a measure of occupational expansion over the estimation window. We analogously define nest-level national income shares  $r_{\mathcal{O}_m} \equiv \sum_{o \in \mathcal{O}_m} r_o$  and their changes  $\hat{r}_{\mathcal{O}_m}$ . Second, we interact these national shocks with predetermined group shares as follows, for each group  $g$ :  $Z_{gm}^{\text{within}} \equiv \sum_o \pi_{go|\mathcal{O}_m} \hat{r}_o$  and  $Z_g^{\text{across}} \equiv \sum_m \pi_{g\mathcal{O}_m} \hat{r}_{\mathcal{O}_m}$ , with employment shares measured in the base period. These objects summarize how exposed a group is to national reallocation pressures within nests and across nests, respectively. More specifically, in the model, specialization responses depend on the within-nest and cross-nest wage indices  $\tilde{\Phi}_{gm}$  and  $\Phi_g$ , which capture the “denominators” of  $\hat{\pi}_{go|\mathcal{O}_m}$  and  $\hat{\pi}_{g\mathcal{O}_m}$ . The shift-share measures  $Z_{gm}^{\text{within}}$  and  $Z_g^{\text{across}}$  are designed to approximate these indices:<sup>21</sup> groups with larger baseline exposure to occupations (or nests) that expand nationally face stronger outside-option shifts, and therefore exhibit predictable changes in specialization. Importantly, we confirm instrument relevance in the first-stage results in Appendix Table C.1.

We are estimating the labor-supply elasticities, and our instrument is therefore valid if it captures labor-demand shocks that are orthogonal to group-by-occupation supply shocks. Since we have a limited number of occupations, we have to follow the “exogenous shares” (Goldsmith-Pinkham, Sorkin, and Swift, 2020) exclusion restriction, which requires group-by-occupation productivity shocks to be unrelated to baseline shares, after conditioning on controls:

$$\mathbb{E}[\varepsilon_{go} | \pi_{go|\mathcal{O}_m}, X_g] = 0, \quad \mathbb{E}[\varepsilon_{go} | \pi_{g\mathcal{O}_m}, X_g] = 0,$$

where  $X_g$  includes fixed effects for the demographic components that define groups (education, gender, age bin, and Census division).

In contrast to reduced-form work disentangling the impact of a specific shock,<sup>22</sup> our instruments are allowed to reflect *any* aggregate or occupation-level labor-demand disturbance (trade, technology, product-demand shifts, etc.) as long as these shocks are not related to group-specific labor-supply/productivity movements. Below, we further corroborate the validity of our exclusion restriction with an analysis of pre-trends, as suggested by Goldsmith-Pinkham et al. (2020). Since the “exogenous shares” assumption ultimately boils down to a continuous difference-in-differences setup, corroborating the parallel-trends assumption is central to our identification strategy.

---

<sup>21</sup>For more detail on related approximation strategies, but in settings without nests, see Galle et al. (2023) and Galle and Lorentzen (2024).

<sup>22</sup>In those setups, it is often more advisable to focus on the “exogenous shifts” identification assumption (Borusyak, Hull, and Jaravel, 2022), since it is common for baseline shares to correlate with multiple national-level demand shocks.

## 6.2 Estimation results

Table 2 has our main estimation results, with columns 1-3 estimating different versions of Equation (15). All specifications include occupation fixed-effects, to absorb the change in wages in that occupation ( $\ln \hat{w}_o$ ). In the first two columns, we use OLS, both without and with controls. In both cases, the implied values for  $\kappa$  and  $\mu$  have the expected sign, but are very high. For instance, the lowest value we find for  $\mu$  is  $\mu = 16.6$ , in the specification with controls. These high values reflect the upward bias in the OLS estimation. Indeed, once we employ IV estimation, the estimated reallocation elasticities drop substantially (column 3). Specifically, we estimate  $\mu = 3.4$  and  $\kappa = 5.3$ , with respective standard errors of 1.13 and 3.1.<sup>23</sup> The first stage is sufficiently strong, with a Kleibergen-Paap F-statistic of 11.1.

Worker mobility is likely to differ across demographic groups. To shed light on this heterogeneity, we proceed with estimating different reallocation elasticities by age group.<sup>24</sup> Specifically, in our empirical setup we have three different age groups (young, middle, and old), and we estimate a separate  $\kappa_g$  and  $\mu_g$  for each of these age groups (see columns 4, 5, and 6 respectively).

Interestingly, both estimated reallocation elasticities decline with age, implying that older workers are less mobile than younger ones. Specifically, we find that  $\mu_g$  declines from a high value of 5.4 for young workers, over 3.0 for middle aged, to the low value of 1.4 for old workers. The within-nest reallocation elasticity  $\kappa_g$  analogously declines from 4.7 for young to 2.9 for old workers, with an intermediate value of 3.6 for the middle group. The standard errors (SEs) are most precise for old workers, especially for the  $\mu_g$  estimate ( $SE = 0.37$ ), but also for their  $\kappa_g$  ( $SE = 1.4$ ). For younger and middle-aged workers, the SEs range from 2.0 to 3.0.

## 6.3 Sensitivity analysis

**Pre-trends** Our exclusion restriction posits that productivity shocks at the group-by-occupation level should not be correlated with a group's employment share. To corroborate this restriction, we analyze if there are pre-trends in group-level income, our dependent variable, correlated with the instruments. Such a correlation would indicate that the instruments are correlated with long-term trends in income, irrespective of period-

---

<sup>23</sup>Standard errors for  $\kappa$  and  $\mu$  are computed using the delta method. Specifically, for  $i = 1, 2$  respectively:  $SE(f(\beta_i)) = SE(\beta_i)|f'(\beta_i)| = SE(\beta_i)/\beta_i^2$ . The non-linear relationship between the estimated value and the structural parameter therefore implies that the confidence intervals on  $\beta_i$  do not simply translate to confidence intervals for the structural parameter.

<sup>24</sup>We also looked into heterogeneity by gender or education level, but that analysis was largely uninformative due to weak first stages or large standard errors of the second stage coefficients.

**Table 2:** Age-specific reallocation elasticities. (Dep. var.:  $\ln \hat{I}_g$ )

	(1) All (OLS)	(2) All (OLS)	(3) All	(4) Young	(5) Middle	(6) Old
$\ln \hat{\pi}_{gO_m}$	-0.015 (0.025)	-0.060*** (0.011)	-0.29*** (0.098)	-0.19** (0.074)	-0.33 (0.22)	-0.72*** (0.18)
$\ln \hat{\pi}_{go O_m}$	-0.019** (0.0080)	-0.0046 (0.0036)	-0.19* (0.11)	-0.21* (0.11)	-0.28 (0.23)	-0.34** (0.16)
Implied $\mu$	66.3 (109.9)	16.6 (3.06)	3.40 (1.13)	5.38 (2.13)	3.02 (1.99)	1.40 (0.36)
Implied $\kappa$	53.6 (23.0)	216.6 (170.8)	5.27 (3.10)	4.74 (2.56)	3.55 (2.96)	2.90 (1.34)
KP F-stat			11.1	5.80	2.22	8.59
Controls	No	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2430	2430	2430	810	810	810

Notes: The table estimates equation (15), namely  $\ln \hat{I}_g = \alpha_o + \beta_1 \ln \hat{\pi}_{go|O_m} + \beta_2 \ln \hat{\pi}_{gO_m} + \varepsilon_{go}$ , where  $\ln \hat{I}_g$  is the log change in average hourly wage in a group, and  $\alpha_o$  is an occupation fixed-effect. Specifications 1 and 2 are estimated with OLS, and the others with IV, with instruments  $\sum_o \pi_{go|O_m} \hat{r}_o$  and  $\sum_m \pi_{gO_m} \hat{r}_{O_m}$ . Specifications 4-6 restrict the sample to young, middle-aged, and old workers respectively. All specifications except the first control for gender FE, education level FE, and Census division FE. Specifications 2 and 3 also control for age-bin FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. Significance levels based on p-values: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors for  $\mu$  and  $\kappa$  based on the delta method:  $SE(f(\beta_i)) = SE(\beta_i) |f'(\beta_i)| = SE(\beta_i) / \beta_i^2$ .

specific shocks to labor demand (see also Goldsmith-Pinkham et al., 2020). We conduct the pre-trend analysis for the period 1990-2000, using the 5% Census samples in IPUMS for those years.

For both instruments, their positive relationship with changes in hourly income disappears during the pre-period, corroborating their validity. For the cross-nest instrument ( $\sum_m \pi_{gO_m} \hat{r}_{O_m}$ ), we actually estimate, if anything, a negative coefficient in the pre-period ( $p = 0.1$ ; see Appendix Table C.2, column 2). For the within-nest instrument ( $\sum_o \pi_{go|O_m} \hat{r}_o$ ), we first control for nest-level expansion in the dependent variable, based on our model's insights from specification (15). After this adjustment, the correlation between the instrument and the outcome is again positive and highly significant in the main period, but close to zero and insignificant in the pre-period (columns 3-4). For completeness, we also provide results using the unadjusted outcome variable; however,

these are less informative in the context of our model (columns 5-6).

**Strength of the first stage** The first stages for the estimations by age group lack some power, as the F-statistics are below 10. In part, this is because we need to employ two instruments that are correlated. We address the strength of the first stage by also estimating an “inverted” specification, now with  $\ln \hat{\pi}_{gO_m}$  as dependent variable and  $\ln \hat{I}_g$  as regressor (see Appendix Table C.3, columns 4-6). Given that we employ an exactly identified IV regression, this inverted specification yields identical estimates of the structural elasticities. Importantly though, two of the three F-statistics in the inverted estimation of the age-specific reallocation elasticities are now above 15, mitigating concerns about potential weak-instrument bias.

**Weighted estimation** In our baseline estimation, all nine occupations receive an equal weight. However, some of these occupations are substantially larger than others in terms of employment shares. For instance, precision production workers make up only 3.4% of employment, while mechanics make up 19%. We therefore also estimate the reallocation elasticities with these national employment shares as weights, and obtain quite similar estimates (see Appendix Table C.4). Indeed, we continue to find that the reallocation elasticities steadily decline with age. In the IV specifications, all point estimates are slightly but not significantly lower than in the baseline. If anything, this implies stronger distributional effects of labor demand shocks.<sup>25</sup>

## 7 Baseline quantification

In this section, we shed light on the quantitative impact of generative AI on the labor market. To this end, we first parametrize our model and calibrate the AI shock. Appendix Table D.1 provides an overview of the inputs in our counterfactual analysis.

### 7.1 Parametrization

Using our above estimates for age-specific reallocation elasticities (Table 2, columns 4-6), we set  $\mu_g$  to 5.4 for young workers, 3 for middle-aged workers, and 1.4 for old workers. For  $\kappa_g$ , the values are 4.8, 3.6, and 2.9, respectively. In turn, our value for the demand elasticity,  $\psi = 1.34$ , is based on the estimate by CJK, and is further supported by BMV’s

---

<sup>25</sup>Unfortunately, the first-stage F-statistics are all below 10 in this weighted estimation. We therefore again also report the results from the “inverted” estimation, with identical elasticity estimates in the IV, where at least the F-stat for old workers is 22 (see Table C.5).

similar finding of  $\psi = 1.78$ . Appendix Table D.1 provides an overview of our calibrated parameters.

