I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17851

The Impact of Labour Demand Shocks When Occupational Labour Supplies Are Heterogeneous

Michael J. Böhm
Ben Etheridge
Aitor Irastorza-Fadrique

APRIL 2025

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17851

The Impact of Labour Demand Shocks When Occupational Labour Supplies Are Heterogeneous

Michael J. Böhm
TU Dortmund University and IZA

Ben Etheridge
University of Essex, IFS and MiSoC

Aitor Irastorza-Fadrique
Institute for Fiscal Studies (IFS)

APRIL 2025

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Impact of Labour Demand Shocks When Occupational Labour Supplies Are Heterogeneous*

As technological advances accelerate and labour demands shift, the ability of workers to reallocate across occupations will be crucial for shaping labour market dynamics, inequality, and effective policy design. In this paper, we develop a tractable equilibrium model of the labour market that incorporates heterogeneous labour supply elasticities to different occupations and across different occupation pairs. Using worker flows from German administrative data, we estimate these elasticities and validate them through external measures such as occupational licensing and task distance. Our model quantifies the heterogeneous impacts of recent labour demand shifts on occupational wages and employment, highlighting the role of cross-occupation effects in shaping market responses to shocks. Finally, we leverage this framework to project employment flows and wage adjustments under future occupational demand shifts that are implied by the latest automation technologies.

JEL Classification: J21, J23, J24, J31

Keywords: heterogeneous labour supply elasticities, occupational substitutability, labour demand shocks, german panel data, job flows, occupational employment and wages, automation technologies, future projections

Corresponding author:
Michael J. Böhm
Technische Universität Dortmund
Vogelpothsweg 87
44227 Dortmund
Germany

E-mail: michael.j.boehm@tu-dortmund.de

* We thank Christian Bayer, Lukas Buchheim, Wolfgang Dauth, Thomas Cornelissen, Emma Duchini, Khalil Esmkhani, Giovanni Gallipoli, Georg Graetz, David Green,Wenchao Jin, Gizem Kosar, Rafael Lalive, Peter Levell, Helen Miller, Nitya Pandalai-Nayar, Itay Saporta-Eksten, Oskar Nordstr¨om Skans, Tom Waters, and Jens Wrona as well as participants at various seminars and conferences for helpful comments. Etheridge thanks support from the ESRC Centre for Micro-Social Change, award number ES/S012486/1. Irastorza-Fadrique acknowledges doctoral scholarship support from the University of Essex Social Sciences and the Institute for Fiscal Studies. Irastorza-Fadrique also thanks support from the ESRC Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant reference ES/M010147/1) and through the grant 'Productivity, wages and the labour market' (grant reference ES/W010453/1). The data basis of this paper is the weakly anonymous Sample of Integrated Labour Market Biographies (version 7519). Remote data access was provided by the Research Data Centre (FDZ) of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB) under contract number 'fdz1935'. All errors remain ours.

1 Introduction

Over the past few decades, the labour market has undergone significant transformations, driven by technological advancements and other structural changes. These shifts are particularly evident in the evolving occupational structure, where some occupations have experienced large declines while others have expanded strongly (Acemoglu & Autor, 2011; Goos et al., 2014). With rapid advancements in automation, digitalisation, and the wider use of artificial intelligence technologies, the pace of these changes is expected to accelerate even further (Agarwal et al., 2024; Brynjolfsson et al., 2025; Cui et al., 2025).

The contribution of this paper is to specify elasticities of labour supply which characterise how labour demand shocks translate into such occupational employment and wage changes in equilibrium. Using the German labour market as a laboratory, we employ these elasticities to study the effects of past occupational supply and demand shifts, and to project the impact of currently predicted shocks on occupational changes and employment flows. Given our model's tractability, we can transparently characterise and quantify the rich pathways.

Our approach takes seriously the idea that labour supplies to different occupations and across different occupation pairs are differentially elastic.[1] This heterogeneity leads to a fuller interpretation of occupational changes in equilibrium than was previously available, and it has several novel implications for labour market policy. We estimate our model by using workers' observed flows across occupations, which identify the relative sizes of the elasticities, and by instrumenting demand shocks with occupations' initial task contents.

A key conceptual contribution of this paper is how we specify the heterogeneous elasticities of occupational labour supply. We generalise a standard static random utility theory of workers' occupational choice by allowing for pairwise occupational transition costs. This setup leads to sufficient statistics for the labour supply elasticities in terms of pairwise switching probabilities. Since empirical transition probabilities vary substantially, the model predicts that supply elasticities are also heterogeneous across occupations. As comes naturally out of the model, we distinguish between 'cross-price' elasticities, which capture the impact on employment of wage changes in a different occupation, and 'own-price' elasticities, which capture the impact of wage changes in the occupation itself.

[1] Among canonical models of occupational labour markets, Acemoglu & Autor (2011) and Acemoglu & Restrepo (2022) specify the labour supply to tasks as perfectly elastic. Autor et al. (2003), Autor et al. (2006) and Cortes (2016) allow for imperfectly elastic market supply but in a restricted manner and without direct quantitative interpretation. See also, as examples of models in the same broad classes, Autor & Dorn (2013); Jung & Mercenier (2014); Goos et al. (2014); Burstein et al. (2019); Gregory et al. (2022).

Our sufficient statistics approach allows the heterogeneity of supply elasticities to be directly estimated from gross worker flows between occupations in steady state. While these baseline flows and the implied elasticities are in principle endogenous, we show that they need not be affected by shifts in occupational wages and find empirically that they remain stable over time. We further discuss extensions to our occupational choice model, including a dynamic version with occupation-specific human capital accumulation and forward-looking decisions under uncertainty, as well as a version with different non-employment states including unemployment. We demonstrate that, in all considered cases, differential worker flow rates remain informative about occupational substitutabilities on the labour supply side, and elasticities can be constructed from similar moments to those in the basic model.

Another key conceptual contribution is to embed heterogeneous labour supply elasticities into an equilibrium model.[2] Specifically, we show how endogenous price changes are determined in terms of fundamental shocks and spillovers from these shocks across occupations. These spillovers exhibit both own- and cross-occupation effects, which we again identify using worker flows and estimate to be substantially heterogeneous. This framework enables a theory-consistent IV estimation of the labour supply side, where demand shock proxies are interacted with their corresponding spillovers to instrument wages. More broadly, our model's full system of simultaneous equations allows for an analysis of standard demand and supply shocks, as well as the novel role of heterogeneity in shaping occupational changes in equilibrium over time.

We apply our model using the German Sample of Integrated Employment Biographies (SIAB), a comprehensive panel dataset tracking workers' careers over time. The SIAB provides a detailed and consistent occupational classification from 1975 to 2010, which we leverage for historical analysis, and an updated classification with greater granularity in emerging occupations, used for predictions from 2022 onward. In our historical sample, we estimate supply elasticities and spillovers using worker transition flows from 1975–1984, conducting our primary analysis over 1985–2010. Own-price elasticities vary significantly across occupations, and while cross-price elasticities are often near zero, they exhibit substantial variation – some occupation pairs show strong substitutability, where wage changes in one lead to significant employment shifts in another. Own-price elasticities tend to be lower in occupations with greater licensing and job restrictions, whereas cross-price elasticities correlate with task distance between occupations. Our empirical elasticities improve on prior measures by providing a cardinal, quantitative

[2]We keep occupational demand relatively standard CES in order to focus on the role of heterogeneity on the labour supply side. We calibrate the elasticity of substitution in production following the literature and our key results are robust to a whole range of values for that demand side parameter.

interpretation. In the prediction sample, we estimate supply elasticities using worker flows in the decade leading up to 2022 and incorporate expert assessments on how recent technologies may replace occupational tasks. Using these elasticities and predicted demand shocks, we compute the model's equilibrium projections of employment, wages, and net flows between occupation pairs.

Our analyses yield three main findings. First, occupational wages and employment changed significantly during the period 1985–2010. We show that heterogeneous labour supply elasticities play a key role in these changes, explaining why some occupations primarily experienced employment shifts in response to shocks, while others saw wage adjustments. This elasticity heterogeneity accounts for more than 20% of the variation in occupational employment growth, while 60% is driven by demand shocks and the remaining 20% by shifts in labour supply across occupations.

Second, cross-price elasticities of labour supply play a particularly important role. We find that labour demand shocks are correlated among relatively cross-price elastic occupations. This reduces opportunities for workers in declining jobs to transition into rising substitutable occupations, lowers overall labour supply elasticities compared to a scenario where shocks and elasticities were randomly distributed, and slows employment adjustment to structural changes. We also see greater dispersion in occupational wage changes and larger increases in wage inequality as a result of correlated shocks.

The third set of results comes from our prediction sample and focuses on the period after 2022. The model projects that occupational demand shocks driven by technological replacement will increase employment in IT and construction-related occupations, raise wages in the health and education sectors, and lead to declining wages in manufacturing-related jobs as well as certain high-skilled occupations, such as accountants and auditors.[3] Cross-elasticities of labour supply again play a crucial role in these dynamics. For instance, manufacturing workers have limited attractive substitute occupations to transition into, whereas IT occupations can draw workers from a broad range of technical and business-related jobs. But the equilibrium determination of wage changes themselves is another key factor, as these changes ultimately set the viable pathways for adjusting to shocks given the occupational cross-elasticities.[4]

[3]To be precise, we scale the technological replacement potentials so that the relatively less affected health (except laboratory medicine), education, and IT occupational fields experience relatively positive demand shocks. The discussed wage or employment changes are also interpreted in relative terms.

[4]Laboratory occupations in medicine, which are highly automatable but closely substitutable with human medicine, provide a striking example. The latter experiences a large positive demand shock, leading laboratory occupations to shrink as workers transition into human medicine. However, equilibrium wage adjustments prevent significant flows into other closely related health occupations.

Our study relates to a large literature on changes in the occupational structure. In particular, our work complements a long line of research examining demand shocks and their impacts on tasks and occupations in the labour market (e.g., Autor et al., 2003; Autor & Dorn, 2013; Goos et al., 2014; Graetz & Michaels, 2018; Dauth et al., 2021; Gregory et al., 2022). We show that a key determinant of these impacts is the heterogeneous capacity of labour supplies to adjust, which leads to the large variations in employment and wage changes that have been observed during periods of structural change. Recent work by Acemoglu & Restrepo (2022) relates wage changes of different demographic groups to their own realised task displacement as well as the ripple effects of task displacement from other groups.[5] Our work is complementary to Acemoglu & Restrepo as it also takes the differential equilibrium effects of shocks to one occupation on others into account but measures the effect heterogeneity from worker flows ex ante, providing a forward-looking perspective on market adjustments.[6]

This ex ante approach complements recent research identifying labour demand shocks driven by rapid technological advancements (e.g., Webb, 2020; Eloundou et al., 2023; Felten et al., 2023; Brynjolfsson et al., 2025). Studies tracking mentions of such technologies in job postings, including Alekseeva et al. (2021), Acemoglu et al. (2022), and Hampole et al. (2025), provide useful insights into how automation and AI reshape labour demand. By incorporating these predicted demand shifts, our model generates forward-looking projections of their effects on wages, employment, and worker flows. Another key advantage of integrating occupational heterogeneity within a market equilibrium framework is its ability to inform more targeted policy interventions. For instance, it complements research on job search guidance and retraining programs (Belot et al., 2019; Altmann et al., 2023) by identifying occupational transitions that align with workers' previous experience and are in high demand under equilibrium conditions. Moreover, our model pinpoints occupational areas with the lowest elasticities – those least adaptable – where active human capital policies are most needed when shocks arise.[7]

[5]See also Böhm (2020) who infers labour demand shocks from the interaction of ex post task reallocations and wage changes across different skill groups. Bhalotra et al. (2023) use women and men's realised employment and wages across task groups for equilibrium identification of labour supplies.

