

The Labor Market Impact of Technical Change: Inspecting the Mechanisms

Fenella Carpena

OsloMet

Simon Galle

BI Norwegian Business School

October 2024*

Abstract

Motivated by the rise of artificial intelligence (A.I.), we write down 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, allowing for asymmetric reallocation and associated ripple effects on wages across occupations. With our estimates in hand, we shed light on the general-equilibrium impact of occupation-specific advances in A.I.: wages in administrative services grow the least, ripple effects on less exposed occupations are substantial, A.I. modestly lowers the returns to education and, on average, benefits the poor more.

JEL: J23, J24, E24, E25

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

*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 Pawel Gola, Daniel Gross, Jochen Güntner, Jonas Hamang, Sverre Jensen, Attila Lindner, Linnea Lorentzen, Espen Moen, Peter Morrow, Plamen Nenov, Andrés Rodríguez-Clare, Ovidijus Stauskas, Alisa Tazhitdinova, Jon Vogel, and seminar audiences at Aarhus, BI, CORA at Goethe University, JKU Linz, OsloMet, and the University of Oslo for insightful comments and valuable feedback. We used ChatGPT-4o for light editing. All errors are our own. Corresponding author: Simon Galle, simon.galle@bi.no.

1 Introduction

As rapid advancements in artificial intelligence (A.I.) unfold and start to affect the economy (Korinek, 2023; Humlum and Vestergaard, 2024), we are once again revisiting the familiar debate between techno-optimists, who emphasize the productivity gains from new technologies, and techno-pessimists, who warn of job displacement and potential income loss for workers. In reality, both effects are typically in place. For instance, already in the early industrial revolution, the Luddites were particularly exposed to mechanisation of their occupation: they faced significant job displacement, while others benefited from cheaper textiles. Today, many occupations are (and will be) affected by the rise of generative A.I., but some, such as Hollywood screenwriters, are notably more concerned about their job security. These examples underscore the need to recognize the heterogeneous impact of technical change across occupations (see e.g. Feigenbaum and Gross, 2024), to precisely pinpoint the incidence of gains and losses from a new technology, in the absence of redistribution.¹

In this paper, we present a simple, tractable, yet general model of the general-equilibrium (GE) impact of occupation-specific technical change on labor market outcomes. The model crystallizes how three fundamental forces shape the way gains and losses from capital-embodied technical change are distributed.

The first key force is the degree of substitutability between labor and machines. While robots are often modeled as perfect substitutes for labor, this assumption is less fitting for production processes involving generative A.I., where human input often remains critical. We therefore focus on imperfect substitutability, potentially even complementarity, between machines and labor in production.

The second fundamental force is the final demand response when an occupation's output becomes cheaper to produce. Indeed, even if the new technology is a substitute

¹The question of optimal policy design to ensure that new technologies are Pareto improving is intriguing and important, but outside the scope of this paper. Korinek and Stiglitz (2018) provide an insightful discussion on this topic.

for human labor, overall labor in that occupation might increase if demand for the occupation's output picks up sufficiently.

Prioritizing clarity and tractability, we include the above two forces in our framework using constant-elasticity setups throughout. A production function with a constant elasticity of substitution (CES) regulates input substitution between labor and machines: this governs the response of the cost share of labor. Another CES function controls final demand across the output of all occupations: this determines the response of the final expenditure share on an occupation. Overall *labor demand* for a specific occupation is then obtained from the product of the two aforementioned shares, and this product will be shocked by technical change. Naturally, our CES framework allows for examining the impact of varying degrees of substitutability or complementarity, as reflected in the corresponding elasticities of substitution.

Turning to *labor supply*, we also model labor reallocation across occupations, our third fundamental force, using a constant-elasticity framework. The elasticity of labor supply is micro-founded by the degree of worker heterogeneity, which governs the scope for labor mobility across occupations. If workers are equally talented in all occupations, then can seamlessly reallocate after displacement, and the distributive implications of technical change on labor income approach zero. Conversely, if worker reallocation is costly, this impact is greater. To capture imperfect labor mobility, we follow the seminal approach by [Autor, Levy, and Murnane \(2003\)](#) – henceforth ALM, and introduce a [Roy \(1951\)](#) model, which is the efficient benchmark for modeling reallocation.²

We generalize the ALM setup by allowing for reallocation across any discrete number of occupations. To maintain tractability, we assume Fréchet heterogeneity in worker productivity across occupations, which results in a constant elasticity of occupational labor supply ([Lagakos and Waugh, 2013](#)). Our setup then spans the two extremes of

²Recent examples include [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 [Bernon and Magerman \(2024\)](#), who study the role of the input-output network in affecting income inequality after sector-specific productivity shocks.

perfectly elastic and perfectly inelastic labor supply to occupations. This way, we highlight the critical role of the labor supply elasticity, since distributional effects are maximal when labor supply is inelastic and non-existent when workers are perfectly mobile (Galle, Rodríguez-Clare, and Yi, 2023).

Our model also tractably incorporates spillover (or ripple) effects from labor reallocation (Acemoglu and Restrepo, 2022; Lorentzen, 2024) where displaced workers are more likely to move to certain occupations over others. For instance, displaced screenwriters might move into other non-routine cognitive occupations rather than to jobs with routine manual labor, thereby unevenly increasing competition and driving down wages in the former occupations. To capture these ripple effects, we introduce a nesting structure in our Roy-Fréchet framework, allowing for different reallocation elasticities within and across subsets of occupations.

In summary, four fundamental forces determine how occupation-specific technical change affects the wage distribution. On the labor demand side, the first is the degree of substitutability between labor and machines in production (σ), and the second is the final demand elasticity (ψ). These two elasticities together shape labor demand. On the labor supply side, the third force is the reallocation elasticity within occupation nests (κ) and the fourth is the reallocation elasticity across them (μ).

This simple but flexible framework of occupation-specific labor demand and labor supply will crystallize how capital-embodied technical change influences the distribution of relative wages. To see this, we focus on the simplest case involving only one reallocation elasticity (κ). In that setting, we find that the wage change for an occupation directly exposed to capital-embodied technical change, relative to any other occupation, is explicitly linked to changes in relative cost changes, with an elasticity equal to $(\sigma - \psi)/(\sigma + \kappa - 1)$.³ Specifically, when one occupation becomes relatively cheaper due to

³In their analysis of the labor market impact of immigration, Burstein, Hanson, Tian, and Vogel (2020) show that an isomorphic 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 machines.

capital-embodied technical progress, the relative wage will also decrease, governed by the difference between the input and the demand substitution elasticity ($\sigma - \psi$). Here, $\psi - \sigma$ measures the gross complementarity of workers and machines, and therefore determines the overall *demand* effect of improvements in machine productivity. On the labor *supply* side, the reallocation elasticity governs the magnitude of the distributional effects: these effects are strongest when $\kappa \rightarrow 1$ (specific factors), while there is no wage inequality when $\kappa \rightarrow \infty$ (workers are perfectly mobile).

We further examine the model's mechanisms through simulation analysis, revealing an interesting interplay between input choice and final demand. Occupation-specific wages create feedback mechanisms between changes in expenditure shares in the final good and cost shares in production. First, for a given technological shock and demand elasticity, the expenditure share on an occupation grows as the *input* substitution elasticity (σ) rises. Second, after the same machine productivity shock, a higher elasticity of final *demand* elasticity (ψ) further deepens the decline in the labor cost share on the input side. This decline in the labor share stems from increased occupation-specific wages, and is therefore benign for workers.

After illustrating the model, we proceed by estimating the reallocation elasticities with a transparent IV approach. In particular, we estimate the within-nest (κ) and across-nest (μ) elasticities 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, worker mobility tends to fall with age, although even old workers remain more mobile within than across occupation nests.

With our estimates in hand, we shed light on the GE impact of occupation-specific advances in A.I. To calibrate this capital-embodied productivity shock, we combine the findings on chatGPT's productivity effect from [Dell'Acqua, McFowland, Mollick, Lifshitz-Assaf, Kellogg, Rajendran, Krayer, Candelon, and Lakhani \(2023\)](#) with the occupational exposure measure to large language models from [Eisfeldt, Schubert, and Zhang](#)

(2023). The occupation most exposed to A.I. is administrative services, while mechanics and transportation workers are least exposed. These two occupations therefore experience the lowest and highest wage growth respectively. Nevertheless, all occupations experience positive real wage growth due to the productivity gains from A.I.