### 7.1.1 Calibration of $\sigma_o$

To our knowledge, no direct estimates of the occupation-specific elasticity of substitution between human labor and large language models (LLMs) currently exist. This is not surprising given the very recent rise of these LLMs. We therefore discipline  $\sigma_o$  indirectly, based on two inputs: (i) a conservative range of plausible elasticities  $[\underline{\sigma}, \bar{\sigma}]$ , (ii) a measure of LLMs' complementarity with human labor ( $C_o$ ), based on highly detailed task-level evidence from one major LLM provider: Anthropic (Handa et al., 2025). These two objects allow us to calibrate  $\sigma_o$  as follows:

$$\sigma_o = \underline{\sigma} + (\bar{\sigma} - \underline{\sigma}) \frac{1 - C_o}{2}. \quad (16)$$

This expression yields higher values of  $\sigma_o$  for occupations with high relative substitutability (low  $C_o$ ) and vice versa. We now explain how we discipline the choice of  $\underline{\sigma}$  and  $\bar{\sigma}$ , and how we calculate  $C_o$ .

We set the range of plausible elasticities to  $[\underline{\sigma} = 0.9, \bar{\sigma} = 3]$ . This choice is, first, guided by historical evidence from earlier waves of technical change,<sup>26</sup> which places occupation-specific equipment-labor elasticities in a disciplined but wide range of [0.65-3.27].<sup>27</sup> The range's upper end is informed by ICT estimates (Eden and Gaggl, 2018), of which LLMs are an, albeit recent, subcomponent. In contrast, the very low estimate of 0.65 comes from broader capital-embodied technical change (Caunedo et al., 2023) that is less directly comparable to LLMs, providing a first motivation of our slightly higher lower bound of  $\underline{\sigma} = 0.9$ .

The range  $[\underline{\sigma} = 0.9, \bar{\sigma} = 3]$  is further supported by early reduced-form evidence on the labor-market effects of LLMs when interpreted through our model. The key object governing whether an AI-driven productivity shock raises or lowers occupation-level la-

---

<sup>26</sup>For a external calibration also focusing on occupation-specific input substitution elasticities, see Bárány and Siegel (2021).

<sup>27</sup>Focusing first on ICT-type capital, Eden and Gaggl (2018)'s implied elasticity of substitution between ICT capital and *routine* labor rises from about 2.14 to 3.27 over their sample period, while the corresponding elasticity between ICT and *non-routine* labor is more modest, around 1.43. Second, estimates for a particularly "automation-like" equipment category—industrial robots—also point to substantial heterogeneity: Adachi (2025) estimates occupation-group-specific elasticities between robots and labor ranging from roughly 1.01 in abstract-intensive occupations up to 2.71 for production occupations. Finally, Caunedo et al. (2023) provide direct occupation-level estimates of the elasticity of substitution between capital and labor, ranging from 0.65 for technicians up to 2.18 for administrative services, with other occupations lying in between. Taken together, these findings support allowing  $\sigma_o$  to vary from values modestly below one (complementarity) to values around three (strong substitutability).

bor demand is gross substitutability,  $(\sigma_o - \psi)$ . With our calibrated final-demand elasticity  $\psi = 1.34$ , reduced-form evidence on whether exposed occupations contract or expand is informative about the sign and plausible magnitude of  $\sigma_o - \psi$ .

Two empirical regularities stand out. First, labor demand declines are most clearly detectable in settings and occupations where LLMs appear more substitutive: payroll data, evidence from job postings, and online freelancing markets all show contraction concentrated in automation-heavy segments (Hui, Reshef, and Zhou, 2024; Brynjolfsson et al., 2025; Chen, Srinivasan, and Zakerinia, 2025; Demirci, Hannane, and Zhu, 2025; Liu, Wang, and Yu, 2025; Teutloff, Einsiedler, Kässi, Braesemann, Mishkin, and del Rio-Chanona, 2025).<sup>28</sup> This evidence motivates a high-substitutability tail with  $\sigma_o > \psi$  in automation-heavy occupations; setting  $\bar{\sigma} = 3$  (the midpoint of the upper-tail estimates in Eden and Gaggl (2018) and Adachi (2025)) captures this strongly contractionary region while remaining within the historical upper range.

Second, for more augmentation-oriented occupations, positive employment responses, when present, tend to be muted rather than strongly expansionary (Brynjolfsson et al., 2025).<sup>29</sup> These muted positive effects motivate  $\underline{\sigma} = 0.9$ : it permits modest gross complementarity for the most augmentation-heavy occupations, while remaining close to near-unit lower-tail estimates reported for ICT- and robot-related margins (Eden and Gaggl, 2018; Adachi, 2025).

Because the labor-market effects of LLMs are highly heterogeneous and closely linked to whether LLM use is primarily automative or augmentative, we use this distinction to locate each occupation's  $\sigma_o$  in  $[\underline{\sigma}, \bar{\sigma}]$ . Following Brynjolfsson et al. (2025), we measure automation versus augmentation use based on the Anthropic Economic Index (Handa et al., 2025), which classifies millions of Claude interactions into automative and augmentative behaviors at the task level. Aggregating these task-level interactions to occupation-level shares allows us to construct an occupation-level *complementarity index* ( $C_o \in [-1, 1]$ ):

---

<sup>28</sup>Using U.S. payroll records, Brynjolfsson et al. (2025) document sizeable relative employment declines for early-career workers in AI-exposed occupations and show that these declines are concentrated where AI use is more “automative” than “augmentative.” Studies using job postings as a proxy for labor demand find that vacancies fall more in occupations with higher AI-substitution vulnerability after ChatGPT’s release (Liu et al., 2025), with related evidence of larger posting declines in highly automation-prone jobs (Chen et al., 2025). Consistent with fast adjustment in spot markets, online freelancing studies find post-ChatGPT declines in jobs and earnings in exposed categories (Hui et al., 2024), and large posting/demand reductions in automation-prone clusters or substitutable-skill postings (Demirci et al., 2025; Teutloff et al., 2025).

<sup>29</sup>Brynjolfsson et al. (2025) emphasize that employment gains are muted in applications where AI augments rather than automates. Some vacancy-based studies report relatively stronger posting growth for augmentation-prone or human–AI-collaboration roles (Chen et al., 2025), and online-market “winners” evidence suggests demand growth is concentrated in complementary (rather than substitutable) skill clusters (Teutloff et al., 2025). Giving more weight to the Brynjolfsson et al. (2025) on US payroll data, we view the evidence as supporting limited complementarity for the most augmentation-prone occupations, but not the strong expansion that would be implied by pushing  $\underline{\sigma}$  toward the lowest historical estimates from broadly defined capital-embodied technical change.

more positive values of  $C_o$  indicate greater complementarity, while more negative values indicate greater substitutability.<sup>30</sup> This way, we have all elements to calibrate  $\sigma_o$  as described in Equation (16), while ensuring that the model’s predictions will stay faithful to lessons from early reduced-form evidence.

## 7.2 Calibration of the shock

We calibrate occupation-level machine-productivity shocks in three steps. First, we use the experimental evidence in Dell’Acqua, McFowland, Mollick, Lifshitz-Assaf, Kellogg, Rajendran, Krayer, Cadelon, and Lakhani (2023) as an anchor for the productivity gain from access to GPT-4. Second, we scale this anchor across occupations using the exposure measure in Eisfeldt et al. (2025), which quantifies the share of tasks within each occupation that can be performed using LLMs such as ChatGPT. Third, we choose the vector of  $\hat{\gamma}_o$  so that, occupation by occupation, the model matches the resulting labor-productivity targets.

In step 1, Dell’Acqua et al. (2023) study strategy consultants at Boston Consulting Group and exploit randomized access to GPT-4. They find that access to GPT-4 increased task-completion speed by 25.1%. This estimate is conservative as a measure of labor-productivity gains, since Dell’Acqua et al. (2023) also document improvements in output quality.<sup>31</sup>

---

<sup>30</sup>The Handa et al. (2025) data are publicly available at <https://huggingface.co/datasets/Anthropic/EconomicIndex>, and contain task-level information that is aggregated from four million Claude conversations. Handa et al. (2025) matched each conversation to a unique O\*NET task. They also classified each conversation as *automative* or *augmentative*. Automative behaviors are directives (completely delegating the task to AI) or feedback loops (AI completes the task, guided by user feedback). Augmentative behaviors consist of task iteration (collaborating with the LLM to refine processes), learning (asking the LLM to explain), and validation (asking the LLM to check one’s work). For each task  $t$ , the data contain  $p_t$ , the share of all Claude conversations involving  $t$  (where  $\sum_t p_t = 1$ ). Importantly, we also observe  $a_t$  and  $g_t$ , the share of conversations within task  $t$  that are *automative* or *augmentative*, respectively (where  $a_t + g_t = 1$ ).

We link tasks  $t$  to nine occupations  $o$  in two steps. First, as in Handa et al. (2025), we merge tasks to the Standard Occupational Classification (SOC) system, using O\*NET’s *Task Statements* mapping. Second, we map SOC codes to nine occupation codes, in the same way as Caunedo et al. (2023).

Letting  $T_o$  be the set of tasks mapped to occupation  $o$ , we define usage weights within occupation  $o$  as  $w_{ot} = p_t / \sum_{\tau \in T_o} p_{\tau}$ , so that  $\sum_{t \in T_o} w_{ot} = 1$ . We then aggregate automation and augmentation shares to the occupation level:  $A_o = \sum_{t \in T_o} w_{ot} a_t$ ,  $G_o = \sum_{t \in T_o} w_{ot} g_t$ . This allows us to define a preliminary complementarity index  $\tilde{C}_o \equiv (G_o - A_o) / (G_o + A_o)$ , which lies in  $[-1, 1]$  and increases with the relative importance of augmentative (compared to automative) LLM use in occupation  $o$ . Since our calibration focuses on *relative* differences in complementarity across occupations, any affine transformation of  $\tilde{C}_o$  is innocuous: it can be absorbed by a corresponding renormalization of the  $[\underline{\sigma}, \bar{\sigma}]$  range. We therefore employ a centered, unit-free version  $C_o$  obtained by subtracting the cross-occupation mean of  $\tilde{C}_o$  and dividing its cross-occupation range, so that  $C_o = 0$  corresponds to the average occupation and the extremes correspond to the most augmentation- and automation-heavy occupations in the data.

<sup>31</sup>Noy and Zhang (2023) obtain similar results in a randomized writing task for college-educated professionals: exposure to ChatGPT reduced time to completion by 40% and increased output quality by 18%. Their task is narrower (writing only), whereas Dell’Acqua et al. (2023) study a broader set of tasks relevant

In step 2, we extrapolate the benchmark gain to all occupations using the occupational exposure measure in [Eisfeldt et al. \(2025\)](#). They measure exposure at the 5-digit SOC level as the share of tasks in an occupation that can be performed using LLMs such as ChatGPT.<sup>32</sup> As a validation of these exposure measures, [Eisfeldt et al. \(2025\)](#) show that, in the two weeks after ChatGPT’s release, firms with more exposed occupational compositions earned 44 basis points higher daily stock returns than less exposed firms.