[6]Several studies use job flows to predict workers' outside options or wage spillovers across employers, including Carlsson et al. (2016); Arnold (2021); Bassier (2024); Schubert et al. (2024); Green et al. (2024). Our approach focuses on occupational structural changes and embeds them within an explicit supply-and-demand framework to jointly analyse employment outcomes and wage adjustments to shocks over time. Research on equilibrium sectoral effects of shocks, primarily from the trade literature, includes Caliendo et al. (2019); Traiberman (2019); Humlum (2021); Adão et al. (2024). While these studies examine the dynamics of adjustment to shocks, they typically impose relatively restrictive assumptions on labour supply elasticities. Here our work is most closely aligned with Bocquet (2024), who develops an explicit network model of occupational job search.

[7]This aspect also relates to studies designing or evaluating policies targeted at specific worker groups, such as wage support programs or trade adjustment assistance (Hyman et al., 2024).

This paper continues as follows. Section 2 presents a partial equilibrium model with perfect information that provides a tractable framework for labour mobility decisions under frictions. In Section 3, we add the demand side to the model to characterise labour market equilibrium. Section 4 discusses the data and describes the components of the estimated own- and cross-price elasticities of occupations' labour supply. In Section 5, we estimate the full equilibrium model using instrumental variables, and discuss robustness of our estimate of the aggregate supply parameter. Section 6 uses the equilibrium model to extract the effects of supply heterogeneity from demand and supply shocks in decomposition and performs counterfactual analyses. Section 7 contains details of our prediction exercises. Section 8 concludes.

2 The Model of Labour Supply

We adopt a random utility model of worker preferences that characterises occupation-specific labour supply functions. This builds on Cortes & Gallipoli (2018) and Hsieh et al. (2019), who adapt the environment in Eaton & Kortum (2002) to occupational choices, and Card et al. (2018) who study the selection of workers into firms. We first present our baseline static model with perfect information, which provides a tractable framework for labour mobility decisions under frictions. We then discuss how this model can be extended to include further features affecting occupational choice. Labour demand and market equilibrium are modelled in the next section.

2.1 Environment and Occupational Choice

There is a continuum of workers $\omega \in \Omega$ and a finite set of N occupations. The number of employers in each occupation is large, such that labour demand is competitive and there is no strategic wage setting. Every worker is initially and predeterminedly assigned to an occupation i. Workers subsequently choose occupations to maximise their utility, which can be interpreted as a total lifetime payoff and is occupation-combination as well as individual specific. It includes wages as pecuniary benefits, a specific cost of switching between occupations i and j, and an idiosyncratic preference for working in occupation j.

The indirect utility of worker ω with initial occupation i choosing occupation j is given by:

$$u_{ij}(\omega) = \theta p_j + a_{ij} + \varepsilon_j(\omega) \quad (1)$$

where θp_j is the general pecuniary payoff to occupation j. The component p_j can be interpreted as the log occupational price or wage rate offered to all workers per unit of their skill and θ as their pecuniary preference or 'wage elasticity' parameter. The occupation–

combination-specific term a_{ij} summarises potential pecuniary and non-pecuniary costs of selecting occupation j for individuals initially assigned to occupation i. These can include lower payoffs as switchers may need to learn new tasks in j or institutional barriers. We will return to these costs further below. The final summand $\varepsilon_j(\omega)$ is an idiosyncratic preference shock for working in occupation j, which may, for example, include non-pecuniary match components with occupation-specific amenities or types of coworkers. We assume $\varepsilon_j(\omega)$ is independently drawn from a type I extreme value (i.e. Gumbel) distribution.[8] Draws, including for the current occupation, occur at the beginning of the period. Based on realised shocks, switching costs, and log occupational prices, workers decide whether to stay in their occupation or switch to a different one.

By standard arguments (McFadden, 1973), the assumptions on eq. (1) imply that workers' occupational choice probabilities take the form:

$$\pi_{ij}(\mathbf{p}) = \frac{\exp(\theta p_j + a_{ij})}{\sum_{k=1}^{N} \exp(\theta p_k + a_{ik})}, \tag{2}$$

where $\mathbf{p}$ is the vector of N log occupational prices. We follow the convention that, by the law of large numbers, π_{ij} is the fraction of workers switching from occupation i to j. Choice probabilities are occupation–combination-specific and they may involve staying in the current occupation ($i = j$). Intuitively, eq. (2) says the more attractive occupation j is relative to all other occupations, and the lower the cost of switching to it from i, the higher is the fraction of workers who will move to that occupation. Since they are aggregated over idiosyncratic shocks, the probabilities are not individual-specific and we therefore omit the index ω from now on.

2.2 Price Elasticities of Occupational Employment

Let τ_i denote the share of the working population originating in occupation i, such that $\sum_i \tau_i = 1$. One can think of $\{\tau_i\}$ as the stationary distribution of employment in a baseline period. Further, let $E_j(\mathbf{p})$ be the fraction ending up working in occupation j as a function of log occupational prices. This implies

$$\begin{aligned} E_j(\mathbf{p}) &= \sum_i \tau_i \pi_{ij}(\mathbf{p}) \\ &= \tau_j \text{ if } \mathbf{p} = \mathbf{p}^* \end{aligned} \tag{3}$$

[8] Gumbel location μ and scale δ are general because equation (1) can always be recast as $u_{ij}(\omega) = \frac{\theta}{\delta} p_j + \frac{a_{ij}}{\delta} + \frac{\varepsilon_j(\omega) - \mu}{\delta}$, yielding the same choice probabilities (see Card et al., 2018). In that sense, θ can be thought of as scaling the importance of wages relative to idiosyncratic shocks.

with $\mathbf{p}^*$ the vector of baseline log occupational prices. From now, we simplify our language by using 'prices' to mean log occupational prices as described in eq. (1).

2.2.1 Own- and Cross-Price Elasticities

Our interest centres on (own- and cross-occupation) price elasticities, that is, the elasticity of occupation j's employment with respect to any occupation k's price (including $k = j$). Writing $e_j \equiv \ln E_j(\mathbf{p})$, and differentiating eq. (3), we obtain:

Remark 1 (Elasticities and Job Flows) *The short-term partial derivative of occupation j's log employment share with respect to k's log price is equal to:*

$$\frac{\partial e_j(\mathbf{p})}{\partial p_k} = \theta d_{jk} \tag{4}$$

with

$$d_{jk} = \begin{cases} \frac{\sum_i \tau_i\left(\pi_{ij}\left(1-\pi_{ij}\right)\right)}{\tau_j} & \text{if } j = k \\ -\frac{\sum_i \tau_i\left(\pi_{ij}\pi_{ik}\right)}{\tau_j} & \text{otherwise} \end{cases} \tag{5}$$

Appendix A.1 contains the derivation.

The remark shows how these price elasticities can be computed using transition probabilities (job flows under the law of large numbers), baseline employment shares (which can also be estimated from the job flows), and an unobserved pecuniary parameter θ. We return to the estimation of θ in Section 5 and now focus our attention on eq. (5).

Element d_{jk} in eq. (5) can be thought of as a constituent of an $N \times N$ matrix capturing the heterogeneity of price elasticities, D, which we also refer to as the 'elasticity matrix' throughout the paper. With a slight abuse of notation, we thus refer to elements d_{jj} and d_{jk} as own- and cross-price elasticities, respectively. One immediate property of these elasticities is that $d_{jj} = -\sum_{k \neq j} d_{jk}$ and that D is of rank $N-1$. This feature does not preclude separate effects in the empirical analysis below, however, as own- and cross-elasticities are weighted by differing price changes across occupations.

We obtain d_{jk} from differentiating the choice probabilities (2) with respect to occupational prices. The economic intuition in (5) is that more observed switches between i and j (high π_{ij}) also indicate more workers at the margin of indifference between choosing the two occupations. In terms of the cross-elasticities, d_{jk} are high when occupations i (including j and k themselves) tend to send many workers to both occupations at the same time. Therefore, employment in j can be thought of as strongly reacting to price changes in k when many workers across the set of occupations $\{i\}$ (i.e. weighted by the size τ_i) are indifferent between j and k. This intuition that more switches indicate a larger share of workers that can be redirected via a price change also holds for other models of occupational choice. In fact, in Section 2.3 we study extensions of our baseline model to show that they yield substantively similar elasticity specifications as in eq. (5).[9] The intuition also naturally extends to own-elasticities, which constitute an aggregation of an occupation's cross-elasticities as we shall see again next.

We gauge the content of eq. (5) further by deriving an alternative formulation in terms of moments of job flows. First, and as standard, let $E_\tau x \equiv \sum \tau_i x_i$ be the average of vector elements x_i weighted by the stationary employment distribution $\{\tau_i\}$. Then define $\tilde{\pi}_{iq} \equiv \frac{\pi_{iq}}{\tau_q}$, such that $\tilde{\pi}_{iq}$ gives normalised job flows, with $E_\tau \tilde{\pi}_{iq} = 1$. Normalising the transition probabilities in this way yields moments that are invariant to occupation size. In parallel, let $Cov_\tau(x, y) \equiv \sum \tau_i (x_i - E_\tau x)(y_i - E_\tau y)$. This leads us to the following result:

Remark 2 (Occupational Substitutabilities) *For all $j \neq k$, the off-diagonal elements of D can be expressed as:*

$$-d_{jk} = \underbrace{\tau_k}_{\substack{\text{occupational}\\ \text{importance}}} \times \underbrace{Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)}_{\substack{\text{occupational}\\ \text{substitutability}}} + \underbrace{\tau_k}_{\substack{\text{price}\\ \text{index}}} \tag{6}$$

where we examine the negative of d_{jk}, rather than d_{jk} itself, so that we can interpret higher elasticities by larger positive numbers. For all $j = k$, the on-diagonal elements of D can be expressed as:

$$d_{jj} = \underbrace{\sum_{k \neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)}_{\substack{\text{aggregate}\\ \text{substitutability}}} + \underbrace{1}_{\text{direct}} - \underbrace{\tau_j}_{\substack{\text{price}\\ \text{index}}} \tag{7}$$

Appendix A.2 contains the derivation.

We start by analysing expression (6). The key difference to a more standard labour supply specification is the heterogeneous occupational substitutability term. In particular, if $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) > 0$ then the occupation pair j and k are 'competing' for workers, and the larger is this covariance of normalised job flows, the higher is the cross-price elasticity (i.e. the responsiveness of employment in occupation j to changes in the price of occupation k). The substitutability term is weighted by an 'occupational importance' term τ_k that depends on the size of the occupation of the price change and which reflects that price increases in a smaller competing occupation will have smaller percentage ripple

[9] The elasticities in Remark 1 are short-term in the sense that, on top of π_{ij}, next period occupation sizes would also change. Section 2.3 discusses a dynamic generalisation of the model that allows for longer-run elasticities. We also show the robustness of our empirical results at different frequencies in Appendix G.1.

effects than price increases in a larger occupation. The last term in expression (6) is an occupation-specific intercept which captures occupation k's contribution to an overall price index. As will be seen further below, the price index effect is already a feature of the more standard homogenous model and quantitatively unimportant in terms of the variability of elasticities across occupations.