Although A.I. indeed affects non-routine occupations more than previous waves of technical change (Autor, 2024), we find that A.I. still leads to higher average wage growth in the non-routine nest than in the routine nest. This is in large part because the most exposed occupation, administrative services, belongs to the routine nest, and the relative wage decline in that occupation ripples over to the entire nest. The presence of these ripple effects highlights the relevance of appropriately modeling labor reallocation.

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 7.1%. 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 A.I. show a strongly pro-poor distribution across demographic groups. Specifically, income changes exhibit a negative correlation of 41% with initial income levels. Moreover, A.I. has a modest negative effect on the education premium: compared to workers with a college degree, high-school dropouts experience a larger increase in earnings by 1.1 percentage points. These findings indicate that A.I. has 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. ALM; 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, the Roy-Fréchet setup tractably generalizes the classic Ricardo-Roy realloca-

tion 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.⁴ 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, Morales, and Vogel (2019), and Caunedo, Jaume, and Keller (2023) – henceforth BMV and CJK respectively – as well as on the framework in Galle and Lorentzen (2024).⁵ 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).⁶

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 A.I. 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

⁴This 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.

⁵BMV 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 (2024) 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.

⁶Another 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.

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).⁷ In the context of 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.⁸

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 combining the output from occupations according to a CES production function (as in BMV):

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

⁷Relatedly, 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).

⁸Specifically, 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.

where Y_o is the output produced by occupation o , ν_o is a demand shifter for this occupation, and $\psi > 0$ is the elasticity of substitution across occupations' output. This final good is used for consumption and for all machines in the economy $M = \sum_o M_o$. The final good has a price P .

Each occupation produces an output

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

with inputs M_o as machines and Z_o the effective units of labor supplied to an occupation, and with σ the elasticity of substitution between machines and labor.^{9 10} Below, we will focus on technical change as shocks to γ_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 ,¹¹ the marginal cost for a unit of Y_o is

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

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

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

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:

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

⁹Importantly, we can allow σ_o to vary across occupations, but we drop the subscript o to reduce the notational burden.

¹⁰As 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.

¹¹Factor demand is: $M_o = \frac{\delta_o \gamma_o^{\sigma-1} P^{-\sigma} Y_o}{[\delta_o (P/\gamma_o)^{1-\sigma} + (1-\delta_o) w_o^{1-\sigma}]^{\frac{\sigma}{\sigma-1}}}$, $Z_o = \frac{(1-\delta_o) w_o^{-\sigma} Y_o}{[\delta_o (P/\gamma_o)^{1-\sigma} + (1-\delta_o) w_o^{1-\sigma}]^{\frac{\sigma}{\sigma-1}}}$.

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, Hurst, Jones, and Klenow (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.

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).¹² 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 κ_{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 μ_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

¹²Isomorphic 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. However, this more flexible pattern of reallocation elasticities is empirically relevant, as we see below.

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

This expression clarifies that the Fréchet scale parameters govern the cross-group pattern of comparative advantage, while the shape parameter (κ_{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

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

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: $\bar{z}_{\mathcal{O}_m} \tilde{\zeta}_{gm} \tilde{\Phi}_{gm}$. 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 μ_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 Φ_g summarizes all the endogenous variation in average income across groups.

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:

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

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 ($\hat{A}_{go}, \hat{A}_{gm}$) – all specific to an occupation (or nest). 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.

Labor demand and relative wages Our primary focus is on the impact of capital-embodied technical change ($\hat{\gamma}_o$), which exogenously shifts labor demand. Here, we take a closer look at the components of shifts in labor demand. Combining the expressions for the labor share (1) and the expenditure share (2), and setting $\hat{\delta}_o = \hat{\nu}_o = 1$, the change in labor demand can be written as

$$\hat{\omega}_o \hat{\beta}_o = \hat{w}_o^{1-\sigma} \hat{P}_o^{\sigma-1} \hat{P}_o^{1-\psi}. \quad (11)$$

The right hand side of this equation consists of three terms. First, the wage component embodies the standard wage substitution effect as function of σ . The middle term reflects input substitution: it determines the degree of substitution toward (or away from) labor when the cost of production in an occupation changes holding the cost of labor (incorporated in the first term) constant. Finally, $\hat{P}_o^{1-\psi}$ reflects a scale effect arising from demand substitution: ψ regulates how expenditure shares ($\hat{\beta}_o$) change when prices (\hat{P}_o) shift.

When the input substitution and the scale effect exactly cancel out (i.e. when $\sigma = \psi$), there will be a constant wage change across all occupations: the only change in real wages arises from the common increase in purchasing power. In contrast, when $\sigma > \psi$, the input substitution effect dominates the scale effect, and the decline in \hat{P}_o associated with a positive productivity shock puts downward pressure on labor demand, leading to a decline in real wages (vice versa when $\sigma < \psi$).

To crystallize the impact of shifts in labor demand on the wage distribution, we focus on a special case with a single nest and a single group ($G = 1$).

Proposition 1. *Assume there is a single group and a single nest, with reallocation elasticity κ . If*

there are only capital-embodied technology shocks ($\hat{\gamma}_o$), then log relative wage changes are given by

$$\ln \left(\frac{\hat{w}_o}{\hat{w}_n} \right) = \left(\frac{\sigma - \psi}{\sigma + \kappa - 1} \right) \ln \left(\frac{\hat{P}_o}{\hat{P}_n} \right). \quad (12)$$

Appendix Section A.2 has the short, three-step proof.^{13 14} Intuitively, $\sigma - \psi$ governs the net impact of input substitution versus scale effects, as explained above. Second, in the denominator, $\sigma - 1$ regulates the partly offsetting change in labor demand resulting from the wage change and associated input substitution (see Equation (11)). Finally, the reallocation elasticity κ governs the size of the relative wage changes: the more elastic occupational labor supply, the smaller the wage changes needed to reach the new labor market equilibrium.

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 param-

¹³Burstein et al. (2020) obtain a closely related result in the context of immigration. Their model has an analogous structure for labor supply and final demand, but models production within occupation with inputs of native and immigrant workers, instead of with machines and workers in general.

¹⁴CJK, referring to Hicks (1932) and Robinson (1933), provide the expression for the cross-price elasticity of labor demand: how occupational labor demand responds to changes in the price of capital:

$$-\frac{d \ln Z_o}{d \ln \gamma_o} = \frac{\kappa(\psi - \sigma)(1 - \omega_o)}{\psi + \kappa + (\sigma - \psi)(1 - \omega_o)} \quad (13)$$

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

eter values are $\sigma = 2.18$, $\psi = 1.34$, $\mu = 1.4$, and $\kappa = 2.9$. For now, these values merely serve an illustrative purpose, but in Section 6.1 we justify the relevance of these values.


3.1 Role of the within- and cross-nest reallocation elasticities κ and μ

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, 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 (κ) and the cross-nest (μ) 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.

Figure 1: Role of the reallocation elasticities

(a) Within-nest wage convergence with κ



(b) Cross-nest wage convergence with μ



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. 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.

3.1.1 Role of input substitution elasticity (σ)

Changes in the input substitution elasticity (σ) 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 σ after a positive productivity shock to machines (see Figure 2, panel a). Indeed, given that $\hat{\omega}_o = \left(\hat{w}_o/\hat{P}_o\right)^{1-\sigma}$, the cost share of labor falls with σ when the production cost in an occupation (\hat{P}_o) declines. Moreover, we notice that \hat{w}_o/\hat{P}_o actually falls with σ (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 (panel b), arising from the decline in the occupation’s price (Appendix Figure 2, panel b). Perhaps surprisingly, an increase in the production-side parameter σ further strengthens this demand-side effect. Since the productivity effect is constant with σ , the output price continues to fall due to the decline in the occupation’s wage (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. 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 σ , 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](#)).