Using “management analysts” from [Dell’Acqua et al. \(2023\)](#) as the benchmark occupation, we compute for each 5-digit occupation its “AI exposure” relative to management analysts. We then aggregate these relative exposure measures to our nine 1-digit occupation categories using employment-share weights. The final column of Table 1 reports the resulting exposure measures: cognitively oriented occupations (e.g., administrative services or scientific technicians) are the most exposed, while occupations with high manual content (e.g., mechanics and transportation or low-skill services) are the least exposed.

Our calibration target for each of the nine occupations is the benchmark productivity gain from [Dell’Acqua et al. \(2023\)](#) (25.1%) multiplied by the occupation’s relative exposure measure. We then calibrate  $\hat{\gamma}_o$  so that the model matches these implied labor-productivity increases.<sup>33</sup> The calibrated shocks are strongly aligned with exposure: they have a correlation of almost 94% with an occupation’s AI exposure (Appendix Figure D.1).

Finally, as noted in Section 5, our counterfactual results are essentially insensitive to  $\omega_o$ , for which we lack direct data. Because the calibration fixes labor-productivity targets, the implied technology shocks adjust to alternative  $\omega_o$  values, leaving predicted wages and employment effectively unchanged.<sup>34</sup>

### 7.3 Results

We present results in two layers. First, we report occupation-level wage and employment responses to the calibrated AI shocks. Second, we aggregate these responses to

---

for an occupation’s output.


<sup>32</sup>For each occupation, [Eisfeldt et al. \(2025\)](#) score tasks by whether ChatGPT can achieve a 50% reduction in completion time, either with or without additional tools built for ChatGPT. The exposure measure is the share of tasks meeting this criterion, with tasks requiring additional tools entering the numerator at 50% weight. This strategy closely follows [Eloundou et al. \(2024\)](#). In a recent contribution, [Smeets, Tian, and Traiberman \(2025\)](#) follow a related calibration strategy to ours.

<sup>33</sup>In the model, changes in labor productivity are  $\widehat{L}Po = \frac{\widehat{R}_o}{\widehat{P}_o \widehat{Z}_o}$ . Noting that  $\omega_o R_o = w_o Z_o$  and that  $\widehat{\omega}_o = \left(\frac{\widehat{w}_o}{\widehat{P}_o}\right)^{1-\sigma}$ , it follows that  $\widehat{L}Po = \widehat{\omega}_o^{\frac{\sigma}{1-\sigma}}$ .

<sup>34</sup>To corroborate this, we ran the counterfactuals for both the baseline model and the extension in Section 8 under two values of  $\omega_o$  (0.24 and 0.76). In the baseline model, results are exactly identical. In the extended model, results are nearly identical: the correlation is nearly perfect and the maximum difference in predicted wage changes is 0.000005.

distributional outcomes across demographic groups, emphasizing returns to education, gender earnings gaps, and age-earnings profiles. Throughout, we highlight the role of reallocation and ripple effects: wage incidence is shaped not only by direct exposure, but also by equilibrium spillovers from reallocation across occupations.

**Figure 4:** Labor market impact of the rise of generative AI



Notes: The figure displays the impact of increases in LLM productivity at the occupation level, for the shock as calibrated in Section 7.2. Panel (a) shows percentage increases in real wages for each occupation, while panel (b) displays percentage changes in employment shares by occupation.

Our calibration of the productivity shocks implies that eight out of nine occupations experience a productivity increase (Appendix Figure D.1, panel a). In turn, this leads to a substantial increase in aggregate productivity, driving up real wages in all occupations (see Figure 4, panel a). Still, differential labor displacement leads to unequal real wage changes. Scientific technicians and administrative services experience the lowest wage growth, since these two occupations have the highest exposure to LLMs (Figure D.1, panel c). Even though scientific technicians' exposure is lower, they still experience lower wage growth than administrative services due to their larger input substitution elasticity ( $\sigma_o = 2.51$  for the technicians vs.  $\sigma_o = 2.1$  for administrative services).

Within each nest, the manual-labor intensive occupations experience higher wage growth, because they have the lowest exposure to advances in LLMs. In contrast, the more cognitively oriented occupations experience lower wage growth. For instance, in the routine nest, machine operators, performing manual work, experience the highest wage growth (4.6%) while wages in administrative services grow by around 3.6%. Analogously, in the non-routine nest, wages for "mechanics and transportation workers" grow by 4.8%, while wages for scientific technicians grow by 3.5%.

Accounting for labor reallocation and associated ripple effects is a strength of our

general-equilibrium framework, and these reallocation effects are sizable. For instance, the employment share of scientific technicians falls by 2.5%, while the mechanics and transportation occupation grows by almost 1.5% (Figure 4, panel b). These sizable reallocation effects explain why the inequality in AI's wage impact is muted. Since we found reallocation across occupations to be fairly elastic, employment is a stronger margin of adjustment than wages.

Capturing these reallocation effects through the model also complements reduced-form approaches that measure the impact of a new technology on the labor market through exposure measures. Our setup corroborates these exposure-measure approaches, since we obtain a negative correlation between AI exposure and the wage changes of 84% (see Appendix Figure D.1, panel c), which simultaneously implies that  $(1 - 0.84^2) = 29.4\%$  of the wage variation remains unaccounted for by these exposure measures. Moreover, a general-equilibrium model is necessary to determine the *level* of the wage gains, and, while doing so, account for the spillover effects of a shock in one occupation on wages in the other occupations through (potentially substantial) sectoral reallocation (Beaudry et al., 2012; Galle and Lorentzen, 2024; Lorentzen, 2024).

## 8 Extended model with intensive and extensive margin

Occupational labor-demand shocks may also affect the unemployment rate and hours worked per worker. To capture these additional adjustment margins, we extend the baseline model by adding a parsimonious search-and-matching block and an intensive-margin hours decision. To preserve tractability in this richer environment, we adopt a nested-Fréchet formulation for productivity draws, which implies within-nest reallocation elasticities exceed the cross-nest elasticity. The demand side remains identical to the baseline model.

### 8.1 Theory

This section adds an intensive and an extensive margin to the labor supply side. The resulting extended model nests our baseline model, except for one, more minor, assumption. Specifically, we drop the two-step productivity draws for nests and occupations, and instead simply assume nested-Fréchet productivity draws. With the nested Fréchet, we preserve tractability throughout, but it comes with the restriction that the within-nest reallocation elasticities are higher than the cross-nest one.<sup>35</sup>

---

<sup>35</sup>Using the nested Fréchet, this model setup remains tractable – building on the results in Kim and Vogel (2021). If we instead were to assume the two-step productivity draws, the uncertainty about occupational productivities in step 2 would limit the tractability on the combination of the sorting pattern and the inten-

The other new aspects of the model are pure extensions. First, we add frictional unemployment by introducing a bare-bones search-and-matching framework, as well as, second, an intensive margin decision arising from a standard trade-off between consumption and disutility from working. Both extensions are modeled as in the highly tractable [Kim and Vogel \(2021\)](#) framework. Specifically, workers apply to the occupation that maximizes their expected utility, knowing their productivity in each occupation. Vacancies in each occupation are posted by employers at a cost  $c_g$  expressed in terms of the final good – our numeraire. Matching between vacancies and applicants is governed by a Cobb-Douglas matching function, with a matching elasticity  $\chi_g$ . After being hired, workers make a decision on how many hours to work, trading off leisure versus consumption. Once these hours worked have been supplied and output is realized, employers and employees engage in Nash bargaining over the match surplus. This results in a share  $\theta_g$  of the surplus going to the hired employee. At the same time, unmatched and therefore unemployed workers receive a real income of zero. Finally, employers post vacancies as long as the expected net benefit is weakly positive, which results in a zero-profit condition for the equilibrium vacancy posting.

We derive the model in detail in Appendix Section [A.4](#), and summarize it here. Employment shares in the formal occupations are just as before, as are the within-nest and cross-nest wage indices ( $\tilde{\Phi}_{mg}$  and  $\Phi_g$  respectively). Given the cross-nest wage index, we show that, given workers' decisions on the intensive margin, average hours per worker in group  $g$  are

$$h_g = \tilde{\eta}_g (\delta_g \theta_g)^{\frac{1}{\xi_g}} \Phi_g^{\frac{1}{\xi_g}},$$

where  $1/\xi_g$  is the intensive margin elasticity,  $\delta_g$  is the demand shifter for consumption, and  $\tilde{\eta}_g \equiv \Gamma(1 - 1/(\xi_g \mu_g))$ . Intuitively, as the wage index  $\Phi_g$  increases, average hours per worker ( $h_g$ ) also increase.

In this tractable setup, the employment rate ( $e_g$ ) also becomes a function of the cross-nest wage index:

$$e_g \propto \Phi_g^{\frac{\chi_g(1+\xi_g)}{(1-\chi_g)\xi_g}},$$

recalling that  $\chi_g$  is the matching or employment elasticity. Total income generated by a group is therefore also a function of the wage index, amplified by the employment and intensive-margin elasticity:

$$I_g \propto \Phi_g^{\frac{1+\xi_g}{(1-\chi_g)\xi_g}} L_g. \tag{17}$$

---

sive margin decision.

**Equilibrium** Compared to the baseline model, we have updated the labor supply side, but total payments to labor in an occupation are still measured by  $\sum_g I_{go} = \sum_g \pi_{go} I_g$ . At the same time, the labor demand side has remained identical. Hence, the expression for the equilibrium system of equations remains as before:

$$\omega_o \beta_o \sum_n \sum_g \frac{\pi_{gn} I_g}{\omega_n} = \sum_g \pi_{go} I_g.$$

We present the associated system of hat equations for the counterfactual equilibrium in Appendix A.4.

## 8.2 Counterfactual analysis

**Parametrization** The parametrization and calibration of the model remain as in the baseline model, except we set  $\mu_g = 4.739$  for young workers, instead of  $\mu_g = 5.38$ , such that all  $\mu_g < \kappa_g$ .<sup>36</sup> In addition, we set the value of the two new parameters, the employment elasticity and the intensive margin elasticity, based on the estimates in [Galle and Lorentzen \(2024\)](#), which are in turn closely in line with standard values in the literature. Specifically, the employment elasticity is set to  $\chi = 0.3$ , which is very close to the estimates in [Shimer \(2005\)](#) and [Barnichon and Figura \(2015\)](#), and we set the intensive margin elasticity to  $\xi = 2.5$ , which is in line with the estimates in [Chetty \(2012\)](#).