Analogous to the cross-elasticities, the key term in expression (7) for j's own-elasticity is its 'aggregate substitutability', which constitutes a size-weighted average of the occupation's pairwise substitutabilities with all other occupations. This captures the fact that a unit increase in the price of occupation j is equivalent to an equal and opposite price decline in all other occupations. Expression (7) also contains 'direct' and price-index effects. These terms contribute to the average level of the elasticities but little to the observed variability.[10]

Before returning to questions of measurement, we note that the occupational substitutabilities highlighted in Remark 2 have two more attractive properties. First, the pairwise substitutabilities are naturally symmetric between j and k, in contrast to the cross-elasticities which depend on the size of the occupation of the price change. It is worth emphasising, however, that these symmetries are entirely consistent with occupational hierarchies.[11] Second, we have formulated $Cov_{\tau}\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ and $\sum_{k \neq j} \tau_k Cov_{\tau}\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ using normalised job flows such that they are also invariant to the fineness of the occupational classification (see Appendix A.2 for details).

Equation (5) gives rise to a direct measurement of occupational elasticities via the observed moments of job flows. One can in principle use the transition probabilities

[10] Appendix A.2 derives an alternative version of eq. (7) as

$$d_{jj} = \underbrace{-\tau_j Var_{\tau}\left(\tilde{\pi}_{.,j}\right)}_{\substack{\text{job-flow} \\ \text{dispersion}}} + \underbrace{1}_{\text{direct}} - \underbrace{\tau_j}_{\substack{\text{price} \\ \text{index}}}$$

Here, the aggregate substitutability is reformulated as a 'job-flow dispersion' term, reflecting how dispersed or concentrated the inflows to occupation j are: Sectors hiring from a diversity of sources (in this case, a *small* $Var_{\tau}\left(\tilde{\pi}_{.,j}\right)$) are more elastic. To see why, note that inflows are typically concentrated if the diagonal element of the transition matrix is close to 1 (meaning everyone remains in the current occupation) and the off-diagonal elements are close to 0. In this case, $Var_{\tau}\left(\tilde{\pi}_{.,j}\right)$ is large, the job-flow dispersion component is more negative, and d_{jj} is lower, indicating a lower own-price elasticity.

[11] To provide an example to illustrate this point, suppose that doctors can become plumbers, but that plumbers can never become doctors. If the wage of plumbers increases then the flow out of medicine increases. If the wage of doctors increases, then the flow into plumbing declines. When these occupations are of equal size, the model predicts that these flow changes will be identical.

to identify features of the occupational transition costs.[12] This approach is developed in studies employing job flows in combination with worker and employer fixed effect wage estimations (Sorkin, 2018; Morchio & Moser, 2024). Since disentangling amenities, or occupational access and transition costs, from wages is not the focus of this paper, we stick with the sufficient statistics approach in our analysis. We do, however, show empirical correlates of the cross-price and own-price elasticities in Section 4, which may be interpreted as capturing aspects of the underlying occupational transition costs.

Finally, one can relax the measurement requirement for using the job flows in steady state equilibrium. In case we do not start out in equilibrium, short-run elasticities at the beginning of our analysis period would still depend on the same measured switching probabilities π_{ij}, but departing from initial employment shares. While τ_i are attractive as they are a function only of steady state flows, the empirical results below do not substantively change when we use those E_i instead ($Corr(E_i, \tau_i) \approx 0.85$) in calculation (5).

2.2.2 Occupational Employment Changes

With these properties of the occupational price elasticities in hand, we now generalise the formulation given in eq. (4). In particular, the response of the vector of employment shares to a change in the vector of prices can be approximated by:

$$\Delta \mathbf{e} \approx \frac{\nabla \mathbf{e}}{\nabla \mathbf{p}} \Delta \mathbf{p} = \theta D \Delta \mathbf{p} \tag{8}$$

where $\Delta \mathbf{e}$ represents the change of the $N \times 1$ vector of log employment shares, $\{e_j\}$, and $\frac{\nabla \mathbf{e}}{\nabla \mathbf{p}}$ the $N \times N$ matrix of partial derivatives $\frac{\partial e_j(\mathbf{p})}{\partial p_k}$ $\forall$ j, k. Given some demand-side shock and ensuing shock to prices, which we discuss below, the change to employment shares can be approximated by eq. (8). This is exact for marginal changes in prices.

Equation (8) shows how the model traces out a supply curve vector, $\mathbf{e}(\mathbf{p})$, of log employment shares. With a view to our empirical analysis, we rewrite the product of elasticity matrix D with the vector of price changes as follows:

$$\Delta e_j \approx \theta \mathbf{d_j} \Delta \mathbf{p}$$

[12] Cross-occupational switching costs can be identified relative to amenities of staying from $\ln\left(\frac{\pi_{ij}}{\pi_{jj}}\right) = a_{ij} - a_{jj}$. One can then identify all the a_{ij} terms up to scale, θ, by normalising $a_{11} = 0$ and using information on wages via the relationship $\ln\left(\frac{\pi_{jj}}{\pi_{11}}\right) = a_{jj} + \theta\left(p_j - p_1\right)$.

$$= \theta\Bigg(\underbrace{d_{jj}\Delta p_j}_{\substack{\text{own-occupation}\\ \text{effect}}} + \underbrace{\sum_{k\neq j} d_{jk}\Delta p_k}_{\substack{\text{total cross-}\\ \text{-occupation effect}}} \Bigg), \tag{9}$$

where $\mathbf{d_j}$ is the jth row of matrix D, and, in the bottom line, we separate the effects of on-diagonal elements in D from those of all off-diagonal elements. To summarise the intuition, the own-occupation effect in eq. (9) represents the part of occupations' employment changes that is due to their own price changing. The total cross-occupation effect captures the effect of heterogeneity in price changes across all other occupations: Intuitively, large price changes in occupations that are very substitutable with j (i.e. $d_{jk} \ll 0$) will have potentially important spillovers on j's employment share. We provide additional formal details in Appendix A.3.

2.3 Generalisations

We conclude this section by discussing key extensions of our baseline labour supply model. In all of these, the moments that characterise the heterogeneity of occupational elasticities are substantively similar to those in eq. (5). We intuitively describe the extensions here and provide formal details in Appendix B.

Dynamic overlapping generations model – In Appendix B.1, we develop a life-cycle extension of our main model. This generalisation includes occupational choice at two career stages (labour market entry of young workers and switching at middle-age) and the associated varying human capital accumulation across occupations. Uncertainty about the next period prevails at the individual (preference shock) as well as aggregate (price changes, which we allow to be persistent) level and young workers make forward-looking decisions accordingly. We derive heterogeneous supply elasticities in early and late career again as a function of young and middle-aged workers' occupational choices.

The dynamic model illustrates how occupational elasticities in response to a price change at time t are composed of the reactions at multiple career stages (relatedly, see the elasticities by skill group discussed in the next paragraph). It also clarifies how longer-run elasticities – at $t+1$ when previously young workers become middle-aged and middle-aged workers retire – would look like. Finally, the dynamic model shows how, given rich data on pre-labour market training or skills, one would estimate the labour supply elasticities also at the beginning of the career.[13] Our main takeaway from this analysis,

[13]In the current paper, we instead focus on supporting our model's mechanisms during the career. In particular, Section 5 reports that θ estimated from changing choice probabilities during the career is similar to that estimated from the changes of total employment overall.

however, is that, under reasonable conditions, the occupational elasticities generated by the dynamic framework are similar to those derived from our benchmark model.

Varying elasticities by skill group – Consistent with the notion of different elasticities between younger and older workers, one can extend the framework to allow for varying mobility rates during the career across population subgroups. We show that in this case, the respective aggregate elasticity (overall employment change in j for a price change in k) is the average of each subgroup's elasticity weighted by the subgroup's prevalence in occupation j. In our empirics, splitting by subgroup yields near identical aggregate elasticities to those computed in the main text because (1) in the case of age, relative elasticities are highly correlated across groups and (2) in the case of educational groups, occupations are fairly segmented such that aggregate elasticities mainly reflect either high- or low-skilled workers' mobility in each occupation. See Appendix B.2 for details.

Non-employment sector(s) – One can also extend the analysis to allow for transitions into and out of different non-employment states during the career. In Appendix B.3, we compute the respective occupational elasticities with non-employment and conceptually show how these can be used in estimation even if price (or utility more generally) changes for non-employment sectors are unavailable. Using our panel data, we find that accounting for unemployment transitions does not change the empirical results of the paper.

Further extensions and interpretations – The model allows for data-driven grouping of occupations that are close to one another in terms of elasticities or job flows. In another unreported extension, we further include location choice and explicitly nest it on top of occupation choice in the model. The resulting augmented regional and occupational labour supply elasticities are again expressible solely as functions of job flows within and across nests. Finally, the baseline choice model, eq. (1)–(2), could be interpreted in a more general manner: We already discussed that θ trades off wages with the dispersion of idiosyncratic shocks in destination occupation j; in fact, this dispersion could be made option–j-specific (derivations in Appendix A.1 go through allowing for θ_j). In terms of parameters a_{ij}, one could recast the model so that they reflect the level of idiosyncratic draws rather than transition costs or, alternatively, search frictions affecting the probability of receiving an offer from potential occupation j when starting in i.

3 Labour Demand and Equilibrium

We proceed to close the model by specifying an explicit theory of occupational labour demand. We characterise the resulting system of equations for prices and quantities and study its reaction to shocks. We provide a discussion of the main issues here leaving mathematical details to Appendix C. The model we present features perfect competition but the results remain unchanged if we specify monopsonistic behaviour with time-constant markdowns. More generally, keeping the demand side relatively standard allows us to focus on the effects that arise solely from heterogeneous labour supplies.

We consider an economy-wide constant elasticity of substitution (CES) production function

$$Y = A \left(\sum_j \beta_j E_j^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{s.t.} \sum \beta_j = 1 \tag{10}$$

where β_j are the factor intensities of different occupation inputs and $\sigma > 0$ is the elasticity of substitution between occupations in production. Parallel to Remark 2, which focused on supply, under competitive markets, labour demand elasticities take the form:

$$\frac{\partial e_j^d}{\partial p_k} = \sigma \begin{cases} -(1-\tau_j) & \text{if } j = k \\ \tau_k & \text{otherwise} \end{cases} \tag{11}$$

In equation (11), own-elasticities of labour demand are negative but attenuated by an occupation's size. The latter is a standard result of producer theory, since substitutability declines with an occupation's market share, and can be interpreted as the demand for occupation k's contribution to an overall price index.[14] The cross-price elasticities are positive and, after occupation size adjustment, constant. These constant elasticities are a key feature of CES aggregation, which could be relaxed, for example, by nesting the production function in line with the nested extension of the labour supply.[15]

The full supply and demand model allows us to characterise the equilibrium as a system of N simultaneous equations:

$$e_j(\mathbf{b}, \mathbf{s}) = e_j^s(\langle p(\mathbf{b}, \mathbf{s}) \rangle, \mathbf{s}) = e_j^d(\langle p(\mathbf{b}, \mathbf{s}) \rangle, \mathbf{b}) \tag{12}$$

[14] As discussed in Remark 2, price index terms are part of the elasticities on the labour supply side, too. Anticipating the counterfactual analyses of Section 6, when we impose homogenous occupational substitutabilities this leads to $Cov_\tau(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}) = 0 \; \forall j, k$ and then:

$$\frac{\partial e_j^s}{\partial p_k} = \theta \begin{cases} 1-\tau_j & \text{if } j = k \\ -\tau_k & \text{otherwise} \end{cases}$$

That is, what remains in the homogenous labour supply model are also only the price index terms due to potentially varying occupation sizes.