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 σ reaches 2.¹⁵

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

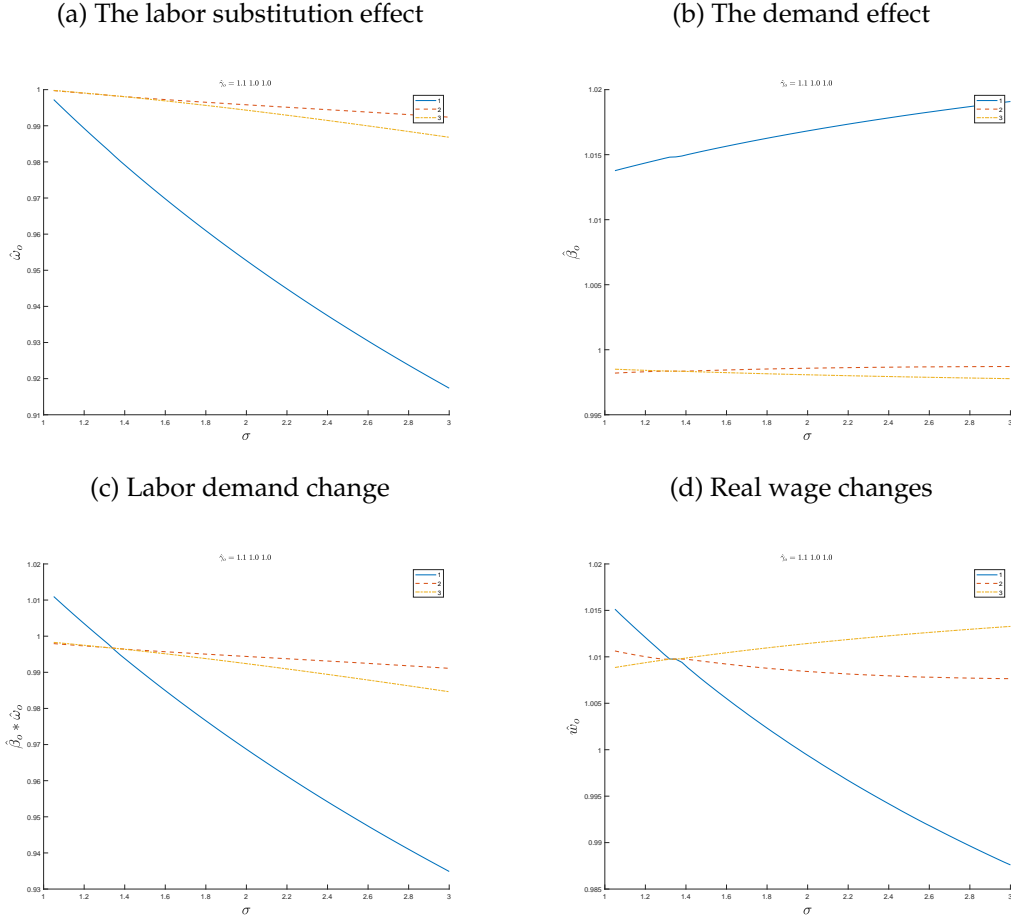
¹⁵Whether and for which σ the model generates negative real wage changes for a positive productivity shock depends on how large the expenditure share is on an occupation. In the current setup, we have low expenditure shares (=10%) for the directly affected occupation, which results in a modest aggregate productivity increase. In contrast, when all expenditure shares are set at 1/3, we only obtain negative real wage changes for high σ and small productivity shocks (see Appendix Figures B.3, B.4, and B.5).

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 (panel d).

Broader GE effects For the non-shocked, non-co-nested occupation, the real wage change increases with σ . This is due to the shocked occupation's price falling more and more and its expenditure share increasing, which raises real wages in the other occupations. Since real wages increase with σ , 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 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).

Figure 2: Role of σ



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.

3.1.2 Role of demand substitution elasticity (ψ)

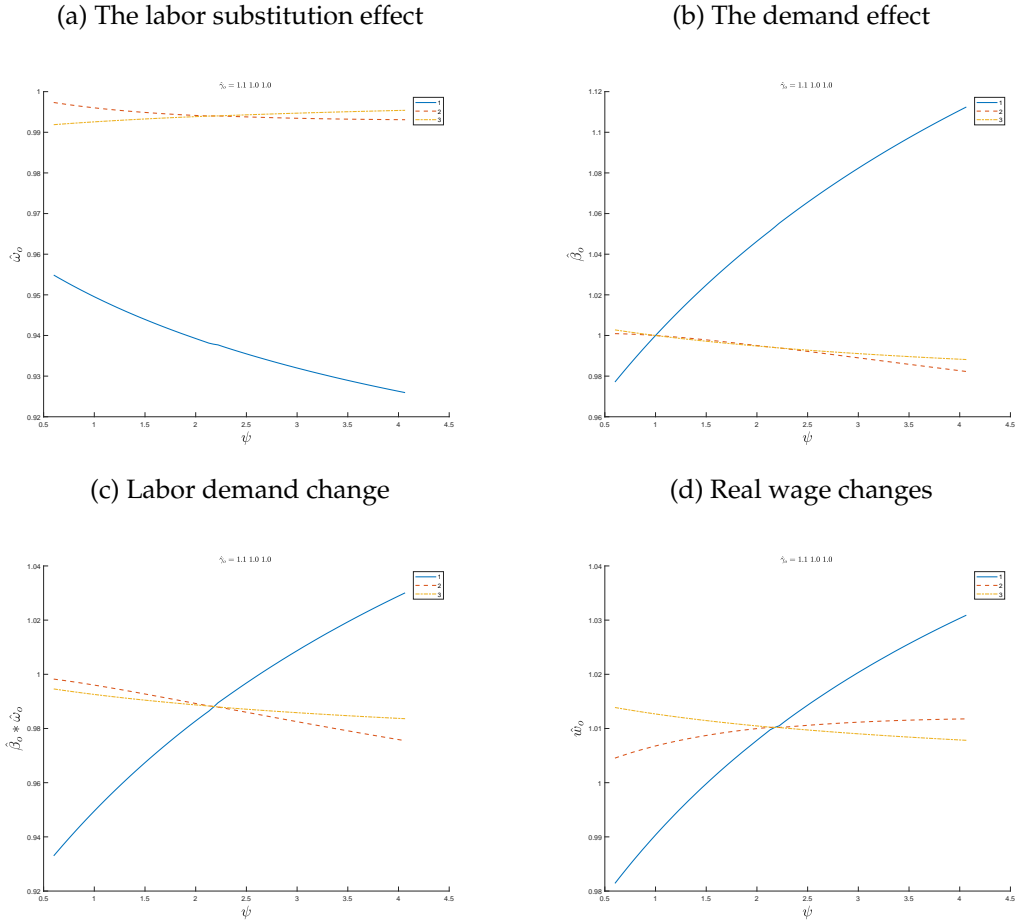
By definition, an increased demand elasticity (ψ) leads to stronger expenditure switches toward the good that experiences the positive productivity shock (Figure 3, panel b). In turn, this increase in demand for output from the occupation then leads to upward pressure on the occupation price (see Appendix 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 ψ (panel a). In addition, and perhaps surprisingly, this *supply-side* substitution effect deepens when the *demand* elasticity increases. Indeed, the larger is ψ , the stronger is the demand effect (recall panel b), which puts increasingly upward pressure on the occupation's wage (panel f), thereby further reducing the labor share. While the level drop in the labor share reflects a displacement effect from technical change, the further decline with ψ arises from a fully benign demand-side mechanism.

This is related to what we see in panel (d): real wage changes for the shocked occupation are low, even negative, when ψ is close to unity, but become larger and larger when ψ grows. As before, this is a reflection of the relative size of the labor substitution effect (governed by σ) and the demand effect (governed by ψ) on net labor demand changes (panel c). The higher ψ , the more important the demand effect becomes relative to the substitution effect.

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

Figure 3: Role of ψ



Notes: 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.