**Counterfactual results** The forces leading to labor market clearing in the extended model are closely aligned with the mechanisms in the baseline model. The main change is that certain groups' impact on labor supply is amplified compared to others, due to a relative increase in their employment rate and average hours worked. While this leads to more dispersion in the groups' real income effects, it has a minimal impact on the calibrated shocks and the equilibrium wages. Indeed, we find that both have a correlation of 99.9% with their counterpart in the baseline model (see e.g. Appendix Figures D.2 and D.3).

For the US in the aggregate, real income increases by 8.8% due to the rise of generative AI (see Table 3). These aggregate gains are unequally distributed, with a maximum group-level gain of 9.5% and a minimum gain of 8.1%. Given our parametrization, the change in hourly income accounts for close to 50% of the aggregate gains, the changes in

---

<sup>36</sup>Even though  $\mu_g > \kappa_g$  is not consistent with this extended version of the model, the simulated model still converges. When we instead set  $\mu_g = 5.38$  for young workers, as in the baseline framework, there are only very minor differences with the results presented in this section. This is driven by (i) the economic difference between the two  $\mu_g$  values being minor, (ii) cross-nest reallocation of young workers alone not being critical for the general equilibrium results.

**Table 3:** Labor's adjustment margins after the rise of generative AI

	Aggregate	Mean	SD	Min.	Max.
$\hat{I}_g$	8.75	8.79	0.31	8.13	9.48
$\hat{i}_g$	4.26	4.30	0.15	3.99	4.63
$\hat{h}_g$	1.70	1.70	0.06	1.58	1.83
$\hat{e}_g$	2.56	2.56	0.09	2.37	2.76

Notes: The table shows summary statistics, in percentage terms, for groups' income changes ( $\hat{I}_g$ ), and how they are broken down across hourly income changes ( $\hat{i}_g$ ), hours worked ( $\hat{h}_g$ ), and the employment rate ( $\hat{e}_g$ ).

the employment rate and hours worked for respectively 30% and 20%.

Previous episodes of technical change typically positively affected the returns to education – commonly categorized as skill-biased technical change (Acemoglu, 1998; Krusell et al., 2000; Hémous and Olsen, 2022). For generative AI, this pattern is reversed. Indeed, one of the subgroups gaining the most are high-school drop-outs, who gain 9.3% in the aggregate.<sup>37</sup> This is driven by their disproportionately high employment share of over 55% in the two occupations with the highest wage gains (low-skill services and mechanics & transport), compared to 26% for the overall population. In contrast, workers with the highest level of education gain 0.73 percentage points less in the aggregate.

**Table 4:** Income changes due to generative AI across demographic groups

	Aggregate	Mean	SD	Min.	Max.
All groups	8.75	8.79	0.31	8.13	9.48
Young	8.71	8.77	0.32	8.13	9.48
Middle aged	8.74	8.80	0.31	8.38	9.46
Old	8.77	8.80	0.31	8.37	9.44
Male	8.85	8.95	0.33	8.13	9.48
Female	8.56	8.63	0.18	8.35	9.07
Less than high school	9.28	9.16	0.22	8.84	9.48
High school	9.02	8.91	0.31	8.49	9.26
Some college	8.77	8.71	0.27	8.37	9.08
College	8.62	8.61	0.08	8.38	8.76
Post-graduate degree	8.55	8.54	0.08	8.13	8.64


Notes: The table presents the distribution of income changes ( $\hat{I}_g$ ), split by demographic group, for the model with frictional unemployment and an intensive-margin adjustment.

Again in contrast to previous episodes of technical change, which have been pro-rich, we find that on average, initially poorer groups gain more from generative AI than

<sup>37</sup>These findings echo the earlier insightful analysis of Bloom, Prettner, Saadaoui, and Veruete (2024), with our theoretical setup being more general.

initially richer groups. Given that advances in AI tend to benefit workers in manual-labor intensive occupations most, this may not be surprising. Specifically, we obtain a negative correlation of 27% between groups' log income and their income changes (see Figure 5). Particularly at the bottom of the income distribution, most groups experience above-median income changes. Finally, note that with a correlation of -27%, 93% of the variation in income changes remains unexplained by initial income, indicating the importance of incorporating a detailed analysis of occupation-specific technical change across demographic groups, going beyond simply low- versus high-educated workers.

Figure 5: Income changes due to AI along the income distribution



Notes: The figure shows how the distribution of income changes ( $\hat{I}_g$ ), on the vertical axis, are related to the distribution of initial income ( $\ln I_g$ ), on the horizontal axis. The solid line represents the linear fit between the two variables.

## 9 Conclusion

We have written down a model of the impact of technical change that crystallizes the impact of three elasticities on the wage distribution: the input substitution elasticity, the demand substitution elasticity, and the labor supply (reallocation) elasticities. We leverage our model to examine the general-equilibrium impact of LLMs on labor market outcomes. While the most exposed occupations, such as administrative services, gain less than manual-labor intensive occupations, we find that ripple effects on wages, primarily arising from within-nest reallocation, dampens inequality in wage changes. Lower educated workers experience higher wage growth than higher educated workers.

### **Declaration of generative AI and AI-assisted technologies in the manuscript preparation process**

During the preparation of this work, but after the derivation and illustration of the core model components and after the estimation of the reallocation elasticities, the authors used ChatGPT 5.1 and 5.2 for editing and polishing. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

## References

Acemoglu, D. (1998). Why do new technologies complement skills? Directed technical change and wage inequality. *The Quarterly Journal of Economics* 113(4), 1055–1089.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, Volume 4, pp. 1043–1171. Elsevier.

Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy* 128(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2022). Tasks, automation, and the rise in US wage inequality. *Econometrica* 90(5), 1973–2016.

Adachi, D. (2025). Elasticity of substitution between robots and workers: Theory and evidence from Japanese robot price data. *Journal of Monetary Economics*, 103782.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American economic review* 103(5), 1553–1597.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review* 103(6), 2121–68.

Autor, D. H., F. Levy, and R. J. Murnane (2003, November). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118(4), 1279–1333.

Bárány, Z. L. and C. Siegel (2021). Engines of sectoral labor productivity growth. *Review of Economic Dynamics* 39, 304–343.

Barnichon, R. and A. Figura (2015). Labor market heterogeneity and the aggregate matching function. *American Economic Journal: Macroeconomics* 7(4), 222–49.

Beaudry, P., D. A. Green, and B. Sand (2012). Does industrial composition matter for wages? A test of search and bargaining theory. *Econometrica* 80(3), 1063–1104.

Berlingieri, G., F. Boeri, D. Lashkari, and J. Vogel (2024). Capital-skill complementarity in firms and in the aggregate economy. *NBER working paper*.

Bernon, B. and G. Magerman (2024). A theory of equality and growth. *ECARES working paper*.

Bloom, D. E., K. Prettner, J. Saadaoui, and M. Veruete (2024). Artificial intelligence and the skill premium. *NBER working paper*.

Borusyak, K., P. Hull, and X. Jaravel (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies* 89(1), 181–213.

Brynjolfsson, E., B. Chandar, and R. Chen (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. *Stanford Digital Economy Lab working paper*.

Burstein, A., J. Cravino, and J. Vogel (2013). Importing skill-biased technology. *American Economic Journal: Macroeconomics* 5(2), 32–71.

Burstein, A., G. Hanson, L. Tian, and J. Vogel (2020). Tradability and the labor-market impact of immigration: Theory and evidence from the united states. *Econometrica* 88(3), 1071–1112.

Burstein, A., E. Morales, and J. Vogel (2019). Changes in between-group inequality: computers, occupations, and international trade. *American Economic Journal: Macroeconomics* 11(2), 348–400.

Caselli, F. and A. Manning (2019). Robot arithmetic: new technology and wages. *American Economic Review: Insights* 1(1), 1–12.

Caunedo, J., D. Jaume, and E. Keller (2023). Occupational exposure to capital-embodied technical change. *American Economic Review* 113(6), 1642–1685.

Chen, W. X., S. Srinivasan, and S. Zakerinia (2025). Displacement Or Complementarity? The Labor Market Impact of Generative AI. *Harvard Business School working paper*.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica* 80(3), 969–1018.

Curuk, M. and G. Vannoorenberghe (2017). Inter-sectoral labor reallocation in the short run: The role of occupational similarity. *Journal of International Economics* 108, 20–36.

Dell'Acqua, F., E. McFowland, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Cadelon, and K. R. Lakhani (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (24-013).

Demirci, O., J. Hannane, and X. Zhu (2025). Who is AI replacing? The impact of generative AI on online freelancing platforms. *Management Science*.

Eden, M. and P. Gagl (2018). On the welfare implications of automation. *Review of Economic Dynamics* 29, 15–43.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang (2025). Generative AI and firm values. *Journal of Finance (Forthcoming)*.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384(6702), 1306–1308.

Feigenbaum, J. and D. P. Gross (2024). Answering the call of automation: How the labor market adjusted to mechanizing telephone operation. *The Quarterly Journal of Economics, qjae005*.

Fort, T. C., J. R. Pierce, and P. K. Schott (2018). New perspectives on the decline of US manufacturing employment. *Journal of Economic Perspectives* 32(2), 47–72.

Fortin, N. M. and T. Lemieux (2015). Changes in wage inequality in Canada: An interprovincial perspective. *Canadian Journal of Economics/Revue canadienne d'économique* 48(2), 682–713.

Galle, S. and L. Lorentzen (2024). The unequal effects of trade and automation across local labor markets. *Journal of International Economics* 150, 103912.

Galle, S., A. Rodríguez-Clare, and M. Yi (2023). Slicing the pie: Quantifying the aggregate and distributional effects of trade. *The Review of Economic Studies* 90(1), 331–375.

Giupponi, G., R. Joyce, A. Lindner, T. Waters, T. Wernham, and X. Xu (2024). The employment and distributional impacts of nationwide minimum wage changes. *Journal of Labor Economics* 42(S1), S293–S333.

Gola, P. (2021). Supply and demand in a two-sector matching model. *Journal of Political Economy* 129(3), 940–978.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review* 110(8), 2586–2624.

Grossman, G. M. and E. Oberfield (2022). The elusive explanation for the declining labor share. *Annual Review of Economics* 14(1), 93–124.

Handa, K., A. Tamkin, M. McCain, S. Huang, E. Durmus, S. Heck, J. Mueller, J. Hong, S. Ritchie, T. Belonax, et al. (2025). Which economic tasks are performed with AI? evidence from millions of claudie conversations. *arXiv preprint arXiv:2503.04761*.

Hémous, D. and M. Olsen (2022). The rise of the machines: Automation, horizontal innovation, and income inequality. *American Economic Journal: Macroeconomics* 14(1), 179–223.

Hicks, J. (1932). *The theory of wages*. London:MacMillan.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The allocation of talent and us economic growth. *Econometrica* 87(5), 1439–1474.

Hui, X., O. Reshef, and L. Zhou (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science* 35(6), 1977–1989.

Humlum, A. (2022). Robot adoption and labor market dynamics. *Rockwool Foundation working paper*.

Karabarbounis, L. (2024). Perspectives on the labor share. *Journal of Economic Perspectives* 38(2), 107–136.