[15] Berger et al. (2022) and Lamadon et al. (2022) also nest their models into labour markets defined by combinations of industry and region. Within market, they then impose perfect substitutability of firms' outputs. The equivalent $\sigma \to \infty$ here would lead to β_j-shocks fully compensated by commensurate wage increases (see eq. (13) below) and supply shocks fully feeding through to employment (eq. (14)). Our estimates of the supply-side parameter θ are not much affected even by very large σ.

where **b** is the vector of relative productivities (i.e. demand shifters $\left(\ln \frac{\beta_i}{1-\beta_i}\right)$), **s** is a vector of supply shifters, that, intuitively-speaking, move supply curves vertically in parallel. The occupations are indexed by j as before, and both supply (s) and demand (d) curves depend on the full set of prices.

Our focus is on how this system responds to shocks to the structural parameters, given by changes to $\left(\ln \frac{\beta_j}{1-\beta_j}\right)$ and **s**.[16] We obtain:

Remark 3 (Equilibrium Response to Shocks) *The response of the vectors of prices and employment shares to a change in the vectors of supply and demand shocks can be approximated by:*

$$\Delta \mathbf{p} \approx V\Delta \mathbf{b} - \frac{1}{\sigma} V \Delta \mathbf{s} \tag{13}$$

and

$$\Delta \mathbf{e} \approx \theta D V \Delta \mathbf{b} + V \Delta \mathbf{s} \tag{14}$$

where $V = \left(\frac{\theta}{\sigma} D + I\right)^{-1} (I - W)$ *and* W *is the matrix of stacked occupation sizes with* j,k*th element* τ_k*. Appendix C.1 contains the derivation.*

Equations (13) and (14) mirror expressions from a standard model with homogeneous supply elasticities: given the structure of D and V, positive demand shocks increase both prices and employment, while supply shocks increase employment but reduce prices.

Given matrix V's central role in the solution of the equilibrium model, it is worth discussing some of its properties here. In terms of its mathematical features, it has rank $N-1$, just like matrix D, and each row sums to 0 across columns. Additionally, just like matrix D, it has non-negative eigenvalues, which ensure, roughly speaking, that shocks move prices and employment in the expected direction.

In terms of economic properties, first note that both V and D govern the dissipation of shocks across the economy. In our model of the labour supply curve, V can be interpreted as an upstream, and closely related, matrix to D which specifies the initial spillovers from demand shocks to prices. One way to summarise V's effect is to examine its diagonal elements: The next section will report that in the data these are almost perfectly negatively correlated with those of D, such that the diagonal elements of V tend to be *lower* for more

[16]Potential underlying sources of shocks to supply, Δs_j, could be level shifts in the attractiveness of the occupation (a_j component of a_{ij} changing) or exogenous shocks to employment sizes in eq. (3) directly. In the data, the first would lead to higher switching rates toward j from all occupations within the labour market, while the latter may, for example, be an exogenous migration shock that affects j differentially. We keep this reduced-form and allow for both interpretations of supply shocks in the following discussion.

elastic occupations. As implied by eq. (13), for these elastic occupations, *ceteris paribus*, demand shocks induce relatively muted changes to prices.[17]

The matrix product DV then governs the effect of demand shocks on employment. This similarly has rank $N-1$ with all non-negative eigenvalues. Importantly, and as expected, its diagonal elements are *positively* correlated with those of D and *negatively* with those of V. Accordingly, while *ceteris paribus* demand shocks cause a smaller change in prices for more elastic occupations, they induce a *larger* increase in employment implied by eq. (14).[18] Parallel effects to those just discussed can be traced through shocks to supply.

Finally, we generalise eq. (8) to obtain our equilibrium regression equation for the labour supply side:

Remark 4 (Second-Stage Labour Supply Equation) *The equilibrium relationship between the vectors of price and employment changes on the labour supply side can be approximated by:*

$$\Delta\mathbf{e} \approx \theta D \Delta\mathbf{p} + \Delta\mathbf{s} \tag{15}$$

where the $N \times 1$ vector of labour supply shocks $\Delta\mathbf{s}$ represents the regression error. The result follows directly from combining eq. (13) *and eq.* (14)*, see also Appendix E.3.*

Remarks 3 and 4 characterise equilibrium outcomes when own- and cross-occupational labour supplies are heterogeneous. As discussed, these relationships are tractable and intuitive.[19] The remarks further indicate how the model can be used to empirically analyse the role of heterogeneity on the labour supply side. For estimation, we will instrument $\Delta\mathbf{p}$ in eq. (15) by proxies for demand shocks that are interacted, model-consistently, with their impact on prices according to spillover matrix V. In the absence of supply shocks (i.e. $\Delta\mathbf{s} = 0$), OLS is sufficient. The logic of requiring the IV is that supply shocks contribute to, and so are correlated with, $\Delta\mathbf{p}$. For counterfactuals, we will insert different, more homogenous, versions of D (and accordingly, V) into eq. (13)–(14).

[17]We also study the off-diagonal elements of V, which govern the spillovers of demand shocks to prices in other occupations. In contrast to D, many of these off-diagonal elements are positive. Intuitively, a positive shock to demand can create a relative scarcity in labour not only in the given sector but also in close substitute (cross-elastic) occupations. As indicated by eq. (13), this scarcity can then lead to an increase in prices in *both* occupations.

[18]In terms of off-diagonal elements of matrix DV, which govern the spillovers of demand shocks to quantities in other occupations, these are all negative in the data. Following through the example just given in footnote 17, a positive demand shock has two opposing effects on close substitute occupations: First, as discussed above, a possible increase in prices draws workers in from the rest of the labour market; second, however, is the direct effect of the shock which pulls workers in from these close occupations to the occupation of the positive shock itself. Overall, the second effect dominates and, as given by eq. (14), this cross-effect always reduces employment.

[19]They are also modular, in the sense that different elasticity matrices D (with according upstream V) can be inserted, following the discussion in Section 2.3.

4 Data and Descriptives

This section first presents our data sources for the historical analysis. We then describe the estimated own- and cross-price elasticities as well as show their correlations with occupational outcomes during 1985–2010.

4.1 Main Data Sources and Historical Sample

We use the Sample of Integrated Labour Market Biographies (SIAB, Frodermann et al., 2021), a 2% sample of social security records in Germany dating back to 1975. The SIAB contains workers' complete employment histories and daily wages. It is representative of all individuals covered by the social security system, roughly 80% of the German workforce. The SIAB's panel structure allows us to compute careers and job flows over long periods, while its administrative nature ensures the high quality of all variables.

Occupations in the SIAB are consistently coded, with no removals or additions, during 1975–2010. After this, a structural break occurs and a new occupational classification is introduced, which is fully consistent again from 2012 onward. We now use the historical data until 2010 for our analysis of occupational outcomes over the past decades. In Section 7, we will employ the data from 2012 onward to make predictions about the potential impacts of labour demand shocks moving forward.

We condition the historical analysis on men aged 25–59 who are working full-time in West Germany. The first restriction is primarily due to the old occupation classification, which was devised with male employment in mind (Paulus et al., 2013). The restriction to full-time work allows us to use consistent wage and employment samples. In the later analysis, post-2012, we will relax all these restrictions. We transform the daily spell structure of the SIAB into a yearly panel by using the longest spell in a given year. Our historical sample consists of approximately 600,000 unique individuals and 9 million individual $\times$ year observations for the whole period 1975–2010.[20]

We use the SIAB to compute worker flows (sufficient statistics for the elements of D), changes in occupational employment ($\Delta\mathbf{e}$), and changes in occupational prices ($\Delta\mathbf{p}$). For the latter, we follow the literature on this, which emphasises that raw wages need to be corrected for changing composition of workers' skills (Cavaglia & Etheridge, 2020; Böhm et al., 2024), and use occupation stayers' (i.e. workers who do not switch occupation from one year to the next) wage growth as the main estimate of changes in occupational

[20] We further drop spells of workers with missing information on occupation or wage, and wages below the limit for which social security contributions have to be paid. In preparing the data, we impute censored wages above the upper earnings threshold for social security contributions (Dustmann et al., 2009; Card et al., 2013) and correct for the wage break in 1983–1984 (Fitzenberger, 1999; Dustmann et al., 2009).

prices.[21] The SIAB also allows computing further occupational characteristics – e.g. workers' mean age or share of workers with university degrees by occupation – which we relate to our elasticity measures.

To obtain task information in occupations, we use the Qualifications and Career Surveys (QCS, Hall et al., 2012). The QCS consist of cross-sectional surveys with 20,000–35,000 individuals in each wave. Respondents report on the tasks performed in their occupations, which we categorise into analytical, routine, and manual tasks. Averaging responses from pooled QCS data in 1979 and 1985/1986, we compute initial task intensities among those three categories by occupation. Finally, to obtain measures of occupational licensing, we use the indicators for standardised certification and degree of regulation developed by Vicari (2014). More details on the data, variable construction, and descriptive statistics are presented in Appendix D.

4.2 Estimated Elasticities and Spillovers

Our occupational own- and cross-price elasticities are computed from baseline worker flows according to eq. (5). Specifically, we use the transition rates across all occupation pairs at the endpoints of five-year periods within 1975–1984. The flow of switchers from origin occupation i to destination occupation j (which includes staying in occupation i) is defined as the number of individuals who are employed in occupation i in year t and employed in occupation j in year $t+5$. Dividing each element by total flows from origin occupation i we obtain the transition probability matrix Π, which is of size 120×120, with element π_{ij} representing the empirical probability that a worker employed in origin occupation i switches to j in five years' time. The transition probability matrix also implies a steady state vector τ of size 120×1, with element τ_i representing occupation i's size as a share of total employment. Using these objects, we compute the matrices D and V following Remarks 1 and 3, respectively.[22]

Table 1 lists occupations at different quantiles of the elasticity distribution. Panel A shows own-price elasticities, which range from 0.07 among Physicians and pharmacists to 0.80 among personnel in various social, medical, and hospitality service occupations. While own-price elasticities (d_{jj}) are more or less symmetrically distributed over their range, cross-price elasticities ($-d_{jk}$) between occupation pairs are strongly skewed and highest among closely related occupations. Panel B of Table 1 shows such relative spill-

[21] In Appendix G.2, we show the robustness of our results using an alternative price estimation that corrects for worker–occupation–spell fixed effects (Cortes, 2016).