4 Data

For our empirical analysis, we utilize data from IPUMS USA, focusing on private sector employees aged 24-60, excluding the primary sector. We categorize these workers into 270 demographic groups based on gender, five education levels, three age groups, and the nine regional Census divisions.¹⁶

Following [Caunedo et al. \(2023\)](#), we classify employment into nine 1-digit level occupations (see [Table 1](#)). This level of aggregation minimizes the impact of sampling noise on the employment shares in the 270 worker groups observed in IPUMS. In addition, we group these occupations into either a routine or a non-routine nest, according to the [Autor and Dorn \(2013\)](#) classification. Their evidence on polarization in labor market outcomes across these two nests demonstrates that reallocation across these two nests is far from perfectly elastic, which makes it a natural choice for the nesting structure.

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

For the counterfactual analysis on the rise of generative A.I., we use the 2022 5-year ACS, the most recent IPUMS sample available, to obtain the necessary income and employment data.

¹⁶The 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	β_o	mean π_{go}	SD(π_{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 (β_o), the fifth the average employment share (π_{go}), the sixth the standard deviation across groups of the π_{go} , and the final column has the relative exposure to generative artificial intelligence as measured in [Eisfeldt et al. \(2023\)](#).

5 Estimation of occupational labor supply

5.1 Estimation strategy

5.1.1 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), 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

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

Rearranging and using that $\pi_{go|\mathcal{O}_m} = \frac{\pi_{go}}{\pi_{g\mathcal{O}_m}}$, we obtain:

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

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:

$$\ln \hat{I}_g = \alpha_o + \beta_1 \ln \hat{\pi}_{go|\mathcal{O}_m} + \beta_2 \ln \hat{\pi}_{g\mathcal{O}_m} + \varepsilon_{go}, \quad (14)$$

where $\alpha_o \equiv \ln \hat{w}_o$, $\beta_1 \equiv -\frac{1}{\kappa}$, $\beta_2 \equiv -\frac{1}{\mu}$, and $\varepsilon_{go} \equiv \ln \left(\hat{A}_{go}^{\frac{1}{\kappa_{gm}}} \hat{A}_{gm}^{\frac{1}{\mu_g}} \right)$. Here, $\ln \hat{\pi}_{go|\mathcal{O}_m}$ and $\ln \hat{\pi}_{g\mathcal{O}_m}$ measure reallocation, or occupational expansion more precisely, conditional on the occupational wage change \hat{w}_o . Abstracting from the error term, if a group's reallocation term is higher compared to its counterpart in the average group, then this group is more negatively exposed to the national-level wage changes.¹⁷ The two parameters then estimate the change in earnings associated with the change in specialization across nests (μ) and within-nests (κ).

The error term ε_{go} consists of occupation-by-group productivity shocks. To the extent that these shocks are correlated with demographic characteristics, we can 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.

¹⁷Note that this mechanism holds across all occupations within a given group. How is it possible for all occupations of a group to have higher percentage growth than the average group? The answer is simply that groups with above-average reallocation terms experience concentrated growth in the occupations where they are initially smaller, and vice versa. In other words, these groups expand in occupations where they have a comparative *dis*advantage and shrink in the ones where they have a comparative advantage.

5.1.2 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 β coefficient estimates, and since the implied reallocation elasticity is the negative inverse of the coefficient, there is also an upward bias for these implied elasticities.

As instruments, we employ model-based proxy variables for the within- and cross-nest wage indices, which are the denominators of the employment share changes ($\hat{\Phi}_{mg}$ and $\hat{\Phi}_g$). This approach builds on the insight from Galle et al. (2023) that the shift-share variable $\sum_o \pi_{go}\hat{r}_o$ very closely approximates $\sum_o \pi_{go}\hat{w}_o$, where $r_o \equiv \frac{w_o Z_o}{\sum_n w_n Z_n}$ is the national income share of an occupation. In our current setting, with nested employment shares, we employ $\sum_o \pi_{go|\mathcal{O}_m}\hat{r}_o$ and $\sum_m \pi_{g\mathcal{O}_m}\hat{r}_{\mathcal{O}_m}$ as instruments. Importantly, we confirm the relevance of these instruments in a first-stage analysis in Appendix Table C.1.

In terms of instrument validity, the IV estimation requires the error term, which is a productivity shock in a group-occupation cell, to be orthogonal to the instrument. Importantly, the national-level reallocation ($\hat{r}_o, \hat{r}_{\mathcal{O}_m}$) is allowed to be driven by any type of labor demand shock, as long as these are orthogonal to labor supply shocks. Hence, the exclusion restriction in our setup is less demanding than in papers focusing on, e.g., a single trade or technology shock (Autor et al., 2013; Acemoglu and Restrepo, 2020). We can further crystallize our exclusion restriction following the “exogenous shares” (Goldsmith-Pinkham, Sorkin, and Swift, 2020) or the “exogenous shocks” (Borusyak, Hull, and Jaravel, 2022) framework. Since the latter approach requires a sufficiently high number of different shocks, and since we have only nine occupations, it is more natural for us to focus on the former framework. Formally, our exclusion restriction is then $\mathbb{E}(A_{go}|\pi_{go}, \mathbf{X}_g) = 0$, with \mathbf{X}_g the vector of control variables. This restriction is violated if locations with a higher share also systematically have stronger (either positive or

negative) productivity shocks. Below, we further corroborate the validity of our exclusion restriction with an analysis of pre-trends, as suggested by Goldsmith-Pinkham et al. (2020).

5.2 Estimation results

Table 2 has our main estimation results, with columns 1-3 estimating different versions of Equation (14). 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 estimate the parameters using OLS, both without and with controls. In both cases, the implied values for κ and μ have the expected sign, but are very high. For instance, the lowest value we find for μ 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. 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.¹⁸ Specifically, in our empirical setup we have three different age groups (young, middle, and old), and we estimate a separate κ_g and μ_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 μ_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 κ_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

¹⁸We 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.

μ_g estimate (SE = 0.37), but also for their κ_g (SE = 1.4). For younger and middle-aged workers, the SEs range from 2.0 to 3.0.

The first stages for the estimations by age group lack some power, as the F-statistics are below 10. We address this issue 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.2, 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.

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

	(1)	(2)	(3)	(4)	(5)	(6)
	All (OLS)	All (OLS)	All	Young	Middle	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 μ	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 κ	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 (14), 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 α_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. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Sensitivity analysis

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.¹⁹

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. In line with the diagnostics proposed by Goldsmith-Pinkham et al. (2020), we analyze if there are pre-trends in group-level income, our dependent variable, correlated with the instruments. This would indicate that the instruments are correlated with long-term trends in income, irrespective of period-specific shocks to labor demand. 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 disappear during the pre-period, corroborating their validity. For the cross-nest instrument ($\sum_m \pi_{g\mathcal{O}m} \hat{r}_{\mathcal{O}m}$), we actually estimate, if anything, a negative coefficient in the pre-period ($p = 0.1$; see Appendix Table C.3, column 2). For the within-nest instrument ($\sum_o \pi_{go|\mathcal{O}m} \hat{r}_o$), we first control for nest-level expansion in the dependent variable, based on our model’s insights from specification (14). 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

¹⁹Unfortunately, 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).

completeness, we also provide results using the unadjusted outcome variable; however, these are less informative in the context of our model (columns 5-6).

6 Baseline quantification

In this section, we shed light on the quantitatively impact of generative A.I. on the labor market. To this end, we first parametrize our model and calibrate the A.I. shock.

6.1 Parametrization

Using our above estimates for age-specific reallocation elasticities, we set μ_g to 5.4 for young workers, 3 for middle-aged workers, and 1.4 for old workers. For κ_g , the values are 4.8, 3.6, and 2.9, respectively (see Table 2, columns 4-6). 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 similar finding of $\psi = 1.78$.