Katz, L. F. and K. M. Murphy (1992, February). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics* 107(1), 35–78.

Kim, R. and J. Vogel (2021). Trade shocks and labor market adjustment. *American Economic Review: Insights* 3(1), 115–130.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica* 68(5), 1029–1053.

Lagakos, D. and M. E. Waugh (2013). Selection, agriculture, and cross-country productivity differences. *The American Economic Review* 103(2), 948–980.

Liu, Y., H. Wang, and S. Yu (2025). Labor Demand in the Age of Generative AI: Early Evidence from the US Job Posting Data. Available at SSRN 5504741.

Lorentzen, L. (2024). Domino effects: Understanding sectoral reallocation and its wage implications. *UiO working Paper*.

Noy, S. and W. Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654), 187–192.

Ocampo Díaz, S. (2022). A task-based theory of occupations with multidimensional heterogeneity. *CHCP Working Paper*.

Robinson, J. (1933). *The economics of imperfect competition*. New York: St. Martin's Press.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers* 3(2), 135–146.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95(1), 25–49.

Smeets, V., L. Tian, and S. Traiberman (2025). Field choice, skill specificity, and labor market disruptions.

Teutloff, O., J. Einsiedler, O. Kässi, F. Braesemann, P. Mishkin, and R. M. del Rio-Chanona (2025). Winners and losers of generative AI: Early Evidence of Shifts in Freelancer Demand. *Journal of Economic Behavior & Organization*, 106845.

Tschopp, J. (2015). The wage response to shocks: The role of inter-occupational labour adjustment. *Labour Economics* 37, 28–37.

Willén, A. (2021). Decentralization of wage determination: Evidence from a national teacher reform. *Journal of Public Economics* 198, 104388.

Zárate, R. D. (2022). Spatial misallocation, informality, and transit improvements: Evidence from Mexico City. *working paper*.

## Appendix A Theory

### A.1 System of hat equations

The full system of hat equations is a system of  $O$  equations that allows us to solve for  $O$  unknowns: the wage changes  $\hat{w}_o$ . Setting the final good price as the numeraire, all the hat variables are a function of the data, the productivity shock  $\hat{\gamma}_o$  and the wage changes:

$$\begin{aligned}
\omega_o \beta_o \hat{\omega}_o \hat{\beta}_o \sum_n \sum_g \frac{\pi_{gn} \hat{\pi}_{gn} I_g \hat{I}_g}{\omega_n \hat{\omega}_n} &= \sum_g \pi_{go} \hat{\pi}_{go} I_g \hat{I}_g \\
\hat{P}_o &= \left[ (1 - \omega_o) \hat{\delta}_o \hat{\gamma}_o^{\sigma-1} + \omega_o \frac{(1 - \hat{\delta}_o \delta_o)}{(1 - \delta_o)} \hat{w}_o^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \\
\hat{\beta}_o &= \hat{\nu}_o \hat{P}_o^{1-\psi}, \\
\hat{\omega}_o &= \frac{(1 - \hat{\delta}_o \delta_o)}{(1 - \delta_o)} \left( \frac{\hat{w}_o}{\hat{P}_o} \right)^{1-\sigma}, \\
\hat{\pi}_{go} &= \hat{\pi}_{go|\mathcal{O}_m} \hat{\pi}_{g\mathcal{O}_m} \\
\hat{\pi}_{go|\mathcal{O}_m} &= \frac{\hat{A}_{go} \hat{w}_o^{\kappa_{gm}}}{\hat{\Phi}_{mg}^{\kappa_{gm}}}; \quad \hat{\pi}_{g\mathcal{O}_m} = \frac{\hat{A}_{mg} \hat{\Phi}_{mg}^{\mu_g}}{\hat{\Phi}_g^{\mu_g}}, \\
\hat{\Phi}_{mg} &= \left( \sum_{n \in \mathcal{O}_m} \pi_{gn|\mathcal{O}_m} \hat{A}_{go} \hat{w}_n^{\kappa_{gm}} \right)^{\frac{1}{\kappa_{gm}}} \\
\hat{I}_g &= \hat{\Phi}_g = \left( \sum_m \pi_{g\mathcal{O}_m} \left( \hat{A}_{mg} \sum_{o \in \mathcal{O}_m} \hat{\Phi}_{mg} \right)^{\mu_g} \right)^{1/\mu_g}.
\end{aligned} \tag{18}$$

### A.2 Proof of Proposition 1

When we have a single group, the counterfactual labor market equilibrium (10) can be written as:

$$\hat{w}_o^{1-\sigma} \hat{P}_o^{\sigma-\psi} = \frac{\pi_o \hat{\pi}_o}{\omega_o \beta_o \sum_n \frac{\pi_n \hat{\pi}_n}{\omega_n \hat{\omega}_n}},$$

From the initial equilibrium (9), with a single group we have that  $\frac{\pi_o}{\omega_o \beta_o} = \sum_n \frac{\pi_n}{\omega_n}$ , which we can substitute into the above, and rearrange:

$$\frac{\hat{w}_o^{1-\sigma} \hat{P}_o^{\sigma-\psi}}{\hat{\pi}_o} = \frac{\sum_n \frac{\pi_n}{\omega_n}}{\sum_n \frac{\pi_n \hat{\pi}_n}{\omega_n \hat{\omega}_n}}.$$

Since we have one nest (with reallocation elasticity  $\kappa$ ), we can write this as

$$\hat{w}_o^{1-\sigma-\kappa} = \hat{P}_o^{\psi-\sigma} \frac{\hat{I}^\kappa \sum_n \frac{\pi_n}{\omega_n}}{\sum_n \frac{\pi_n \hat{\pi}_n}{\omega_n \hat{\omega}_n}}.$$

The second term on the right-hand side is constant across all occupations, implying that all relative differences in wage changes are perfectly correlated with changes in occupational prices. Specifically, the log difference between the wage changes in any two occupations  $o$  and  $n$  becomes :

$$\ln \left( \frac{\hat{w}_o}{\hat{w}_n} \right) = \left( \frac{\sigma - \psi}{\kappa + \sigma - 1} \right) \ln \left( \frac{\hat{P}_o}{\hat{P}_n} \right). \quad (19)$$

### A.3 Proof of the Relative-Price Comparative Static

This proof will demonstrate that for any two occupations  $o, n$ , the ratio of their marginal costs ( $\hat{P}_o/\hat{P}_n$ ) is strictly decreasing in their relative machine productivity ( $\hat{\gamma}_o/\hat{\gamma}_n$ ). To show this, we first establish that  $\hat{P}_o$  is decreasing in  $\hat{\gamma}_o$ . Here, it is intuitive that an increase in machine productivity lowers marginal cost. Indeed, since this marginal cost is a CES cost index of the cost of machine inputs and the cost of labor input, the only way for marginal cost to decline when machine productivity increases, would be when an occupation's wage increases substantially. Importantly then, Proposition 1 has established a direct relationship between an occupation's wage and its marginal cost. The proof exploits this link to demonstrate that an occupation's marginal cost has to fall when its machine productivity increases. Once this declining relationship is established for a single occupation, it is straightforward to establish the relationship between the cross-occupation ratios  $\hat{P}_o/\hat{P}_n$  and  $\hat{\gamma}_o/\hat{\gamma}_n$ .

#### 1. Single-occupation expression

Recall that for each occupation  $j \in \{o, n\}$

$$\hat{P}_j^{1-\sigma} = (1 - \omega_j) \hat{\gamma}_j^{\sigma-1} + \omega_j \hat{w}_j^{1-\sigma},$$

while from the proof of Proposition 1, we know that

$$\hat{w}_j = c \hat{P}_j^{\frac{(\sigma-\psi)}{\sigma+\kappa-1}},$$

where  $c \equiv \hat{I}^\kappa \sum_n \frac{\pi_n \hat{\pi}_n}{\omega_n \hat{\omega}_n}$ . Substituting the latter relationship into the former, we obtain

$$\hat{P}_j^{1-\sigma} = (1 - \omega_j) \hat{\gamma}_j^{\sigma-1} + \omega_j c^{1-\sigma} \hat{P}_j^{\frac{(\sigma-\psi)(1-\sigma)}{\sigma+\kappa-1}}. \quad (1)$$

## 2. Own-productivity effect

Totally differentiating Equation (1) with respect to  $\hat{\gamma}_j$  gives

$$(1 - \sigma) \hat{P}_j^{-\sigma} d\hat{P}_j = (1 - \omega_j)(\sigma - 1) \hat{\gamma}_j^{\sigma-2} d\hat{\gamma}_j + \omega_j c^{1-\sigma} \frac{(\sigma - \psi)(1 - \sigma)}{\sigma + \kappa - 1} \hat{P}_j^{\frac{(\sigma-\psi)(1-\sigma)}{\sigma+\kappa-1}-1} d\hat{P}_j.$$

Collecting the  $d\hat{P}_j$  terms and dividing yields,

$$\frac{d\hat{P}_j}{d\hat{\gamma}_j} = -\frac{(1 - \omega_j) \hat{\gamma}_j^{\sigma-2}}{\hat{P}_j^{-\sigma} - \omega_j c^{1-\sigma} \frac{\sigma - \psi}{\sigma + \kappa - 1} \hat{P}_j^{\frac{(\sigma-\psi)(1-\sigma)}{\sigma+\kappa-1}-1}}. \quad (2)$$

We now sign this derivative.

- Sign of the numerator:  $(1 - \omega_j) > 0$  and  $\hat{\gamma}_j^{\sigma-2} > 0$ . Hence, the leading minus sign implies the numerator is negative.
- Sign of the denominator: Factor  $(1/\hat{P}_j)$  from the denominator of Equation (2) and use the single-occupation expression (1) to substitute  $\hat{P}_j^{1-\sigma}$ . Then the denominator becomes:

$$\frac{1}{\hat{P}_j} \left[ (1 - \omega_j) \hat{\gamma}_j^{\sigma-1} + \omega_j c^{1-\sigma} \left( 1 - \frac{\sigma - \psi}{\sigma + \kappa - 1} \right) \hat{P}_j^{\frac{(\sigma-\psi)(1-\sigma)}{\sigma+\kappa-1}} \right].$$

Because  $\kappa > 1$  and  $\psi > 0$ ,  $\left( 1 - \frac{\sigma - \psi}{\sigma + \kappa - 1} \right) = \left( \frac{\psi + \kappa - 1}{\sigma + \kappa - 1} \right) > 0$ . The denominator therefore only consists of positive terms, and is thereby positive as a whole.

We established that in Equation (2), the numerator is negative while the denominator is positive. Therefore,

$$\frac{d\hat{P}_j}{d\hat{\gamma}_j} < 0 \quad \text{for } j = o, n.$$

Intuition: an increase in an occupation's machine productivity monotonically lowers the marginal cost of production in that occupation.