[22] For V, we calibrate $\frac{\theta}{\sigma} = 2.3$, which is justified in Section 5. Even large perturbations in $\frac{\theta}{\sigma}$ have little effect on results. Our findings also remain consistent whether we use two-year or ten-year period lengths for the flows. The resulting analysis period 1985–2010 is similar to Card et al. (2013) and Böhm et al. (2024).

overs of wage on employment changes. These are around 0.18–0.46 for Home wardens and social work teachers on Nursery teachers and child nurses, Non-medical practitioners on Medical receptionists, and Office specialists on Stenographers and data typists. Cross-price elasticities fall off quickly from the top and become an order of magnitude smaller than any own-price elasticities even at the 90th percentile.

Figure 1 shows important correlates of the elasticities. In fact, we plot substitutabilities, which are the key varying components of d_{jj} and d_{jk},[23] against observable occupational characteristics. Panel a reports that occupations with a higher degree share, more analytical tasks, and higher certification and regulation requirements are less substitutable with other occupations and, by extension, less own-price elastic. Aggregate substitutability can thus be directly related to other proxies for occupational flexibility (also in contrast to transition costs, which are defined among occupation pairs and not naturally aggregated).

Panel b of Figure 1 plots individual substitutabilities against occupational task distance, which we constructed as in Gathmann & Schönberg (2010); Cortes & Gallipoli (2018). Pairwise substitutabilities are a natural counterpart to the equally symmetric task distances, again in contrast to a_{ij} or d_{jk}. The relationship is significantly negative, such that the higher is the distance in task content between two occupations, the lower the substitutability. While this clear relationship is reassuring, there are a couple of advantages of working with our elasticity components: task distance is essentially an ordinal concept derived from the subset of tasks that are reported in surveys. In contrast, substitutabilities, or cross-price elasticities, capture all information implied by realised worker flows and they have a natural quantitative interpretation. This can be seen in the noted skewness of Figure 1b, and accordingly by the better fit of Spearman's rank coefficient with task distance than that of standard linear correlation.

We show further summary statistics for matrices Π, D, and V in Appendix Table E.3. As discussed in Section 3, the main diagonal elements of upstream matrix V are negatively correlated with those of D as, for own-price elastic occupations, stronger employment changes go in hand with more muted price changes.[24] Also in contrast to D, the off-diagonal elements of V (spillovers from demand shocks to other occupations' prices) are not particularly skewed, and a subset of elements are positive because demand shocks for a given sector can create scarcity of labour in close substitute occupations.

[23] Appendix Table E.4 reports that (aggregate) substitutabilities are the main factors driving the heterogeneity in own- and cross-price elasticities. As such, both panels of Figure 1 look substantively similar using elasticities instead (see Appendix Figure E.1).

[24] Diagonals of the product DV are still related positively with D (negatively with V), since fundamental demand shocks overall have a larger employment impact in more elastic occupations (see eq. (14)).

Table 1: Summary Statistics: Own- and Cross-Price Elasticities

Panel A	Own-price elasticity (d_{jj})	Occupation
Minimum	0.074	Physicians, dentists, veterinary surgeons, pharmacists
10th percentile	0.294	Health or property insurance specialist
25th percentile	0.358	Members of parliament, association leaders, officials
50th percentile	0.430	Stucco workers, plasterers, rough casters, proofers
75th percentile	0.517	Sheet metal pressers, drawers, stampers, metal moulders
90th percentile	0.604	Salespersons
Third highest	0.740	Ancillary hospitality workers
Second highest	0.797	Medical receptionists
Maximum	0.798	Nursery teachers, child nurses
Panel B	Cross-price elasticity ($-d_{jk}$)	Occupation of price change (k) $\rightarrow$ Occupation of employment change (j)
50th percentile	0.001	Paviours, road makers $\rightarrow$ Sheet metal workers
90th percentile	0.009	Miners, shaped brick/concrete block makers $\rightarrow$ Engine fitters
Fifth highest	0.144	Bricklayers, concrete workers $\rightarrow$ Carpenters, scaffolders
Fourth highest	0.182	Restaurant, inn, bar keepers, hotel and catering personnel $\rightarrow$ Ancillary hospitality workers
Third highest	0.185	Office specialists $\rightarrow$ Stenographers, shorthand typists, data typists
Second highest	0.253	Non-medical practitioners, masseurs, physiotherapists $\rightarrow$ Medical receptionists
Maximum	0.464	Home wardens, social work teachers $\rightarrow$ Nursery teachers, child nurses

Notes: Panel A shows statistics from a ranking of the 120 occupations of the 1988 *Klassifikation der Berufe* according to their own-price elasticity (d_{jj}). Panel B comes from a ranking of the 14280 occupation pairs according to their cross-price elasticity. See text for more details.

Figure 1: Elasticity Components: Comparison with External Metrics

(a) Own-Price Component and Occupational Characteristics

(b) Cross-Price Component and Task Distance

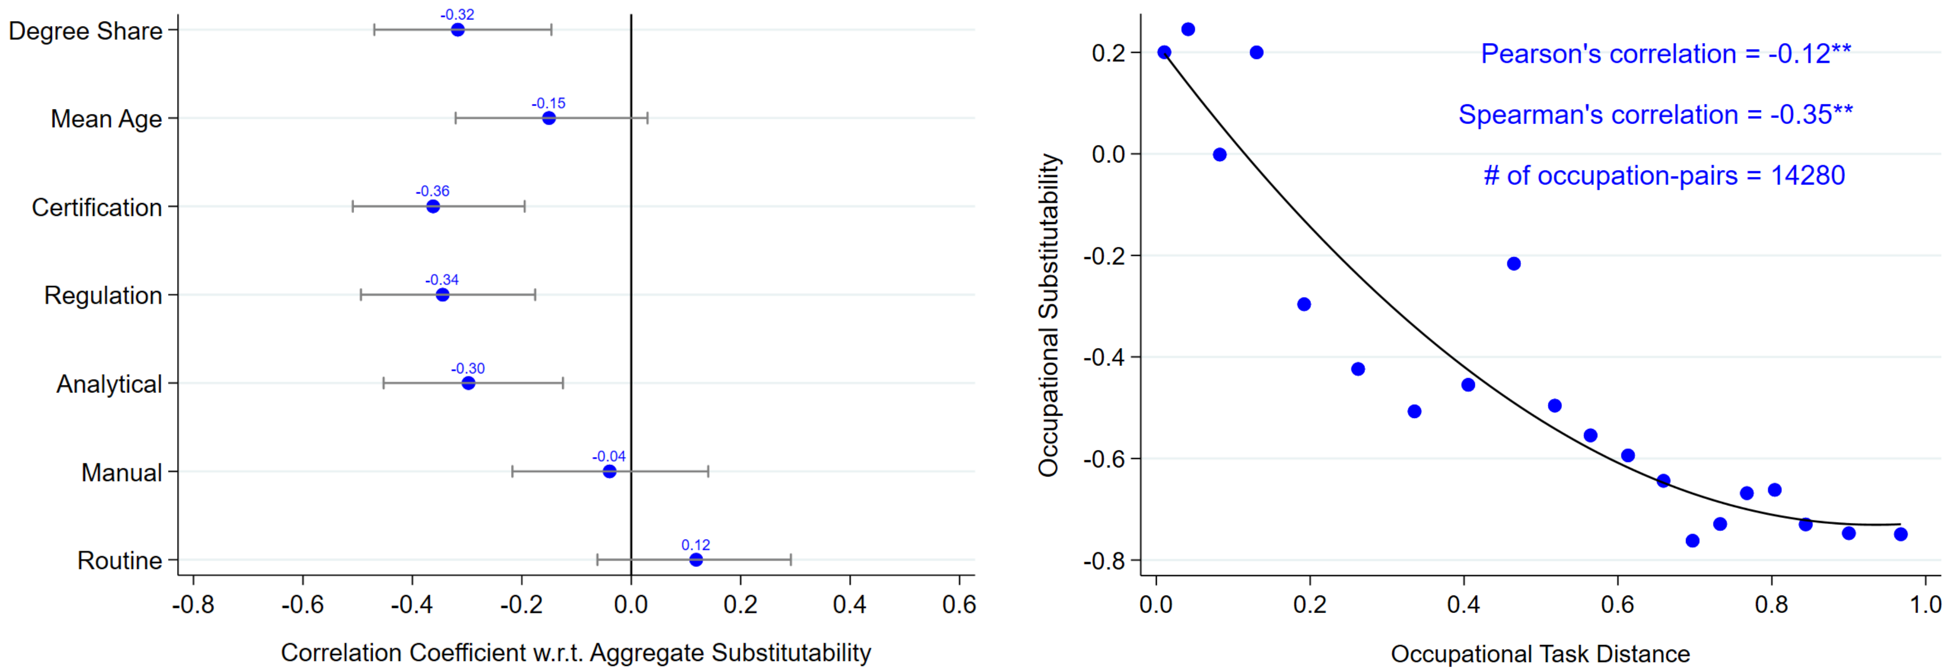

Notes: Panel (a) reports how the aggregate substitutability component of the own-price elasticity, namely $\sum_{k \neq j} \tau_k Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$, correlates with skill requirements across 120 occupations. Occupational certification and regulations come from Vicari (2014). Task content (analytical, manual, and routine) are measured using BiBB, see Appendix D.2. Correlations weighted by initial employment in each occupation. Panel (b) shows the relationship (with a quadratic fit) between the occupational substitutability component of the cross-price elasticity, namely $Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$, and occupational task distance measured as in Cortes & Gallipoli (2018). Appendix Figure E.1 does the same plots for d_{jj} and $-d_{jk}$ instead.

Finally, Table E.3 reports a high autocorrelation of matrix D (and accordingly V) when computed at different points in time during 1975–2010. This stability is empirically in line with findings of Gathmann & Schönberg (2010) on the autocorrelation of occupational task distance over time, and it supports our treatment of the elements d_{jj} and d_{jk} essentially as fixed.[25] In fact, we checked that recomputing elasticities based on flows in the final ten years of the period (2000–2010), and redoing all our analyses with this, does not change the empirical results of the paper.

5 Estimation in the Historical Sample

In this section, we first discuss occupational wage and employment trends during 1985–2010 through the lens of our model. We then use two alternative and complementary approaches to estimate its remaining aggregate parameter on the labour supply side.

5.1 Occupational Changes over 1985–2010

Figure 2a plots annualised occupational changes in employment against changes in occupational prices based on stayers' wage growth. The latter clearly lines up with employment growth, consistent with earlier work (Cavaglia & Etheridge, 2020; Böhm et al., 2024). Yet, there is a significant amount of variation in the movements of employment and wages across occupations. For example, the explicitly labelled occupation of Physicians and pharmacists has high occupational wage growth (over five log points per year) but rather small employment growth, while Assistants exhibit high employment but hardly any wage growth. Data processors have both substantial employment and wage growth.