Given this value for ψ , it is crucial to ensure that the impact of changes in machine productivity align with the observed impacts of A.I. productivity gains at the occupation level; specifically, increased productivity and significant decreases in employment – see e.g., Dell'Acqua et al. (2023) and Hui, Reshef, and Zhou (2023).²⁰ To match these patterns, σ needs to be significantly higher than ψ . We therefore set $\sigma = 2.18$, the estimated σ_o in CJK for administrative services, the occupation most exposed to generative A.I. (see Table 1). Our assumption then is that A.I. technology exhibits the same degree of substitutability with labor in general as the previous episode of technical change has had with labor in administrative services.²¹

²⁰Specifically, Hui et al. (2023) focus on the labor market outcomes of freelancers on a large online platform after the release of ChatGPT. They find that freelancers in highly affected occupations experience a reduction in employment on both the intensive and the extensive margin and a 5.2% decline in earnings. At the same time, Dell'Acqua et al. (2023) find a productivity increase of at least 25.1% in an RCT randomizing access to GPT-4 among BCG consultants. Similar productivity increases are also found by Brynjolfsson, Li, and Raymond (2023) and Noy and Zhang (2023).

²¹The median estimate for σ_o in CJK is 1.32, though the confidence interval on that estimate can barely reject $\sigma_o = 2$. Importantly however, a value for σ below or close to $\psi = 1.34$ is inconsistent with the empirically observed increase in productivity but decline in employment and earnings after the introduction of A.I. This implies that previous episodes of capital-embodied technical change, including the introduction

6.2 Calibration of the shock

We calibrate the productivity shocks in our model to the empirically observed impact of exposure to ChatGPT. Specifically, we focus on the findings from [Dell'Acqua et al. \(2023\)](#), who randomly assigned access to GPT-4 among strategy consultants at Boston Consulting Group. They find that access to GPT-4 increased the speed by which tasks were completed by 25.1%. Note that this is a conservative measure of the increase in labor productivity in their study, as [Dell'Acqua et al. \(2023\)](#) also observe an increase in quality when using GPT-4.²²

We then extrapolate the productivity increase among strategy consultants to all occupations using the occupational exposure measure from [Eisfeldt et al. \(2023\)](#). This paper quantifies the share of tasks within each occupation that can be performed using large language models such as ChatGPT, providing exposure measures for 5-digit SOC occupational codes.²³ Using "management analysts" from the [Dell'Acqua et al. \(2023\)](#) study as our benchmark, we calculate the ratio of AI exposure for each 5-digit occupation relative to these management analysts. These relative exposure measures are then aggregated to our nine 1-digit occupational categorizations, using employment share weights. The final column of [Table 1](#) lists the exposure measures to AI across occupations. The most exposed occupations are cognitively oriented (e.g., administrative services, or scientific technicians), and the least exposed have a high manual labor content (e.g., mechanics and transportation, or low-skill services).

Our calibration target for each of the nine occupations is then the product of the rel-

of personal computers, often exhibited stronger complementarity with labor than A.I.

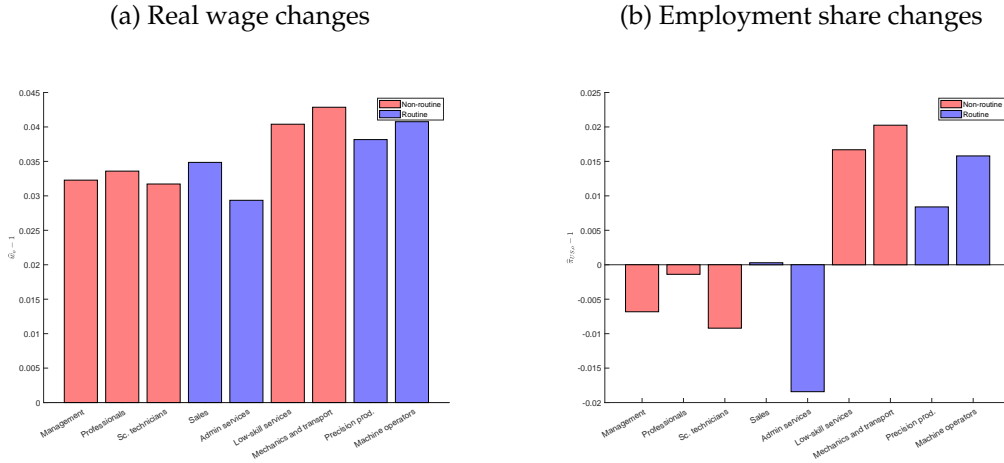
²²In their RCT, [Noy and Zhang \(2023\)](#) obtain similar results: in a writing task for college-educated professionals, exposure to ChatGPT decreased time to completion by 40% and increased output quality by 18%. However, their study focuses on a more narrow task, namely writing only, whereas [Dell'Acqua et al. \(2023\)](#) study the impact of generative A.I. on a broader set of tasks that are all relevant for a given occupation's output.

²³For each occupation, [Eisfeldt et al. \(2023\)](#) score all the tasks involved in each occupation, on whether a 50% reduction in completion time can be achieved using ChatGPT, either with or without the aid of additional tools built for ChatGPT. The measure of occupational exposure is then the share of tasks within an occupation where this completion speed can be achieved, with a 50% weight, in the numerator, for the tasks where additional tools are necessary. This measurement strategy closely follows the procedure in [Eloundou, Manning, Mishkin, and Rock \(2024\)](#), but we were not able to obtain the data from the latter paper.

ative exposure measure with the productivity increase from Dell’Acqua et al. (2023) (i.e., 25.1%), and we calibrate the vector of $\hat{\gamma}_o$ such that the model matches these productivity increases.²⁴ The calibrated shocks have a correlation of almost 94% with an occupation’s A.I. exposure (see Appendix Figure D.1).

6.3 Results

Figure 4: Labor market impact of the rise of generative A.I.



Despite the fact that generative A.I. affects non-routine occupations much more compared to previous waves of technical change, the growth of generative A.I. still leads to larger wage gains among non-routine occupations than among routine occupations (see Figure 4, panel a). This is because the routine occupation of “administrative services” has by far the highest exposure to A.I. (55%), and therefore also the strongest growth in machine productivity, as well as the lowest wage growth (see Appendix Figure D.1). The within-nest ripple effects of this lower wage growth then lead to subdued wage growth for all occupations in the routine nest. Associated with the decline in relative wages in the routine nest, there is reallocation out of this nest toward the non-routine nest, with

²⁴In the model, we measure changes in labor productivity as $\widehat{LP}_o = \frac{\hat{R}_o}{\hat{P}_o \hat{Z}_o}$. Noting that $\omega_o R_o = w_o Z_o$ and that $\hat{\omega}_o = \left(\frac{\hat{w}_o}{\hat{P}_o}\right)^{1-\sigma}$, we have that $\widehat{LP}_o = \hat{\omega}_o^{\frac{\sigma}{1-\sigma}}$.

highest growth in low-skill services and “mechanics and transportation.”

All occupations experience positive real wage growth, due to the aggregate increase in productivity. Within each nest, the manual-labor intensive occupations experience higher wage growth, because they have the lowest exposure to advances in A.I. 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 (4.1%) while wages in administrative services grow by around 2.9%. Analogously, in the non-routine nest, wages for “mechanics and transportation workers” grow by almost 4.3%, while wages in management grow by 3.2%. Overall, these wage changes are substantial, particularly considering that they only reflect the productivity gains from access to GPT-4 – a lower bound on the full productivity gain from recent advances in generative A.I.

Accounting for labor reallocation and associated ripple effects is a strength of our general-equilibrium framework, which reduced-form analysis of the impact of technical change typically do not capture. Indeed, reduced-form approaches often measure the impact of a new technology on the labor market through exposure measures. Interestingly, in the current case, the negative correlation between A.I. exposure and the wage changes remains very high, at 99% (see Appendix Figure D.1, panel c). Nevertheless, 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).

7 Extended model with intensive and extensive margin

Shocks to labor demand across occupations lead to occupational reallocation, but can also affect hours worked per worker and the unemployment rate. To capture these additional margins of labor adjustment, we here extend the model with an intensive margin

decision and a bare-bones search-and-matching model, which will also allow us to have a more comprehensive quantification of the labor market impacts of technical change.

7.1 Theory

We add an intensive and an extensive margin to the labor supply side, but the labor demand side stays identical to its baseline setup. This extended model nests our baseline model, except for one 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.²⁵

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 χ_g . After being hired, workers make a decision on how many hours to work. 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 ν_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

²⁵Using 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 intensive margin decision.