### 3. Relative-price effect

Define these ratios:

$$R_P \equiv \frac{\hat{P}_o}{\hat{P}_n}, \quad R_\gamma \equiv \frac{\hat{\gamma}_o}{\hat{\gamma}_n}.$$

In the previous section, we established that  $\hat{P}_j = f(\hat{\gamma}_j)$  with  $f'(\cdot) < 0$ . We rewrite and define

$$R_P = \frac{f(\hat{\gamma}_o)}{f(\hat{\gamma}_n)} = \frac{f(\hat{\gamma}_n R_\gamma)}{f(\hat{\gamma}_n)} =: g(R_\gamma).$$

Differentiate w.r.t.  $R_\gamma$ :

$$\frac{dg(R_\gamma)}{dR_\gamma} = \frac{\hat{\gamma}_n f'(\hat{\gamma}_n R_\gamma)}{f(\hat{\gamma}_n)} < 0,$$

since  $f'(\cdot) < 0$  and  $f(\hat{\gamma}_n) > 0$ .

### 4. Relation between relative prices and relative productivity shocks

Because  $\frac{dg(R_\gamma)}{dR_\gamma} = \frac{d(\hat{P}_o/\hat{P}_n)}{d(\hat{\gamma}_o/\hat{\gamma}_n)} < 0$ ,

$R_P = \frac{\hat{P}_o}{\hat{P}_n}$  is strictly decreasing in  $R_\gamma = \frac{\hat{\gamma}_o}{\hat{\gamma}_n}$ .

Thus an increase in occupation  $o$ 's machine productivity relative to  $n$ 's ( $R_\gamma \uparrow$ ) lowers its relative price index ( $R_P \downarrow$ ), and vice-versa.

### 5. Relation between relative wages and relative productivity shocks

From Proposition 1, we know that

$$\frac{\hat{w}_o}{\hat{w}_n} = \left( \frac{\hat{P}_o}{\hat{P}_n} \right)^{\frac{\sigma-\psi}{\sigma+\kappa-1}}.$$

Therefore,

$$\frac{d(\hat{w}_o/\hat{w}_n)}{d(\hat{\gamma}_o/\hat{\gamma}_n)} = \left( \frac{\sigma-\psi}{\sigma+\kappa-1} \right) \left( \frac{\hat{P}_o}{\hat{P}_n} \right)^{\frac{\sigma-\psi}{\sigma+\kappa-1}-1} \frac{d(\hat{P}_o/\hat{P}_n)}{d(\hat{\gamma}_o/\hat{\gamma}_n)}$$

Determining the sign of this derivative. We know that:

- $\left( \frac{\hat{P}_o}{\hat{P}_n} \right)^{\frac{\sigma-\psi}{\sigma+\kappa-1}-1} > 0$
- $\frac{d(\hat{P}_o/\hat{P}_n)}{d(\hat{\gamma}_o/\hat{\gamma}_n)} < 0$  (from step 4)

Since one of these two terms is negative while the other is positive, we conclude that:

$$\operatorname{sgn}\left(\frac{d(\hat{w}_o/\hat{w}_n)}{d(\hat{\gamma}_o/\hat{\gamma}_n)}\right) = \operatorname{sgn}\left(\frac{\psi - \sigma}{\kappa + \sigma - 1}\right).$$

Hence, relative wages are increasing in relative machine productivity when the scale effect, regulated by  $\psi$  dominates the input substitution effect, governed by  $\sigma$ . In this case, workers and machines are gross complements. Conversely, workers and machines are gross substitutes when  $\sigma > \psi$  since relative wages decline when relative machine productivity increases.

#### A.4 Extension with intensive and extensive labor-supply margins

Here, we explain how to extend the baseline model by adding an intensive and an extensive margin to the labor supply side, while the labor demand side stays identical its baseline setup. The main paper has the non-technical overview of this model extension.

**Intensive margin** In their intensive margin decision, workers optimize their expected utility which, conditional on working in occupation  $o$ , is given by

$$U(C, H; g) = \delta_g C - \frac{H^{1+\xi_g}}{1 + \xi_g},$$

where consumption  $C$  of the final good is funded by a worker's earnings,  $H$  are the number of hours they decide to work, and we restrict  $\xi_g > 0$ . Conditional on being hired in occupation  $o$ , a worker has real earnings  $w_o z_o$  per hour worked. (Note that  $w_o$  are now real instead of nominal wages, since the final good price  $P$  is the numeraire.) From maximizing utility, we can then find that the optimal number of hours worked by this worker is:

$$H = (\delta_g \theta_g w_o z_o)^{1/\xi_g}.$$

This choice on hours results in a real income of

$$\delta_g^{\frac{1}{\xi_g}} (\theta_g w_o z_o)^{\frac{1+\xi_g}{\xi_g}}. \quad (20)$$

Recall that unemployed workers have zero income. We then guess and verify below that the employment probability is constant across sectors:  $e_{go} = e_g$ . Expected utility in occupation  $o$  is then

$$\frac{\xi_g}{1 + \xi_g} e_g (\delta_g \theta_g w_o z_o)^{\frac{1+\xi_g}{\xi_g}}. \quad (21)$$

**Sorting across occupations** The nested-Fréchet distribution from which workers draw their productivities is given by the following cumulative distribution of  $\mathbf{z} \equiv \{z_1, \dots, z_O, z_{HP}\}$ :

$$F_g(\mathbf{z}) = \exp \left( - \sum_m \left( \sum_{o \in \mathcal{O}_m} A_{go} z_o^{-\kappa_{gm}} \right)^{\mu_g / \kappa_{gm}} \right),$$

with the restrictive assumption that  $\kappa_{gm} > \mu_g$ . We assume that the home production occupation is in its own, separate nest. Workers sort into occupations knowing their productivity in each occupation, anticipating their intensive margin decision, and the probability of unemployment in each sector. Given that their utility in an occupation is monotonically increasing in  $w_o z_o$ , we can formalize the sorting pattern across occupations as follows. Let  $\mathbf{w} \equiv \{w_1, \dots, w_O, w_{HP}\}$  and define

$$\Omega_o(\mathbf{w}) \equiv \{z \text{ s.t. } w_o z_o \geq w_k z_k \text{ for all } k\},$$

which implies that a worker with productivity vector  $\mathbf{z}$  will work in occupation  $o$  iff  $\mathbf{z} \in \Omega_o(\mathbf{w})$ . We assume that  $w_{HP}$  is exogenous, since income from home production is not determined by the market. Standard properties of the Fréchet then imply that the within-nest ( $\pi_{go|\mathcal{O}_m}$ ) and cross-nest ( $\pi_{g\mathcal{O}_m}$ ) employment shares for market occupations are as in the baseline:

$$\pi_{go|\mathcal{O}_m} = \frac{A_{go} w_o^{\kappa_{gm}}}{\tilde{\Phi}_{mg}^{\kappa_{gm}}}, \quad (22)$$

$$\pi_{g\mathcal{O}_m} = \frac{(\sum_{n \in \mathcal{O}_m} A_n w_n^{\kappa_{gm}})^{\mu_g / \kappa_{gm}}}{\Phi_g^{\mu_g}}, \quad (23)$$

with the within-nest wage index  $\tilde{\Phi}_{mg}$  defined as in Equation (4). The cross-nest wage index now takes the form:

$$\Phi_g \equiv \left( A_{gHP}^{1/\kappa_{HP}} w_{HP}^{\mu_g} + \sum_{m \neq HP} \tilde{\Phi}_{mg}^{\mu_g} \right)^{1/\mu_g}. \quad (24)$$

**Income and welfare** From Equation 13 in Kim and Vogel (2021), we know that

$$E[z_o^b | o] = \Gamma \left( 1 - \frac{b}{\mu_g} \right) \left( \frac{\Phi_g}{w_o} \right)^b.$$

Given the sorting pattern into sectors and given Equation (20), average real income for workers that applied to a sector is therefore

$$\frac{\theta_g w_o Z_{go}}{\pi_{go} L_g} = \eta_g e_g \delta_g^{\frac{1}{\xi_g}} (\theta_g \Phi_g)^{\frac{1+\xi_g}{\xi_g}},$$

which is constant across occupations (a special implication of the Frechet). Here,  $\eta_g \equiv \Gamma\left(1 - \frac{1+\xi_g}{\xi_g \mu_g}\right)$ . Total income generated by a group therefore becomes

$$I_g \equiv \sum_o w_o Z_{go} = \eta_g e_g (\delta_g \theta_g)^{\frac{1}{\xi_g}} \Phi_g^{\frac{1+\xi_g}{\xi_g}} L_g. \quad (25)$$

Given workers' decision on the intensive margin, and given the implied value for  $E[z_s^{1/\xi_g} | o]$ , average hours per worker in group  $g$  is

$$h_g = \tilde{\eta}_g (\delta_g \theta_g)^{\frac{1}{\xi_g}} \Phi_g^{\frac{1}{\xi_g}}.$$

**Matching** To model frictional unemployment, we also follow the parsimonious search-and-matching framework from [Kim and Vogel \(2021\)](#). Employers post vacancies to hire workers and a Cobb-Douglas hiring function matches vacancies to applicants. In equilibrium, the cost of posting a vacancy equals its expected benefit (zero-profit condition). Once a worker is hired, the employer and the employee bargain over the surplus of the vacancy, and a share  $\theta_g$  of the match surplus ends up going to the worker. At the time of hiring, the vacancy cost is sunk, so the expected surplus is equal to the average revenue realized by an applicant  $I_{go} \equiv w_o Z_{go}$ .

Therefore, in group  $g$  and occupation  $o$ , the expected match surplus per worker that an employer receives is  $(1 - \theta_g) I_{go}$ . The cost of posting a vacancy is assumed to be  $c_g$ . Hence, the zero-profit condition (ZPC) for firms entails

$$c_g V_{go} = (1 - \theta_g) I_{go}.$$

The Cobb-Douglas hiring function that matches occupation-specific applicants ( $\pi_{go} L_g$ ) and vacancies ( $V_{go}$ ) is given by

$$H_{go} = A_g^M V_{go}^{\chi_g} (\pi_{go} L_g)^{1-\chi_g}.$$

Market tightness is the ratio of vacancies over applicants:  $\psi_{ogs} \equiv V_{ogs}/(\pi_{ogs} L_{og})$ . This is a useful definition since the employment rate ( $e_{go} \equiv H_{go}/\pi_{go} L_g$ ) is then a function of labor market tightness and the employment elasticity  $\chi_g$ :

$$e_{go} = A_g^M \psi_{go}^{\chi_g}.$$

Starting from the ZPC, and substituting in the expressions for  $I_{go}$  and  $e_{go}$ , we obtain that

$$c_g V_{go} = (1 - \theta_g) A_g^M \psi_{go}^{\chi_g} \eta_g (\delta_g \theta_g)^{\frac{1}{\xi_g}} \Phi_g^{\frac{1+\xi_g}{\xi_g}} \pi_{go} L_g.$$

Solving for labor market tightness:

$$\psi_{go}^{1-\chi_g} = A_g^M \frac{(1 - \theta_g)}{c_g} \eta_g (\delta_g \theta_g)^{\frac{1}{\xi_g}} \Phi_g^{\frac{1+\xi_g}{\xi_g}}.$$

Importantly, the left-hand side is identical for all occupations  $o$ , which implies that labor market tightness and the associated employment rate is indeed constant across  $o$ :  $e_{go} = e_g$ , verifying our earlier conjecture. This employment rate is then equal to:

$$e_g = (A_g^M)^{\frac{1}{1-\chi_g}} \left( \frac{(1 - \theta_g) \eta_g}{c_g} \right)^{\frac{\chi_g}{1-\chi_g}} (\delta_g \theta_g)^{\frac{\xi_g}{\xi_g(1-\chi_g)}} \Phi_g^{\frac{\chi_g(1+\xi_g)}{(1-\chi_g)\xi_g}}. \quad (26)$$

Intuitively, a shock that increases  $\Phi_g$ , increases the return to posting a vacancy in a group and thereby pushes the employment rate up.