An implication of our model is that such heterogeneity can be due to differences in labour supply curves across occupations. To graphically assess this possibility, we first consider individual price changes in isolation and split occupations at the median of own-price elasticities in Figure 2b. That is, we concentrate on the first summand (own-occupation effect) in equation (9). The blue circles, including Physicians and pharmacists, are the occupations for which employment is *ex ante* predicted to be relatively unresponsive to own-price changes, while the red circles, including Assistants, are predicted to

[25] In Section 5, we provide evidence for the mechanism underlying eq. (5), by showing that a price increase in occupation j raises π_{ij}, *ceteris paribus*. The intuition for why this effect need not substantively change the according elements of matrix D is that elasticities depend on overall bilateral (gross) flows among occupations. These do not unambiguously rise or fall in response to price changes and in contrast to net flows. For example, from eq. (5) we can write $d_{jj} = \pi_{jj}(1-\pi_{jj}) + \sum_{i \neq j} \frac{\tau_i}{\tau_j} \pi_{ij}(1-\pi_{ij})$. The first summand decreases when p_j rises, since $\pi_{jj} > 0.5$ in all occupations while the second set of summands increases, since $\pi_{ij} \ll 0.5$ if $i \neq j$. Therefore the impact on overall bilateral flows for j, and thereby d_{jj}, is ambiguous although net flows into the occupation clearly increase.

Figure 2: Occupational Price and Employment Changes (1985–2010)

(a) All 120 Occupations **(b)** By Own-Price Elasticity (Median Split Illustration)

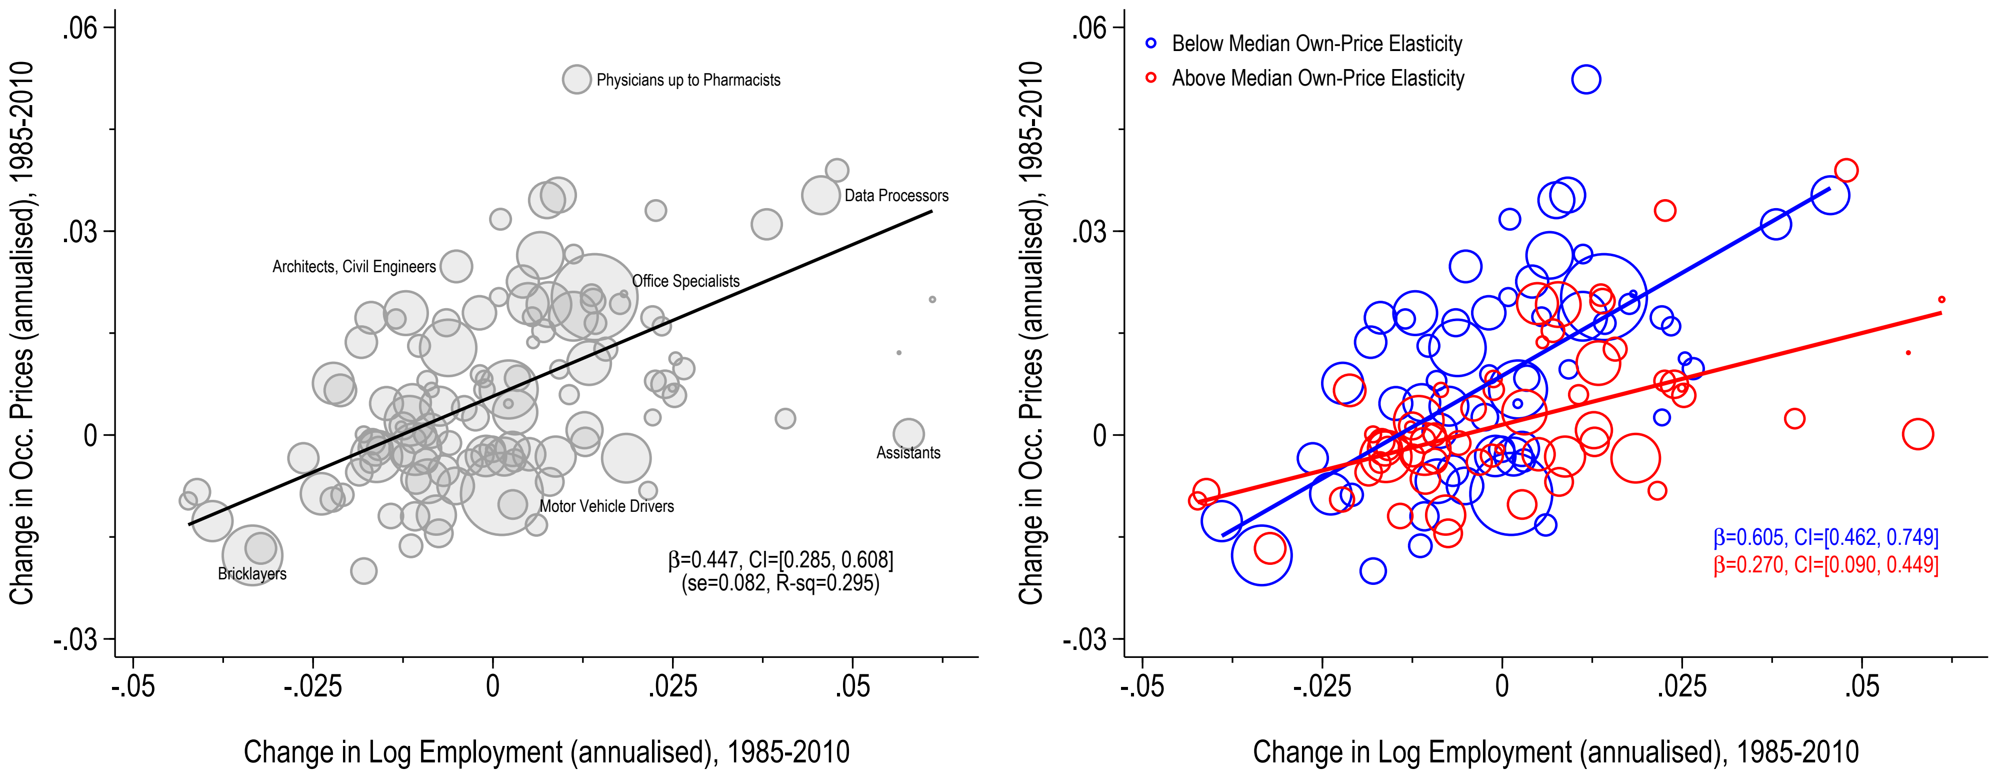

Notes: Panel (a) shows the line from an occupation-size weighted regression of price change on employment change. Panel (b) shows a split by occupations below (blue, inelastic) and above (red, elastic) the median own-price elasticity (d_{jj}). β refers to the slope coefficient, *CI* stands for the 95% confidence interval, *se* refers to standard error, and *R-sq* stands for the R-squared of the regression. Marker size indicates the baseline employment (in 1985) in each occupation.

be relatively elastic. It turns out that, indeed, the relationship between occupational employment and price changes is significantly flatter among the red than among the blue circles. As shown in the plot labels, a 1% increase in wages is on average associated with a $\frac{1}{0.270} \approx 3.7\%$ increase of employment for the group with high predicted elasticities, but only a $\frac{1}{0.605} \approx 1.7\%$ employment increase for the low-elasticity occupations.[26]

The labour supply side relationship described in equation (9) includes additional effects arising from the cross-price elasticities. In order to assess these, we extend the decomposition as follows:

$$\Delta e_j = \alpha + \underbrace{\theta_1 \bar{d}_{diag} \Delta p_j}_{\substack{\text{fixed relationship of} \\ \text{price with employment}}} + \underbrace{\theta_2 (d_{jj} - \bar{d}_{diag}) \Delta p_j}_{\substack{\text{heterogeneity of} \\ \text{own-occupation effect}}} + \underbrace{\theta_3 \sum_{k \neq j} d_{jk} \Delta p_k}_{\substack{\text{total cross-} \\ \text{-occupation effect}}} + \varepsilon_j \tag{16}$$

Equation (16) includes own- and cross-occupation effects but, parallel in spirit to Figure 2, it further splits the own-occupation effect into a fixed relationship that one would obtain when regressing employment onto price changes and the additional effect of the pure heterogeneity in elasticities. This is done by subtracting the mean of matrix D's main diagonal elements in the middle summand, such that the heterogeneity is captured by $d_{jj} - \bar{d}_{diag}$. Using (16) as a regression equation, we also allow for generic coefficients on these effects in order to assess the theoretical prediction that $\theta_1 = \theta_2 = \theta_3 = \theta$. Finally, while the theory analysed a model of employment shares (employment levels in a static population), intercept α now accounts for overall changes in log employment, and the approximation error from eq. (9) is represented by ε_j.

Table 2 reports results from estimating different versions of equation (16). Column (1) shows the regression of Δe_j onto $\bar{d}_{diag} \Delta p_j$ only. As seen in Figure 2a, this fixed relationship of employment with price changes results in a significant positive slope with an R-squared of 0.29. Column (2), which allows for heterogeneity in own-price elasticities d_{jj}, yields an additional positive and significant effect, consistent with Figure 2b. Column (3) then adds the cross effects of price changes in other occupations that may be more or less substitutable. The coefficient on this term is also positive and significant. This is as expected because $d_{jk} < 0$ for $k \neq j$, such that a positive regression coefficient implies that rising prices in other occupations k lead to a decline of employment in occupation j. As noted above, a stronger implication of the model is that coefficients θ_1–θ_3 should all capture the same pecuniary preference parameter. Although econometrically they are allowed to differ, the estimated coefficients turn out to be almost identical. We examine

[26] Appendix Figure E.2 alternatively splits occupations into d_{jj} quartiles. The resulting four regression lines are visibly ranked by predicted labour supply elasticity.

the equality of coefficients more formally in columns (4) and (5). Consistent with $\theta_1 = \theta_2 = \theta_3$ being fulfilled, results hardly change when we run accordingly restricted models.

Table 2: Determinants of Employment Changes: Own- and Cross-Effects (OLS)

		Dependent Variable: Δe_j				
		Unrestricted Model			Restricted Model	
		(1)	(2)	(3)	(4)	(5)
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	1.59 (0.30)	1.79 (0.31)	4.09 (0.89)	1.81 (0.32)	4.15 (0.70)
heterogeneous own effect:	$(d_{jj} - \bar{d}_{diag})\Delta p_j$		1.25 (0.36)	4.07 (1.00)		
total cross effect:	$\sum_{k \neq j} d_{jk}\Delta p_k$			4.02 (1.33)		
R-squared		0.295	0.314	0.394	0.310	0.394
Number of occupations		120	120	120	120	120

Notes: The table presents the estimates from different versions of eq. (16). Regressor in column (4) is $d_{jj}\Delta p_j$. In column (5), the regressor is $\sum_k d_{jk}\Delta p_k$, i.e. including both own- and cross-occupation effects. All regressions include a constant. Observations weighted by occupation j's initial employment size. Period 1985–2010. Standard errors in parentheses; all coefficients shown are significant at the 1% level.

The estimated coefficients in columns (3) and (5) are all substantially larger than those in the other columns (Wald test p-value < 0.01). The reason for this is that, at least over this period, highly cross-elastic occupations tend to experience similar price changes: for $-d_{jk}$ large, Δp_j and Δp_k tend to move together, implying that $Cov(\Delta p_j, d_{jk}\Delta p_k) < 0$. Adding this up for all $k \neq j$, the own-occupation and total cross effects are negatively correlated, and including the latter raises the coefficient on the former in our estimation.[27] Borusyak et al. (2023) find a related result in migration regressions across Brazilian region-industries. They highlight that pecuniary parameters are underestimated when not taking into account that shocks are often correlated between workers' current and potentially substitutable migration options. We will find below that fundamental demand shocks also tend to be correlated among more cross-elastic occupations, which lead to a lower effective labour supply elasticity and subdued adjustment to structural changes.[28]

[27] For clarity, we are here taking the elasticities in D as given and considering the covariance over random draws of price changes. See Appendix E.2 for further discussion on this omitted variable bias.