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.3, 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 Φ_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

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

where $1/\xi_g$ is the intensive margin elasticity, δ_g is the demand shifter for consumption, and $\tilde{\eta}_g \equiv \Gamma(1 - 1/(\xi_g \mu_g))$. Intuitively, as the wage index Φ_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 χ_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:

$$I_g \propto \Phi_g^{\frac{1+\xi_g}{(1-\chi_g)\xi_g}} L_g. \quad (15)$$

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 the appendix.

7.2 Counterfactual analysis

Parametrization The parametrization and calibration of the model remain as in the baseline model, except we set $\mu_g = 4.739$ for middle-aged workers such that all $\mu_g < \kappa_g$. In addition, we set the value of the two new parameters 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 $\xi = 2.5$, 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 group-level 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).

Table 3: Labor's adjustment margins after the rise of generative A.I.

	Aggregate	Mean	SD	Min.	Max.
\widehat{I}_g	7.09	7.21	0.45	6.72	8.20
\widehat{i}_g	3.43	3.54	0.22	3.30	4.02
\widehat{h}_g	1.40	1.40	0.09	1.31	1.59
\widehat{e}_g	2.11	2.11	0.13	1.97	2.39

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

For the US in the aggregate, real income increases by 7.1% due to the rise of generative

A.I. (see Table 3). These aggregate gains are unequally distributed, with a maximum group-level gain of 8.2% and a minimum gain of 6.7%. Given our parametrization, the change in hourly income accounts for roughly 50% of the aggregate gains, the changes in the employment rate and hours worked for respectively 30% and 20%.

Previous episodes of technical change typically negatively affected the returns to education – commonly categorized as skill-biased technical change (Acemoglu, 1998; Krusell et al., 2000; Hémous and Olsen, 2022). For generative A.I., this pattern is reversed. Indeed, one of the subgroups gaining the most are high-school drop-outs, who gain 8% in the aggregate.²⁶ 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, college workers gain 1.1 percentage points less in the aggregate.

Table 4: Income changes due to generative A.I. across demographic groups

	Aggregate	Mean	SD	Min.	Max.
All groups	7.09	7.21	0.45	6.72	8.20
Young	7.15	7.25	0.45	6.72	8.20
Middle aged	7.07	7.20	0.45	6.74	8.15
Old	7.07	7.17	0.46	6.72	8.10
Male	7.20	7.38	0.50	6.72	8.20
Female	6.88	7.03	0.31	6.72	7.72
Less than high school	7.95	7.81	0.25	7.40	8.20
High school	7.56	7.45	0.35	6.94	7.93
Some college	7.20	7.17	0.28	6.76	7.66
College	6.81	6.81	0.05	6.72	6.92
Post-graduate degree	6.79	6.80	0.02	6.72	6.84


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 A.I. than initially richer groups. Given that advances in A.I. tend to benefit workers in manual-

²⁶These findings echo the earlier insightful analysis of Bloom, Prettner, Saadaoui, and Veruete (2024), with our theoretical setup being more general.

labor intensive occupations most, this may not be surprising. Specifically, we obtain a negative correlation of 41% 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 -41%, 83% 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 A.I. along the income distribution



8 Heterogeneous within-nest reallocation elasticities

The extent of within-nest spillovers strongly depends on the within-nest reallocation elasticity. Here, we examine how our findings change when we allow these reallocation elasticities to vary by nest.

8.1 Estimation

Compared to our baseline specification (14), we now allow for heterogeneity in κ_m across the two nests (R for routine and NR for non-routine) by estimating:

$$\ln \hat{I}_g = \alpha_o + \beta_1 \ln \hat{\pi}_{go|\mathcal{O}_m} + \beta_2 1(o \in \text{NonRoutine}) * \ln \hat{\pi}_{go|\mathcal{O}_m} + \beta_3 \ln \hat{\pi}_{g\mathcal{O}_m} + \varepsilon_{go}. \quad (16)$$

In this case, $\beta_1 = -\frac{1}{\kappa_R}$ and $\beta_1 + \beta_2 = -\frac{1}{\kappa_{NR}}$. We estimate this specification in Table 5, column 4.

Focusing first on the cross-nest reallocation elasticity, this value falls from 3.4 in the baseline estimation (column 3) to $\mu = 1.2$ in this specification. This value is close to the theoretical lower bound on the Fréchet dispersion parameter, namely unity, and implies strong distributional effects. After all, a unity value would imply the supply of effective labor units to a nest is inelastic, which makes workers fully exposed to the wage changes in that nest. Interestingly, in specification 4, the estimated reallocation elasticity within the routine nest is even lower, with a point-estimate of $\kappa_R = 0.74$ and a standard error of 0.36. Hence, we cannot reject that $\kappa_R = 1.201$. This is the value we adopt in the counterfactual analysis, since this version of the model requires $\kappa_m > \mu$. In contrast, reallocation within the non-routine nest appears substantially more elastic, with $\kappa_{NR} = 13.0$. This elasticity is not precisely estimated though, with a standard error of 9.0.^{27 28}

8.2 Counterfactual analysis

Allowing for nest-specific reallocation elasticities alters the impact of A.I. on the wage distribution, relevant for within-group inequality, while the cross-group distributional

²⁷Since we obtain a low F-stat in the first-stage of this specification, as in the also estimate the inverted specification, where we obtain an F-stat of 7.6 (see Appendix Table C.6).

²⁸We have also estimated κ_R and κ_{NR} separately by age group. Due to the loss of power in that estimation, the first stage is typically very low. Nevertheless, these estimates continue to suggest that $\kappa_{NR} > \kappa_R$, in line with the results presented here.

Table 5: Cross-nest heterogeneity in κ_m

	(1)	(2)	(3)	(4)
	$\ln \hat{I}_g$	$\ln \hat{I}_g$	$\ln \hat{I}_g$	$\ln \hat{I}_g$
$\ln \hat{\pi}_{g O_m}$	-0.015 (0.025)	-0.060*** (0.011)	-0.29*** (0.098)	-0.83** (0.33)
$\ln \hat{\pi}_{go O_m}$	-0.019** (0.0080)	-0.0046 (0.0036)	-0.19* (0.11)	-1.35** (0.66)
$1(\text{Non-routine nest}) * \ln \hat{\pi}_{go O_m}$				1.28* (0.70)
Implied μ	66.3 (109.9)	16.6 (3.06)	3.40 (1.13)	1.20 (0.47)
Implied κ	53.6 (23.0)	216.6 (170.8)	5.27 (3.10)	
Implied κ_R				0.74 (0.36)
Implied κ_{NR}				13.0 (9.03)
KP F-stat			11.1	2.02
Controls	No	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes
Estimation	OLS	OLS	IV	IV
Observations	2430	2430	2430	2430

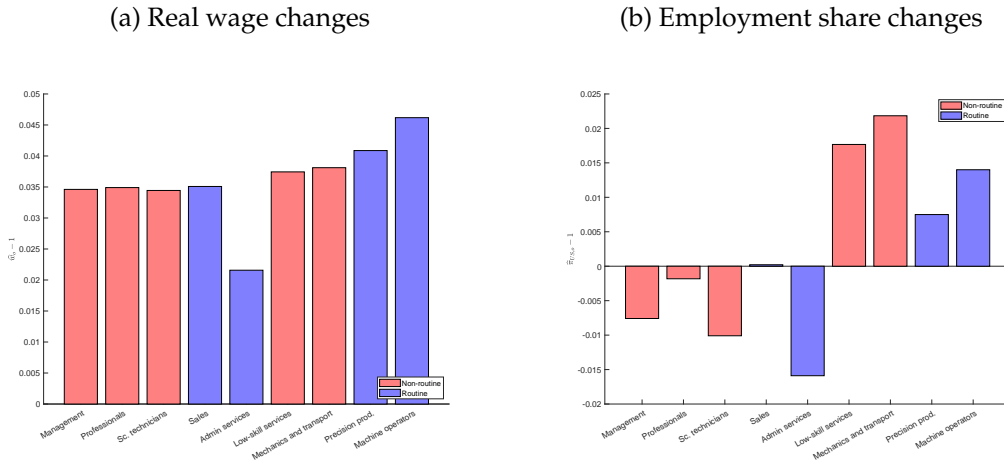
Notes: The table repeats the estimation of equation (14) in columns 1-3, and estimates equation (16) in column 4. $\ln \hat{I}_g$ is measured as the log change in average hourly wage in a group, and α_o is an occupation fixed-effect. Specifications 1 and 2 are estimated with OLS and the others with IV, with instruments $\sum_o \pi_{go|O_m} \hat{r}_o$ and $\sum_m \pi_{g|O_m} \hat{r}_{O_m}$ in column 3, and additionally $1(o \in \text{NonRoutine}) * (\sum_o \pi_{go|O_m} \hat{r}_o)$ in column 4. All specifications except the first control for gender FE, education level FE, and Census division FE and age-bin FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. (see Appendix Table C.6) estimates the inverted specifications, which often have stronger first stages. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