**Equilibrium** Compared to the baseline model, we have updated the labor supply side. However, total payments to labor in an occupation are still measured by  $\sum_g I_{go} = \sum_g \pi_{go} I_g$ . At the same time, the labor demand side has remained identical. Hence, the expression for the equilibrium in all market occupations, so excluding home production where the wage is exogenous, remains as before:

$$\omega_o \beta_o \sum_n \sum_g \frac{\pi_{gn} I_g}{\omega_n} = \sum_g \pi_{go} I_g,$$

and likewise for the counterfactual equilibrium:

$$ELD_o = \omega_o \beta_o \hat{\omega}_o \hat{\beta}_o \sum_n \sum_g \frac{\pi_{gn} \hat{\pi}_{gn} I_g \hat{I}_g}{\omega_n \hat{\omega}_n} - \sum_g \pi_{go} \hat{\pi}_{go} I_g \hat{I}_g. \quad (27)$$

Importantly though, we will need to take into account the intensive and extensive margin in solving for  $\hat{I}_g$ . First note from Equation (26) that

$$\hat{e}_g = \hat{\Phi}_g^{\frac{\chi_g(1+\xi_g)}{(1-\chi_g)\xi_g}},$$

while from (25) the expression for real income changes is updated to

$$\hat{I}_g = \hat{e}_g \hat{\Phi}_g^{\frac{1+\xi_g}{\xi_g}} = \hat{\Phi}_g^{\frac{1+\xi_g}{(1-\chi_g)\xi_g}}.$$


The other hat equations remain similar to before, but now setting the final good price as the numeraire.

$$\begin{aligned}
\hat{\Phi}_g &= \left( \pi_{gHP} + \sum_{m \neq HP} \pi_{g\mathcal{O}_m} \left( \sum_{o \in \mathcal{O}_m} \pi_{go|\mathcal{O}_m} \hat{A}_{go} \hat{w}_o^{\kappa_{gm}} \right)^{\frac{\mu_g}{\kappa_{gm}}} \right)^{1/\mu_g}, \\
\hat{\pi}_{go} &= \hat{\pi}_{go|\mathcal{O}_m} \hat{\pi}_{g\mathcal{O}_m} \\
\hat{\pi}_{go|\mathcal{O}_m} &= \frac{\hat{A}_{go} \hat{w}_o^{\kappa_{gm}}}{\sum_{n \in \mathcal{O}_m} \pi_{gn|\mathcal{O}_m} \hat{A}_{gn} \hat{w}_n^{\kappa_{gm}}}, \\
\hat{\pi}_{g\mathcal{O}_m} &= \frac{\left( \sum_{o \in \mathcal{O}_m} \pi_{go|\mathcal{O}_m} \hat{A}_{go} \hat{w}_o^{\kappa_{gm}} \right)^{\frac{\mu_g}{\kappa_{gm}}}}{\hat{\Phi}_g^{\mu_g}}, \\
\hat{w}_o &= \frac{\hat{w}_o^{1-\sigma}}{(1 - \omega_o) \hat{\gamma}_o + \omega_o \hat{w}_o^{1-\sigma}}, \\
\hat{\beta}_o &= \hat{\nu}_o \hat{P}_o^{1-\psi}, \\
\hat{P}_o &= \left[ (1 - \omega_o) \hat{\delta}_o \hat{\gamma}_o^{\sigma-1} + \omega_o \frac{(1 - \hat{\delta}_o \delta_o)}{(1 - \delta_o)} \hat{w}_o^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.
\end{aligned} \tag{28}$$

Finally,  $\hat{P}$  does not need to be solved for anymore, as it is the numeraire.

## Appendix B Supplementary model illustration


Figure B.1: Role of  $\sigma$




Notes: These figures are generated for a machine productivity shock  $\hat{\gamma}_o = 1.1$  in the first occupation, while the other two occupations are not shocked. Panel (a) shows  $\hat{w}_o / \hat{P}_o$ , which determines the change in the cost share of labor  $\hat{\omega}_o = (\hat{w}_o / \hat{P}_o)^{1-\sigma}$  (Figure 2, panel a). Next, panel (b) shows the change in the real price of an occupation's output ( $\hat{P}_o$ ), which drives the change in the expenditure share on an occupation  $\hat{\beta}_o = (\hat{P}_o)^{1-\psi}$  (Figure 2, panel b). Finally, panel (c) shows the change in employment shares for the occupations.

Figure B.2: Role of  $\psi$


(a) Wage vs. occupation price change



(b) Relative occupation price change




(c) Employment share changes



Notes: These figures are generated for one  $\hat{\gamma}_o = 1.1$  in the first occupation, while the other occupations are not shocked. Panel (a) shows  $\hat{w}_o / \hat{P}_o$ , which determines the change in the cost share of labor  $\hat{\omega}_o = (\hat{w}_o / \hat{P}_o)^{1-\sigma}$  (Figure 3, panel a). Next, panel (b) shows the change in the real price of an occupation's output ( $\hat{P}_o$ ), which drives the change in the expenditure share on an occupation  $\hat{\beta}_o = (\hat{P}_o)^{1-\psi}$  (Figure 3, panel b). Finally, panel (c) shows the change in employment shares for the occupations.

Figure B.3: Role of  $\sigma$  with equal expenditure shares ( $\beta_o = 1/3$ )




Notes: In contrast to their counterpart in the main text, the model for these figures has equal expenditure shares across occupations ( $\beta_o = 1/3$ ). These figures are generated for a machine productivity shock  $\gamma_o = 1.1$  in the first occupation, while the other two occupations are not shocked. Panel (a) shows the change in the cost share of labor  $\hat{\omega}_o = \left(\frac{\hat{w}_o}{\hat{P}_o}\right)^{1-\sigma}$ , while panel (b) displays the change in the expenditure share on an occupation  $\hat{\beta}_o = \left(\hat{P}_o\right)^{1-\psi}$ . Next, panel (c) shows the change in labor demand as a share of total expenditure ( $\hat{\omega}_o \hat{\beta}_o$ ), while panel (d) depicts the real wage changes.

**Figure B.4:** Role of  $\sigma$  with equal expenditure shares ( $\beta_o = 1/3$ )




Notes: In contrast to their counterpart in the main text, the model for these figures has equal expenditure shares across occupations ( $\beta_o = 1/3$ ). The figures are generated for a machine productivity shock  $\gamma_o = 1.1$  in the first occupation, while the other occupations are not shocked. Panel (a) shows  $\hat{w}_o / \hat{P}_o$ , which determines the change in the cost share of labor  $\hat{\omega}_o = (\hat{w}_o / \hat{P}_o)^{1-\sigma}$  (Figure B.3, panel a). Next, panel (b) shows the change in the real price of an occupation's output ( $\hat{P}_o$ ), which drives the change in the expenditure share on an occupation  $\hat{\beta}_o = (\hat{P}_o)^{1-\psi}$  (Figure B.3, panel b). Finally, panel (c) shows the change in employment shares for the occupations.

Figure B.5: Negative wage changes for large  $\sigma$  with equal expenditure shares ( $\beta_o = 1/3$ )




The figure shows wage changes for a larger range on  $\sigma$  and for a smaller productivity shock ( $\hat{\gamma}_o = 1.01$ ), but with equal expenditure shares ( $\beta_o = 1/3$ ) across occupations. For the productivity shock of  $\hat{\gamma}_o = 1.1$ , the wage changes are always positive when all  $\beta_o = 1/3$ , whereas here they are negative for large  $\sigma$ . As documented in Figure 2, note that for  $\hat{\gamma}_o = 1.1$ , the wage changes also become negative for large  $\sigma$  when the expenditure share on the shocked occupation is lower.

Figure B.6: Role of  $\sigma$  when  $\sigma < 1$



Notes: These figures are generated for a machine productivity shock  $\hat{\gamma}_o = 1.1$  in the first occupation, while the other two occupations are not shocked. Panel (a) shows the change in the cost share of labor  $\hat{\omega}_o = \left(\frac{\hat{w}_o}{\hat{P}_o}\right)^{1-\sigma}$ , while panel (b) displays the change in the expenditure share on an occupation  $\hat{\beta}_o = \left(\hat{P}_o\right)^{1-\psi}$ . Next, panel (c) shows the change in labor demand as a share of total expenditure ( $\hat{\omega}_o \hat{\beta}_o$ ), while panel (d) depicts the real wage changes.

**Figure B.7: Role of  $\sigma$  when  $\sigma < 1$**



Notes: The figures are generated for a machine productivity shock  $\hat{\gamma}_o = 1.1$  in the first occupation, while the other occupations are not shocked. Panel (a) shows  $\hat{w}_o/\hat{P}_o$ , which determines the change in the cost share of labor  $\hat{\omega}_o = (\hat{w}_o/\hat{P}_o)^{1-\sigma}$  (Figure B.6, panel a). Next, panel (b) shows the change in the real price of an occupation's output  $(\hat{P}_o)$ , which drives the change in the expenditure share on an occupation  $\hat{\beta}_o = (\hat{P}_o)^{1-\psi}$  (Figure B.6, panel b). Finally, panel (c) shows the change in employment shares for the occupations.