[28] The large increase of R-squared in Table 2 when adding cross-occupation effects (8 compared to 2 percentage points when adding heterogeneous own effects) already indicates the eventual importance of spillovers between occupations in the full model.

5.2 Identification of the Labour Supply Parameter θ

Recall the model's equilibrium relationships captured in Remarks 3 and 4. Now suppose we have access to a variable, which we denote by r_j, that proxies demand shifters $\Delta \ln \frac{\beta_j}{1-\beta_j}$ but is uncorrelated with supply shifters Δs_j. Eq. (13) implies proportionality of the form

$$D\Delta \mathbf{p} \sim DV\mathbf{r} = D\left(\frac{\theta}{\sigma}D + I\right)^{-1} \check{\mathbf{r}} \tag{17}$$

where vector $\check{\mathbf{r}} \equiv (I - W)\,\mathbf{r}$ is the weighted-demeaned version of $\mathbf{r}$. Equation (17) represents an IV first-stage relationship for the relevant regressor, the product of elasticities with price changes. We discuss setting the value of $\frac{\theta}{\sigma}$ when this is used further below and first focus our attention on the key shock vector $\mathbf{r}$.

Our instrument for relative productivity shocks is based on initial task content. As discussed in Section 4, we employ survey information that asks workers which tasks they carry out in their jobs to construct measures of analytical, routine, and manual task intensity across occupations in the late 1970s and early 1980s. Following the literature on routine-biased technical change (RBTC, Autor et al., 2003), research has consistently found that occupations intensive in analytical tasks grew quite strongly, whereas employment in routine-intensive occupations declined in the late 1980s and the 1990s (e.g. Autor et al., 2008; Acemoglu & Autor, 2011). For Germany, Böhm et al. (2024) additionally show that the overall demand shift was negative for manual-intensive occupations; with employment, average wages, and skill prices declining after 1985.[29] We thus approximate occupation j's (negative) demand shocks during 1985–2010 as

$$r_j = (routine_j + manual_j) - analytical_j \tag{18}$$

The idea is that occupations initially scoring high on routine and manual relative to analytical tasks will tend to experience negative demand shocks during the sample period compared to occupations that score low on our measure r_j.

5.2.1 Estimate from the Model Equilibrium Relationship

Following the exposition in the previous section, we illustrate the IV estimation by first implementing the model abstracting from cross-occupation effects. As in Figure 2b, we capture the heterogeneity of own-occupation effects by splitting estimation at the median of d_{jj}. In this case, the instrument $DV\mathbf{r}$ within each sub-sample of 60 occupations reduces

[29] Böhm et al. (2024) also caution that the QCS questionnaires have some difficulty distinguishing between routine and manual job tasks. See Rohrbach-Schmidt & Tiemann (2013) for further details about classifying tasks in the German context.

to a scalar multiple of the pure proxy vector $\mathbf{r}$. Figure 3a displays the relationship between r_j and Δp_j. Overall, it is clearly negative given the negative demand shocks that we proxy. We would also expect the regression line to be flatter among more elastic occupations, which should react to a demand shock relatively less in terms of wages and more in terms of employment. Although not significant at conventional levels, this difference is apparent. Similarly, Appendix Figure E.3 displays the relationship between r_j and Δe_j, and consistently shows that the more elastic occupations present a slightly steeper slope.

Figure 3b then depicts the second stage in this simplified model. Again to parallel the prior figure, and to keep prices on the vertical axis as standard, we display the *inverse* supply curve, with price changes as a function of changes to employment. In this case, the slopes are steeper than those in Figure 2b. This reflects that, in this case, removing shocks to supply also eliminates attenuation of the estimated regression line. What remains the same is that the relationship of wages with employment is substantially steeper among occupations *ex ante* classified as inelastic compared to elastic occupations. These are the relative reactions in terms of employment for a given price change among more versus less elastic occupations. Figure 3 is therefore illustrative of the type of variation employed in our instrumental variables approach.

Implementing the full model requires having some information on the relative elasticities $\frac{\theta}{\sigma}$ in the first stage eq. (17). We choose a calibration based on estimates from the literature. Given a range of $\sigma \in [1.81, 2.10]$ from Burstein et al. (2019) and initial information on the potential value of θ from Table 2, we calibrate $\frac{\theta}{\sigma} = 2.3$ (Appendix Table E.7 shows robustness of results to a wide range of values for $\frac{\theta}{\sigma}$). Then, estimating eq. (15) yields:

$$\begin{aligned} \Delta e_j &= \underset{(1.30)}{4.78}\ \mathbf{d}_j \Delta \mathbf{p} + constant + error_j \\ \mathbf{d}_j \Delta \mathbf{p} &= \underset{(0.0125)}{-0.046}\ \mathbf{d}_j V \mathbf{r} + constant + error_j \end{aligned} \tag{19}$$

In contrast to the illustration of the IV shown in Figure 3b, the theoretical model here is specified again in terms of the standard (rather than inverse) supply curve. The first stage relationship of occupations' task intensities on price changes, multiplied by elasticities $\mathbf{d}_j$ reflecting their implied impact on employment, is negative as expected and displays an F-statistic of 13.5. In the second stage, the estimated pecuniary preference parameter is $\theta = 4.78$ and statistically significant.[30] Before discussing how this number is consistent with findings in the literature, we turn to an alternative approach for estimating θ.

[30] One may also note that it is larger than what would be implied from the OLS regressions in Table 2. This should be so because, if price changes are correlated with supply shocks, OLS will be attenuated and biased downwards. We discuss why the IV is still not too different from the OLS in Appendix E.3.2.

Figure 3: Instrumental Variables Illustration: Median-Split by Own-Price Elasticity

(a) IV Reduced-Form for Wages **(b)** IV Second-Stage: Inverse Supply Curve

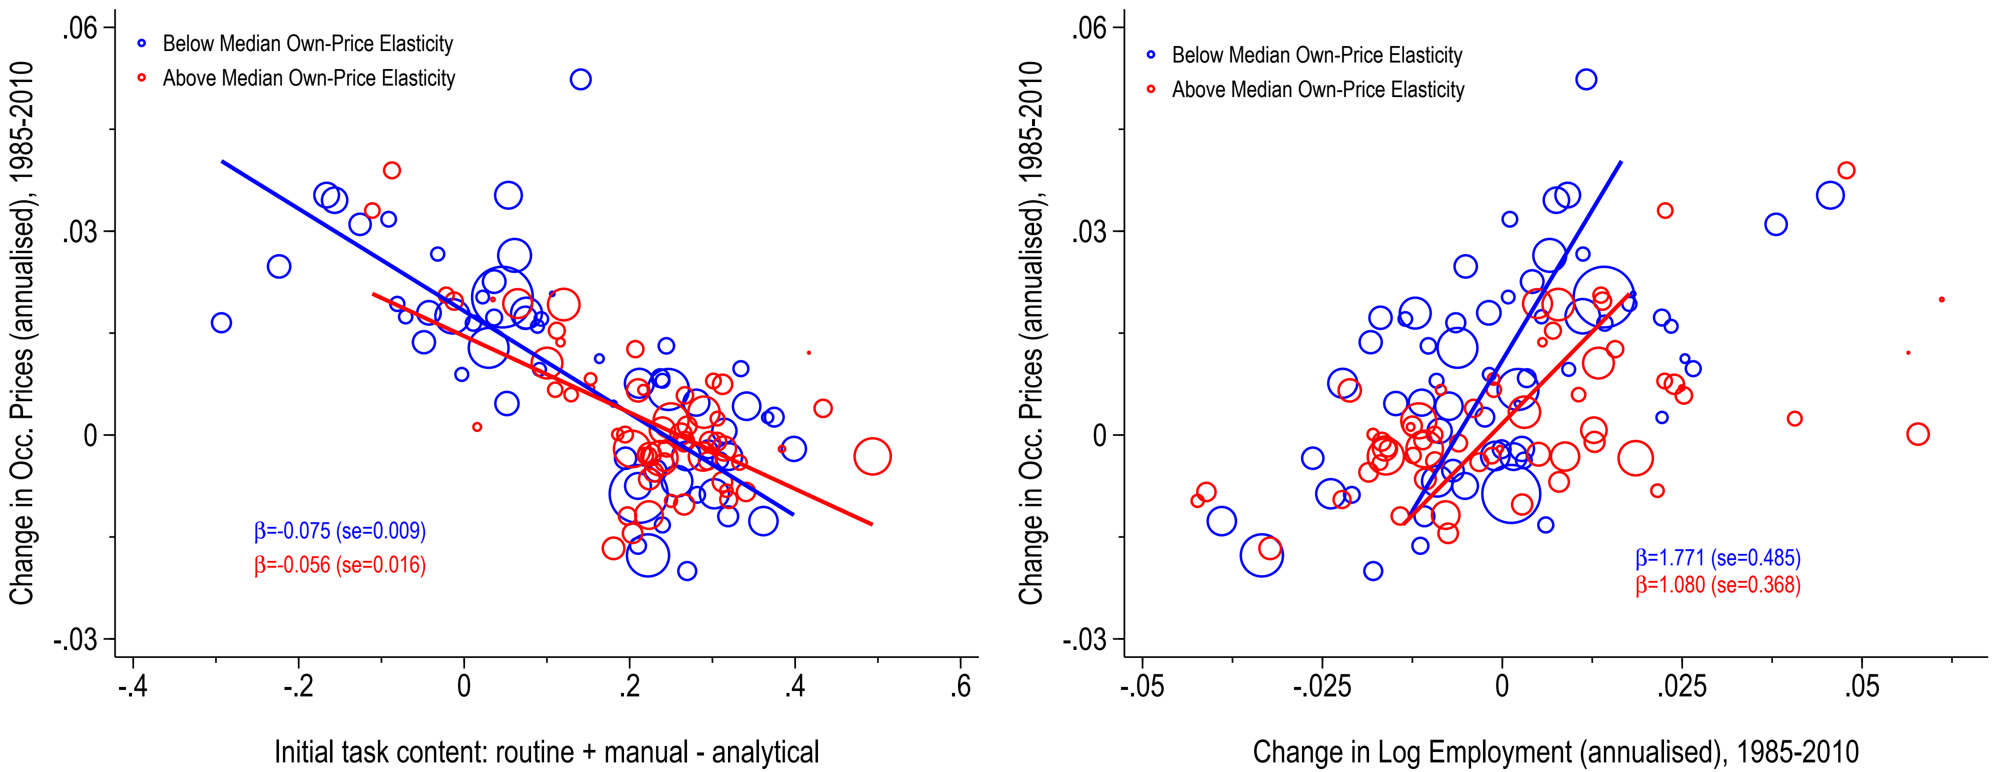

Notes: Panel (a) shows reduced-form regressions of occupations' price changes on their initial task contents r_j. Panel (b) shows second-stage IV-2SLS regressions of occupations' price changes on predicted employment changes using initial task contents as the instrument. Colour codes and linear regression lines are split by occupations below (blue, inelastic) and above (red, elastic) the median own-price elasticity (d_{jj}). β and *se* refer to the slope coefficient and standard error, respectively. Marker size indicates the baseline employment (in 1985) in each occupation.