effects are much less affected. Indeed, comparing the current group-level income changes with those from the previous model, with no differentiation in κ across nests, we find that the correlation is 86% (see Appendix Figure D.4). While the income effects are slightly lower in the current model, overall the differences with the income changes in the previous model are minor (see Appendix Table D.1).

Compared to the baseline model, the main change in the counterfactual results is that

the dispersion in wage changes is lower in the non-routine nest, with all wage changes now around 3.5%, and higher in the routine nest, where the *gap* between administrative services and machine operators grows to 2.6 percentage points (see Figure 6). These changes are the direct result of the high $\kappa_{NR} = 13$ and the low $\kappa_R = 1.201$, leading to strong wage convergence and divergence in the non-routine and routine nests, respectively.

Figure 6: Labor market impact of the rise of generative A.I.



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 A.I. on labor market outcomes. While the most exposed occupations, such as administrative services, gain substantially less than manual-labor intensive occupations, we find that within-nest wage convergence, arising from within-nest reallocation, can be substantial. Across several model specifications, lower educated workers always experience higher wage growth than higher educated workers.

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. (2024). Robots and wage polarization: The effects of robot capital by occupations. Technical report, mimeo.
- Autor, D. (2024). Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research.
- 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.

- 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. Technical report, National Bureau of Economic Research.
- Bernon, B. and G. Magerman (2024). A theory of equality and growth.
- Bloom, D. E., K. Prettnner, J. Saadaoui, and M. Veruete (2024). Artificial intelligence and the skill premium. Technical report, National Bureau of Economic Research.
- 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., D. Li, and L. R. Raymond (2023). Generative ai at work. Technical report, National Bureau of Economic Research.
- 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.
- 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. Candelon, and K. R. Lakhani (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (24-013).
- Eisfeldt, A. L., G. Schubert, and M. B. Zhang (2023). Generative ai and firm values. Technical report, National Bureau of Economic Research.
- 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*.
- 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.
- 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 (2023). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Available at SSRN* 4527336.
- Humlum, A. (2022). Robot adoption and labor market dynamics.
- Humlum, A. and E. Vestergaard (2024). The adoption of chatgpt. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2024-50).
- 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.
- Korinek, A. (2023). Scenario planning for an a (g) i future. *IMF Finance & Development Magazine* 60(4), 30–33.
- Korinek, A. and J. E. Stiglitz (2018). Artificial intelligence and its implications for income

- distribution and unemployment. In *The economics of artificial intelligence: An agenda*, pp. 349–390. University of Chicago Press.
- 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.
- 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. Technical report, 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.
- 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.

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:

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

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 κ), we can write this as

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

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 :

$$\ln \left(\frac{\hat{w}_o}{\hat{w}_n} \right) = \left(\frac{\sigma - \psi}{\kappa + \sigma - 1} \right) \ln \left(\frac{\hat{P}_o}{\hat{P}_n} \right). \quad (17)$$

A.3 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:

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

This choice on hours results in a real income of

$$\delta_g^{\frac{1}{\xi_g}} (\nu_g w_o z_o)^{\frac{1+\xi_g}{\xi_g}}. \quad (18)$$

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

$$\frac{\xi_g}{1 + \xi_g} e_g (\delta_g \nu_g w_o z_o)^{\frac{1+\xi_g}{\xi_g}}. \quad (19)$$

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

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

which implies that a worker with productivity vector z will work in occupation o iff $z \in \Omega_o(w)$. 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:

$$\pi_{go|\mathcal{O}_m} = \frac{A_{go}w_o^{\kappa_{gm}}}{\tilde{\Phi}_{mg}^{\kappa_{gm}}}, \quad (20)$$

$$\pi_{g\mathcal{O}_m} = \frac{(\sum_{n \in \mathcal{O}_m} A_n w_n^{\kappa_{gm}})^{\mu_g / \kappa_{gm}}}{\Phi_g^{\mu_g}}, \quad (21)$$

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

$$\Phi_g \equiv \left(A_{gHP}^{1/\kappa_{HPg}} w_{HP}^{\mu_g} + \sum_{m \neq HP} \tilde{\Phi}_{mg}^{\mu_g} \right)^{1/\mu_g}. \quad (22)$$

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 18, average real income for workers that applied to a sector is therefore

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

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

$$I_g \equiv \sum_o w_o Z_{go} = \eta_g e_g (\delta_g \nu_g)^{\frac{1}{\xi_g}} \Phi_g^{\frac{1+\xi_g}{\xi_g}} L_g. \quad (23)$$

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

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

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 ν_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 $(1 - \nu_g)I_{go}$. The cost of posting a vacancy is assumed to be c_g . Hence, the zero-profit condition (ZPC) for firms entails

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

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 χ_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

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

Solving for labor market tightness:

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

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:

$$e_g = (A_g^M)^{\frac{1}{1-\chi_g}} \left(\frac{(1 - \nu_g) \eta_g}{c_g} \right)^{\frac{\chi_g}{1-\chi_g}} (\delta_g \nu_g)^{\frac{\xi_g}{\xi_g(1-\chi_g)}} \Phi_g^{\frac{\chi_g(1+\xi_g)}{(1-\chi_g)\xi_g}}. \quad (24)$$

Intuitively, a shock that increases Φ_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:

$$ELD_o = \omega_o \beta_o \hat{\omega}_o \hat{\beta}_o \sum_n \sum_g \frac{\pi_{gn} \hat{\pi}_{gn} I_g \hat{I}_g}{\omega_n \hat{\omega}_n} - \sum_g \pi_{go} \hat{\pi}_{go} I_g \hat{I}_g. \quad (25)$$

Importantly though, we will need to take into account the intensive and extensive margin in solving for \hat{I}_g . First note from Equation (24) that

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

while from (23) 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.


$$\begin{aligned} \hat{\Phi}_g &= \left(\pi_{gHP} + \sum_{m \neq HP} \pi_{g\mathcal{O}_m} \left(\sum_{o \in \mathcal{O}_m} \pi_{go|\mathcal{O}_m} \hat{A}_{go} \hat{w}_o^{\kappa_{gm}} \right)^{\frac{\mu_g}{\kappa_{gm}}} \right)^{1/\mu_g}, \\ \hat{\pi}_{go} &= \hat{\pi}_{go|\mathcal{O}_m} \hat{\pi}_{g\mathcal{O}_m} \\ \hat{\pi}_{go|\mathcal{O}_m} &= \frac{\hat{A}_{go} \hat{w}_o^{\kappa_{gm}}}{\sum_{n \in \mathcal{O}_m} \pi_{gn|\mathcal{O}_m} \hat{A}_{gn} \hat{w}_n^{\kappa_{gm}}}, \\ \hat{\pi}_{g\mathcal{O}_m} &= \frac{\left(\sum_{o \in \mathcal{O}_m} \pi_{go|\mathcal{O}_m} \hat{A}_{go} \hat{w}_o^{\kappa_{gm}} \right)^{\frac{\mu_g}{\kappa_{gm}}}}{\hat{\Phi}_g^{\mu_g}}, \\ \hat{w}_o &= \frac{\hat{w}_o^{1-\sigma}}{(1-\omega_o)\hat{\gamma}_o + \omega_o \hat{w}_o^{1-\sigma}}, \\ \hat{\beta}_o &= \hat{\nu}_o \hat{P}_o^{1-\psi}, \\ \hat{P}_o &= \hat{\gamma}_o^{-\nu} [\hat{\gamma}_o(1-\omega_o) + \hat{w}_o^{1-\sigma} \omega_o]^{\frac{1}{1-\sigma}}. \end{aligned} \tag{26}$$

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 σ


(a) Wage vs. occupation price change



(b) Relative occupation price change




(c) Employment share changes




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 $\hat{\omega}_o = \left(\hat{w}_o / \hat{P}_o \right)^{1-\sigma}$ (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 $\hat{\beta}_o = \left(\hat{P}_o \right)^{1-\psi}$ (Figure 2, panel b). Finally, panel (c) shows the change in employment shares for the occupations.