## Appendix C Supplementary estimation results

Table C.1: First-stage results

	(1)	(2)	(3)	(4)
$\ln \hat{\pi}_{go O_m}$				
$\sum_{o \in O_m} \pi_{go O_m} * \hat{r}_o$	-0.75*** (0.080)	-0.99*** (0.12)		
$\sum_{O_m} \pi_{gO_m} * \hat{r}_{O_m}$			-0.40*** (0.14)	-1.70*** (0.32)
T-statistic	9.46	8.02	2.76	5.34
Occupation FE	Yes	Yes	n/a	n/a
Nest FE	n/a	n/a	Yes	Yes
Controls	No	Yes	No	Yes
Estimation	2430	2430	540	540

Notes: The table documents the relevance of the instruments, on the right-hand side in the estimation, for changes in the endogenous regressors in estimation equation (15). The controls include gender FE, education level FE, age-bin FE and Census division FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. P-values: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

**Table C.2:** Reduced forms for actual and pre- period

	$\ln \hat{I}_g$ (1)	$\ln \hat{I}_g + \frac{1}{\hat{\mu}} \ln \hat{\pi}_{gO_m}$ (3)	$\ln \hat{I}_g$ (5)			
$\sum_{O_m} \pi_{gO_m} * \hat{r}_{O_m}$	0.95** (0.43)	-1.11* (0.67)				
$\sum_{o \in O_m} \pi_{go O_m} * \hat{r}_o$		0.19*** (0.042)	0.048 (0.061)	-0.20*** (0.032)	0.080 (0.055)	
Time Period	2000-07	1990-2000	2000-07	1990-2000	2000-07	1990-2000
T-statistic	2.22	1.66	4.46	0.78	6.26	1.46
Nest FE	n/a	n/a	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	270	270	540	540	540	540

Notes: Specification 1,2,5,6 regress changes in group-level hourly income ( $\ln \hat{I}_g$ ) on the instruments used in the IV estimation. Since  $\sum_{o \in O_m} \pi_{go|O_m} * \hat{r}_o$  is a nest-level instrument, we adjust the dependent variable in specifications 3-4, based on the relation in our estimation equation (15). In columns 5-6, we also report the unadjusted income variable, though this variable is less relevant from the perspective of the model. All specifications control for gender FE, education level FE, Census division FE and age-bin FE. In specifications 3-6, standard errors are clustered at the demographic group level. P-values: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

**Table C.3:** Inverted estimation of reallocation elasticities (Dep. var.:  $\ln \hat{\pi}_{gO_m}$ )

	(1) All (OLS)	(2) All (OLS)	(3) All	(4) Young	(5) Middle	(6) Old
$\ln \hat{I}_g$	-0.018 (0.029)	-0.22*** (0.040)	-3.40*** (1.13)	-5.38** (2.13)	-3.02 (1.99)	-1.40*** (0.36)
$\ln \hat{\pi}_{go O_m}$	-0.052*** (0.011)	-0.054*** (0.011)	-0.64*** (0.19)	-1.14*** (0.30)	-0.85*** (0.22)	-0.48*** (0.15)
Implied $\mu$	0.018 (0.029)	0.22 (0.040)	3.40 (1.13)	5.38 (2.13)	3.02 (1.99)	1.40 (0.36)
Implied $\kappa$	0.34 (0.58)	4.10 (0.98)	5.27 (3.10)	4.74 (2.56)	3.55 (2.96)	2.90 (1.34)
KP F-stat			4.54	15.5	1.76	31.9
Controls	No	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2430	2430	2430	810	810	810

Notes: The table estimates the following equation  $\ln \hat{\pi}_{gO_m} = \alpha_o + \beta_1 \ln \hat{I}_g + \beta_2 \ln \hat{\pi}_{go|O_m} + \varepsilon_{go}$ . Specifications 1 and 2 are estimated with OLS and the others with IV, with instruments  $\sum_o \pi_{go|O_m} \hat{r}_o$  and  $\sum_m \pi_{gO_m} \hat{r}_{O_m}$ . Specifications 4-6 restrict the sample to young, middle-aged, and old workers respectively. All specifications except the first control for gender FE, education level FE, and Census division FE. Specifications 2 and 3 also control for age-bin FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. Significance levels based on p-values: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors for  $\kappa$  based on the delta method:  $SE(f(\beta_2)) = SE(\beta_2) |f'(\beta_2)| = SE(\beta_2) / \beta_2^2$ .

**Table C.4:** Weighted estimation of age-specific reallocation elasticities. (Dep. var.:  $\ln \hat{I}_g$ )

	(1) All (OLS)	(2) All (OLS)	(3) All	(4) Young	(5) Middle	(6) Old
$\ln \hat{\pi}_{gO_m}$	-0.046* (0.024)	-0.052*** (0.010)	-0.34*** (0.13)	-0.26** (0.12)	-0.38 (0.29)	-0.75*** (0.21)
$\ln \hat{\pi}_{go O_m}$	-0.022** (0.0095)	-0.00058 (0.0043)	-0.21 (0.14)	-0.25* (0.14)	-0.32 (0.30)	-0.44** (0.21)
Implied $\mu$	21.8 (11.6)	19.1 (3.77)	2.98 (1.13)	3.80 (1.72)	2.64 (1.99)	1.33 (0.37)
Implied $\kappa$	45.6 (19.8)	1735.3 (13048.2)	4.71 (3.01)	4.07 (2.27)	3.13 (2.95)	2.29 (1.11)
KP F-stat			7.93	3.62	1.48	6.33
Controls	No	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2430	2430	2430	810	810	810

Notes: The table estimates equation (15), namely  $\ln \hat{I}_g = \alpha_o + \beta_1 \ln \hat{\pi}_{go|O_m} + \beta_2 \ln \hat{\pi}_{gO_m} + \varepsilon_{go}$ , with national employment shares of the occupations as estimation weights.  $\ln \hat{I}_g$  is the log change in average hourly wage in a group, and  $\alpha_o$  is an occupation fixed-effect. Specifications 1 and 2 are estimated with OLS, and the others with IV, with instruments  $\sum_o \pi_{go|O_m} \hat{r}_o$  and  $\sum_m \pi_{gO_m} \hat{r}_{O_m}$ . Specifications 4-6 restrict the sample to young, middle-aged, and old workers respectively. All specifications except the first control for gender FE, education level FE, and Census division FE. Specifications 2 and 3 also control for age-bin FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. Significance levels based on p-values: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors for  $\mu$  and  $\kappa$  based on the delta method:  $SE(f(\beta_i)) = SE(\beta_i) |f'(\beta_i)| = SE(\beta_i) / \beta_i^2$ .

Table C.5: Weighted, inverted estimation of reallocation elasticities (Dep. var.:  $\ln \hat{\pi}_{gO_m}$ )


	(1) All (OLS)	(2) All (OLS)	(3) All	(4) Young	(5) Middle	(6) Old
$\ln \hat{I}_g$	-0.051** (0.026)	-0.18*** (0.037)	-2.98*** (1.13)	-3.80** (1.72)	-2.64 (1.99)	-1.33*** (0.37)
$\ln \hat{\pi}_{go O_m}$	-0.064*** (0.012)	-0.062*** (0.012)	-0.63*** (0.19)	-0.93*** (0.20)	-0.84*** (0.22)	-0.58*** (0.19)
Implied $\mu$	0.051 (0.026)	0.18 (0.037)	2.98 (1.13)	3.80 (1.72)	2.64 (1.99)	1.33 (0.37)
Implied $\kappa$	0.80 (0.46)	2.95 (0.73)	4.71 (3.01)	4.07 (2.27)	3.13 (2.95)	2.29 (1.11)
KP F-stat			3.82	3.86	1.67	22.7
Controls	No	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2430	2430	2430	810	810	810

Notes: The table estimates the following equation:  $\ln \hat{\pi}_{gO_m} = \alpha_o + \beta_1 \ln \hat{I}_g + \beta_2 \ln \hat{\pi}_{go|O_m} + \varepsilon_{go}$ , with national employment shares of the occupations as estimation weights. Specifications 1 and 2 are estimated with OLS and the others with IV, with instruments  $\sum_o \pi_{go|O_m} \hat{r}_o$  and  $\sum_m \pi_{gO_m} \hat{r}_{O_m}$ . Specifications 4-6 restrict the sample to young, middle-aged, and old workers respectively. All specifications except the first control for gender FE, education level FE, and Census division FE. Specifications 2 and 3 also control for age-bin FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. Significance levels based on p-values: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . Standard errors for  $\kappa$  based on the delta method:  $SE(f(\beta_2)) = SE(\beta_2) |f'(\beta_2)| = SE(\beta_2) / \beta_2^2$ .


## Appendix D Appendix for the counterfactuals

Figure D.1: AI exposure, calibrated shocks, and wage changes for baseline model


(a) AI exposure and calibrated shocks



(b) Calibrated shocks and wage changes



(c) AI exposure and wage changes



Notes: the Figure shows the relationships between an occupation's relative exposure to AI ( $AI_o$ ), measured as in Eisfeldt et al. (2025), the calibrated machine-productivity shocks for an occupation ( $\hat{\gamma}_o$ ) and the resulting counterfactual wage changes ( $\hat{w}_o$ ) for the baseline quantification.

Table D.1: Parametrization, data inputs, and shock calibration for the counterfactuals

Panel A. Elasticities				
	Value	Description	Measurement	Notes
$\kappa_g$	Young: 4.8 Middle: 3.6 Old: 2.9	Within-nest labor re-allocation elasticity by age group $g$	Own estimation	See Table 2, cols. 4–6
$\mu_g$	Young: 5.4 Middle: 3.0 Old: 1.4	Across-nest labor re-allocation elasticity by age group $g$	Own estimation	See Table 2, cols. 4–6
$\psi$	1.34	Demand substitution across occupations' output	Estimate by Caunedo et al. (2023).	BMV find similar value of $\psi = 1.78$ .
$\sigma_o$	$\in [1.65, 2.66]$	Input substitution between machines and labor	$\sigma_o = \underline{\sigma} + (\bar{\sigma} - \underline{\sigma}) \frac{1 - C_o}{2}$ , $C_o$ index corroborated by Brynjolfsson et al. (2025).	

Panel B. Variables				
	Value	Description	Measurement	Notes
$I_g$	—	Income for demographic group $g$	Total labor income	Source: 2022 5-year ACS in IPUMS
$\pi_{go}$	see Table 1	Employment share of group $g$ in occupation $o$	Model implies: $\pi_{go} = I_{go}/I_g$ , with $I_{go}$ a group's earnings in $o$	Source: 2022 5-year ACS in IPUMS
$\beta_o$	see Table 1	Baseline expenditure share on good $o$	Measured as $R_o/\sum_o R_o$ , where $R_o = I_o/\omega_o$	
$\omega_o$	0.76	Baseline cost share of labor in occupation $o$	Estimate by Burstein et al. (2013)	Results are fully insensitive to instead setting $\omega_o = 0.24$ : a consequence of our calibration strategy


  

Panel C. Shock calibration				
	Value	Description	Target	Notes
$\hat{\gamma}$	See Figure D.1	Machine-augmenting productivity shock	Extrapolated productivity increases	labor- (LP) Observed LP increase from Dell'Acqua et al. (2023)


Notes: Section 7.1.1 describes our parametrization (Panel A) in further detail. Section 5 presents the data sources for the variables (Panel B), and Section 7.2 describes the calibration of the AI shock (Panel C).

Figure D.2: AI exposure, calibrated shocks, and wage changes for extended model


(a) AI exposure and calibrated shocks



(b) Calibrated shocks and wage changes




(c) AI exposure and wage changes



Notes: the Figure shows the relationships between an occupation's relative exposure to AI ( $AI_o$ ), measured as in [Eisfeldt et al. \(2025\)](#), the calibrated machine-productivity shocks for an occupation ( $\gamma_o$ ) and the resulting counterfactual wage changes ( $\hat{w}_o$ ) for the extended model with unemployment and an intensive margin.

Figure D.3: Real wage changes in the extended model



Notes: the figure presents changes in real wages for the model with involuntary unemployment and an intensive margin adjustment.