Figure B.2: Role of ψ


(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 $\hat{\omega}_o = \left(\hat{w}_o / \hat{P}_o \right)^{1-\sigma}$ (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 $\hat{\beta}_o = \left(\hat{P}_o \right)^{1-\psi}$ (Figure 3, panel b). Finally, panel (c) shows the change in employment shares for the occupations.

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


(a) The labor substitution effect




(b) The demand effect



(c) Labor demand change




(d) Real wage changes




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 $\hat{\gamma}_o = 1.1$ 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.4: Role of σ with equal expenditure shares ($\beta_o = 1/3$)


(a) Wage vs. occupation price change



(b) Relative occupation price change




(c) Relative occupation price change



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 $\hat{\gamma}_o = 1.1$ 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 $\hat{\omega}_o = \left(\hat{w}_o/\hat{P}_o\right)^{1-\sigma}$ (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 $\hat{\beta}_o = \left(\hat{P}_o\right)^{1-\psi}$ (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 σ with equal expenditure shares ($\beta_o = 1/3$)




The figure shows wage changes for a larger range on σ 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 σ . As documented in Figure 2, note that for $\hat{\gamma}_o = 1.1$, the wage changes also become negative for large σ when the expenditure share on the shocked occupation is lower.

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


(a) The labor substitution effect




(b) The demand effect



(c) Labor demand change




(d) Real wage changes




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 σ when $\sigma < 1$


(a) Wage vs. occupation price change



(b) Relative occupation price change



(c) Relative occupation price change



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 $\hat{\omega}_o = \left(\hat{w}_o / \hat{P}_o \right)^{1-\sigma}$ (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 $\hat{\beta}_o = \left(\hat{P}_o \right)^{1-\psi}$ (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)
	$\ln \hat{\pi}_{g O_m}$	$\ln \hat{\pi}_{g O_m}$	$\ln \hat{\pi}_{gO_m}$	$\ln \hat{\pi}_{gO_m}$
$\sum_{o \in O_m} \pi_{go O_m} * \hat{r}_o$	-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 (14). 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: Inverted estimation of reallocation elasticities (Dep. var.: $\ln \hat{\pi}_{gO_m}$)

	(1)	(2)	(3)	(4)	(5)	(6)
	All (OLS)	All (OLS)	All	Young	Middle	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 μ	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 κ	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. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.3: Reduced forms for actual and pre- period

	$\ln \hat{I}_g$		$\ln \hat{I}_g + \frac{1}{\hat{\mu}} \ln \hat{\pi}_{gO_m}$		$\ln \hat{I}_g$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\sum_{O_m} \pi_{gO_m} * \hat{r}_{O_m}$	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 (14). 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.4: Weighted estimation of age-specific reallocation elasticities. (Dep. var.: $\ln \hat{I}_g$)

	(1)	(2)	(3)	(4)	(5)	(6)
	All (OLS)	All (OLS)	All	Young	Middle	Old
$\ln \hat{\pi}_{g\mathcal{O}_m}$	-0.046* (0.024)	-0.052*** (0.010)	-0.34*** (0.13)	-0.26** (0.12)	-0.38 (0.29)	-0.75*** (0.21)
$\ln \hat{\pi}_{go \mathcal{O}_m}$	-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 μ	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 κ	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 (14), namely $\ln \hat{I}_g = \alpha_o + \beta_1 \ln \hat{\pi}_{go|\mathcal{O}_m} + \beta_2 \ln \hat{\pi}_{g\mathcal{O}_m} + \varepsilon_{go}$, 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 α_o is an occupation fixed-effect. Specifications 1 and 2 are estimated with OLS, and the others with IV, with instruments $\sum_o \pi_{go|\mathcal{O}_m} \hat{r}_o$ and $\sum_m \pi_{g\mathcal{O}_m} \hat{r}_{\mathcal{O}_m}$. 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. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

	(1)	(2)	(3)	(4)	(5)	(6)
	All (OLS)	All (OLS)	All	Young	Middle	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 μ	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 κ	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. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.6: Cross-nest heterogeneity in κ_m , with inverted specifications


	(1)	(2)	(3)	(4)
	$\ln \hat{\pi}_{gO_m}$	$\ln \hat{\pi}_{gO_m}$	$\ln \hat{\pi}_{gO_m}$	$\ln \hat{\pi}_{gO_m}$
$\ln \hat{I}_g$	-0.018 (0.029)	-0.22*** (0.040)	-3.40*** (1.13)	-1.20** (0.47)
$\ln \hat{\pi}_{go O_m}$	-0.052*** (0.011)	-0.054*** (0.011)	-0.64*** (0.19)	-1.62*** (0.25)
$1(\text{Non-routine nest}) * \ln \hat{\pi}_{go O_m}$				1.53*** (0.43)
Implied μ				
Implied κ				
Implied κ_R				
Implied κ_{NR}				
KP F-stat				
Controls				Yes
Occupation FE				
Estimation				
Observations	2430	2430	2430	2430

Notes: The table repeats the estimation of equation: (14) in columns 1-3, and in column 4 estimates $\ln \hat{\pi}_{gO_m} = \alpha_o + \beta_1 \ln \hat{I}_g + \beta_2 \ln \hat{\pi}_{go|O_m} + \beta_3 1(o \in \text{NonRoutine}) * \ln \hat{\pi}_{go|O_m} + \varepsilon_{go}$. 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 $\sum_m \pi_{gO_m} \hat{r}_{O_m}$ in column 3, and additionally $1(o \in \text{NonRoutine}) * (\sum_o \pi_{go|O_m} \hat{r}_o)$ in column 4. All specifications except the first control for gender FE, education level FE, Census division FE and age-bin FE. Standard errors are clustered at the level of the demographic group, defined by gender, education level, age bin, and Census division. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Appendix D Supplementary counterfactual results

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 A.I. (AI_o), measured as in [Eisfeldt et al. \(2023\)](#), the calibrated machine-productivity shocks for an occupation ($\hat{\gamma}_o$) and the resulting counterfactual wage changes (\hat{w}_o) for the baseline quantification.

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 A.I. (AI_o), measured as in [Eisfeldt et al. \(2023\)](#), the calibrated machine-productivity shocks for an occupation ($\hat{\gamma}_o$) 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.

Table D.1: Income changes due to generative A.I. with nest-specific κ

	Aggregate	Mean	SD	Min.	Max.
All groups	7.09	7.12	0.33	6.54	7.84
Young	7.08	7.12	0.31	6.64	7.84
Middle aged	7.09	7.13	0.32	6.60	7.80
Old	7.09	7.11	0.36	6.54	7.83
Male	7.22	7.32	0.27	7.00	7.84
Female	6.83	6.92	0.26	6.54	7.59
Less than high school	7.61	7.51	0.22	6.98	7.84
High school	7.27	7.14	0.40	6.54	7.72
Some college	7.05	6.99	0.33	6.54	7.48
College	6.97	6.94	0.10	6.79	7.06
Post-graduate degree	7.03	7.02	0.03	6.96	7.07

Notes: The table presents the income changes (\hat{I}_g), split by demographic group, for the extension with $\kappa_R = 1.201$ and $\kappa_{NR} = 13$ (Section 8.2).

Figure D.4: Comparison of income changes (\hat{I}_g) across model versions