5.2.2 Alternative Estimate from Changing Choice Probabilities

As a second approach, we implement an identification strategy based on changing choice probabilities over time. In particular, we can rewrite the switching probabilities from equation (2) in logs and relative to the staying probabilities as $\log\left(\frac{\pi_{ij}}{\pi_{ii}}\right) = \theta p_j + a_{ij} - (\theta p_i + a_{ii})$. Differencing this over time yields:

$$\triangle \log\left(\frac{\pi_{ij}}{\pi_{ii}}\right) = \theta\left(\triangle p_j - \triangle p_i\right) + \varepsilon_{ij}, \tag{20}$$

where $\triangle p_j$ is the change in occupation j's log price from 1985 to 2010 that we have used before and $\triangle \log\left(\frac{\pi_{ij}}{\pi_{ii}}\right)$ the according change in the relative choice probabilities. There are two interpretations of the error term in eq. (20): If this reflects measurement errors in the choice probabilities (e.g. due to sampling), OLS regression is consistent. If the ε_{ij} partly also reflect changes in the underlying relative switching costs, and thereby potential endogeneity, we need an instrument for the price changes.

Before continuing, note that eq. (20) is not only a tool to estimating θ but it also reflects the mechanism underlying the labour supply elasticities in Remark 1. That is, occupations whose prices increase, will see higher probabilities of workers switching into them compared to other occupations. The effects will be proportional in relative terms, but stronger in absolute terms in occupations where baseline π_{ij} are large. This is why in Section 2.2 we discussed the relevance of the share of workers who are at the margin of switching.

In order to estimate eq. (20), we again approximate occupation j's demand shocks using their initial task content according to eq. (18). We measure initial choice probabilities during the early period 1975–1985, while eventual probabilities are constructed from the period 2000–2010. Two-stage least squares estimation based on this yields:[31]

$$\begin{aligned} \triangle \log\left(\frac{\pi_{ij}}{\pi_{ii}}\right) &= \underset{(2.54)}{4.03}\ \triangle(p_j - p_i) + constant + error_j \\ \triangle(p_j - p_i) &= \underset{(0.003)}{-0.064}\ (r_j - r_i) + constant + error_j, \end{aligned} \tag{21}$$

where the first stage is strong and the second-stage coefficient $\theta = 4.03$ for the pecuniary preference parameter turns out similar to our main estimate of 4.78. At the same time, the results in eq. (21) support our model's mechanism that workers' relative choices will

[31] The regression has $N \times (N-1) = 14,280$ observations and we cluster standard errors at the level of 120 source occupations i. We replace zeros in π_{ij} with their minimum values following the literature on estimating gravity equations (Head & Mayer, 2013; Cortes & Gallipoli, 2018). When constructing five-yearly changes of transition probabilities and estimating eq. (21) in a panel, results are quantitatively similar and statistically stronger. For details, including results on alternative OLS estimation or with different treatment of zeros, see Appendix E.3.3.

change in response to relative occupational price changes.

Finally, it is worth noting that our values of θ from either estimation approach are also broadly comparable to estimates in the related literature. Cortes & Gallipoli (2018) estimate θ using US wage dispersion and obtain values in the range of 2 to 8.87. As another comparison, the literature on employer wage effects finds that the elasticity of labour supply to the firm is around 2–7 (e.g. see Lamadon et al., 2022, and papers cited therein). Given that switching occupations is likely more costly than switching firms, it seems plausible that our implied own-elasticities fall into the lower end of this range (average $\theta d_{jj} = 2.1$ as $\overline{d}_{diag} = 0.43$). The novelty of our approach, however, lies in the heterogeneity around the average for own-price (from $0.07 \cdot 4.78 = 0.3$ to $0.80 \cdot 4.78 = 3.8$) as well as cross-price elasticities (from essentially 0 to 2.2).

6 Model-Based Decomposition and Counterfactuals

The previous section determined the remaining aggregate paramaters of the full supply and demand model. We now use this to decompose the historical changes in employment and wages into contributions of different factors: shocks to occupational demand, supply, and the heterogeneities in labour supply elasticities that we emphasise.

6.1 Construction of Counterfactuals

We use equations (13)–(14) to express the changes of prices and employment in terms of parameters and exogenous shocks as follows:[32]

$$\Delta \mathbf{p} = \left(\frac{\theta}{\sigma} D + I\right)^{-1} \Delta \mathbf{b} - \frac{1}{\sigma}\left(\frac{\theta}{\sigma} D + I\right)^{-1} \Delta \mathbf{s} \tag{22}$$

$$\Delta \mathbf{e} = \theta D \left(\frac{\theta}{\sigma} D + I\right)^{-1} \Delta \mathbf{b} + \left(\frac{\theta}{\sigma} D + I\right)^{-1} \Delta \mathbf{s} \tag{23}$$

The equilibrium solution treats equations (22) and (23) as equalities and – up to constants representing general wage and employment growth – reproduces the actual changes of $\Delta\mathbf{p}$ and $\Delta\mathbf{e}$ from the data. We manipulate these reduced-form expressions to study the role of labour supply heterogeneity versus occupation-specific shocks for the variation in wages and employment. We provide a summary here leaving details to Appendix F.2.

[32]Equations (22) and (23) are obtained by inserting the solution for V into equations (13)–(14) and then using the fact that demand and supply shocks are weighted mean zero by construction (see Appendix F.1). We use our baseline parameter estimate $\theta = 4.8$ and accordingly $\sigma = 4.8/2.3 = 2.1$. Appendix Table E.7 shows the robustness of θ and results reported below are also similar for the range of σ values in that table.

To study this, we replace D with its matrix equivalents from counterfactual environments with more homogeneous elasticities. Our first counterfactual, matrix D_{own}, considers the case that occupations' aggregate (own-price) elasticities vary but their substitutabilities with other occupations are homogeneous. For example, employment in service occupations may be responsive to price but we suppose that flows of workers into services come equally from any other occupation according to its size. This is consistent with theoretical models often found in the literature on firms (e.g. Card et al., 2018; Lamadon et al., 2022; Berger et al., 2022), where the costs of entering employer j do not depend on the source employer i (that is, $a_{ij} = a_j$ in eq. (1)). Main diagonal elements of D_{own} continue to be the actual own-price elasticities, whereas cross-price elasticities reduce to appropriate fractions of the on-diagonals. We term this the model with 'heterogeneous own-price supply' or, alternatively depending on the context, 'with homogeneous spillovers'.

Another counterfactual imposes completely homogeneous labour supply elasticities. The main diagonal elements of matrix D_{hom} become an average $\bar{d}_{diag}$ and cross-price elasticities a constant fraction of it. This counterfactual is consistent with specifications in the empirical literature that regress occupations' log employment changes on their log wage changes (e.g. Autor et al., 2008; Dustmann et al., 2009; Cavaglia & Etheridge, 2020; Böhm et al., 2024, or column (1) of Table 2). From eq. (15), it leads to a relationship of the form $\Delta e_j = constant + slope \cdot \Delta p_j$, where the slope is proportional to pecuniary preferences θ and the constant is proportional to the average wage growth in the economy. We term this the 'fully homogeneous' model.

As an alternative to the counterfactual D-matrices, we turn off the classic simultaneous equations component. We do this by shutting down supply shocks, using $\Delta \mathbf{s}_{off} = \mathbf{0}$ in equations (22)–(23), which allows us to assess the variation in wages and employment that these shocks account for.

6.2 Results

We begin with a decomposition to uncover the drivers of overall employment reallocation. We do this by following eq. (23) and running regressions of observed employment changes on various components of the right-hand side. The first row in Table 3 shows that demand shocks in the fully homogeneous model (i.e. $\Delta \mathbf{s}$ replaced by $\Delta \mathbf{s}_{off} = 0$ and D replaced by D_{hom}) explain 64% of the variance of employment changes. This is consistent with the literature on job polarisation (e.g. Acemoglu & Autor, 2011; Goos et al., 2014), where demand shocks are the main drivers of occupational changes. But it still leaves room for a substantial role of supply.

Table 3: Decomposition of Overall Employment Changes

	(1) R-sq. between data & model	(2) Remainder explained
Base $\Delta \mathbf{e}$ with $\Delta \mathbf{b}$	0.641	
Adding supply shocks	0.858	60.4%
Adding supply heterogeneity	0.849	57.9%
Full model	1.000	

Notes: This table decomposes the employment changes in our 120 occupations. The first row considers only demand shocks in the fully homogeneous model. The second row adds supply shocks. The third row alternatively adds supply heterogeneity. The final row considers the full model. Column (1) reports the regression R-squared between the data and model. Column (2) gives the percent of remainder explained by either the counterfactual with only supply shocks or with only heterogeneity.

The second row of Table 3 adds supply shocks, still under D_{hom}, to create a new counterfactual employment change according to eq. (23) in the homogeneous model. This explains 86% of the observed employment changes in an R-squared sense, or roughly half of the remaining variance in $\Delta \mathbf{e}$. Similarly, adding heterogeneity of supply under $\Delta \mathbf{s}_{off} = 0$, and using full matrix D with the demand shocks in eq. (23), accounts for 85% of employment changes and again roughly half of the remaining variance.[33] Together, supply shocks and heterogeneity, by construction, explain the full variation in actual employment changes (last row of Table 3).[34] They are thus both important, in addition to demand shocks and equally so, to account for the overall occupational employment changes observed over the past decades.

We provide further insights into this result by constructing proper counterfactuals. Figure 4 displays results where, in keeping with the figures throughout the paper, we relate implied counterfactual price changes $\Delta \mathbf{p}_{cf}$, generated by eq. (22), to implied counterfactual employment changes $\Delta \mathbf{e}_{cf}$ from eq. (23) across different scenarios. We start again with demand shocks in the fully homogeneous model. In this case, all occupational changes emanating from $\Delta \mathbf{b}$ run perfectly along a single supply curve (panel a). We can see from this plot that the explicitly labelled Physicians and pharmacists as well as Data processors are among the occupations with the largest relative demand increases over

[33] It is worth noting why the R-squared is higher when we regress employment changes on demand shocks than when we regress on observed price changes in Table 2. Intuitively, the error terms related to supply shocks in eq. (23), given by $V\Delta \mathbf{s}$, are substantially less dispersed than those in the OLS estimation of eq. (15), $\Delta \mathbf{s}$. In the latter, they are also negatively correlated with the regressor, $\Delta \mathbf{p}$, due to simultaneity, lowering the estimated contribution of prices.

[34] The two standalone contributions sum to $60.4\% + 57.9\% > 100\%$, which implies a -18% interaction effect. This is because eq. (23) is not purely additive.

time. Bricklayers are among the occupations with the largest negative demand shocks.[35]

Panel b shows how supply shocks affect this counterfactual. Here we facilitate interpretation by retaining the regression line from panel a. Switching $\Delta\mathbf{s}$ back on introduces attenuating variation around the price-employment relationship such that the R-squared in a regression of price on employment declines to 64%. The regression line moves *clock-*

Figure 4: Counterfactual Changes of Prices and Employment

(a) Demand shocks only; Fully-homog. supply

(b) ... with supply shocks

(c) ... with supply shocks; Het. own-price supply

(d) Both shocks; Fully heterogeneous supply

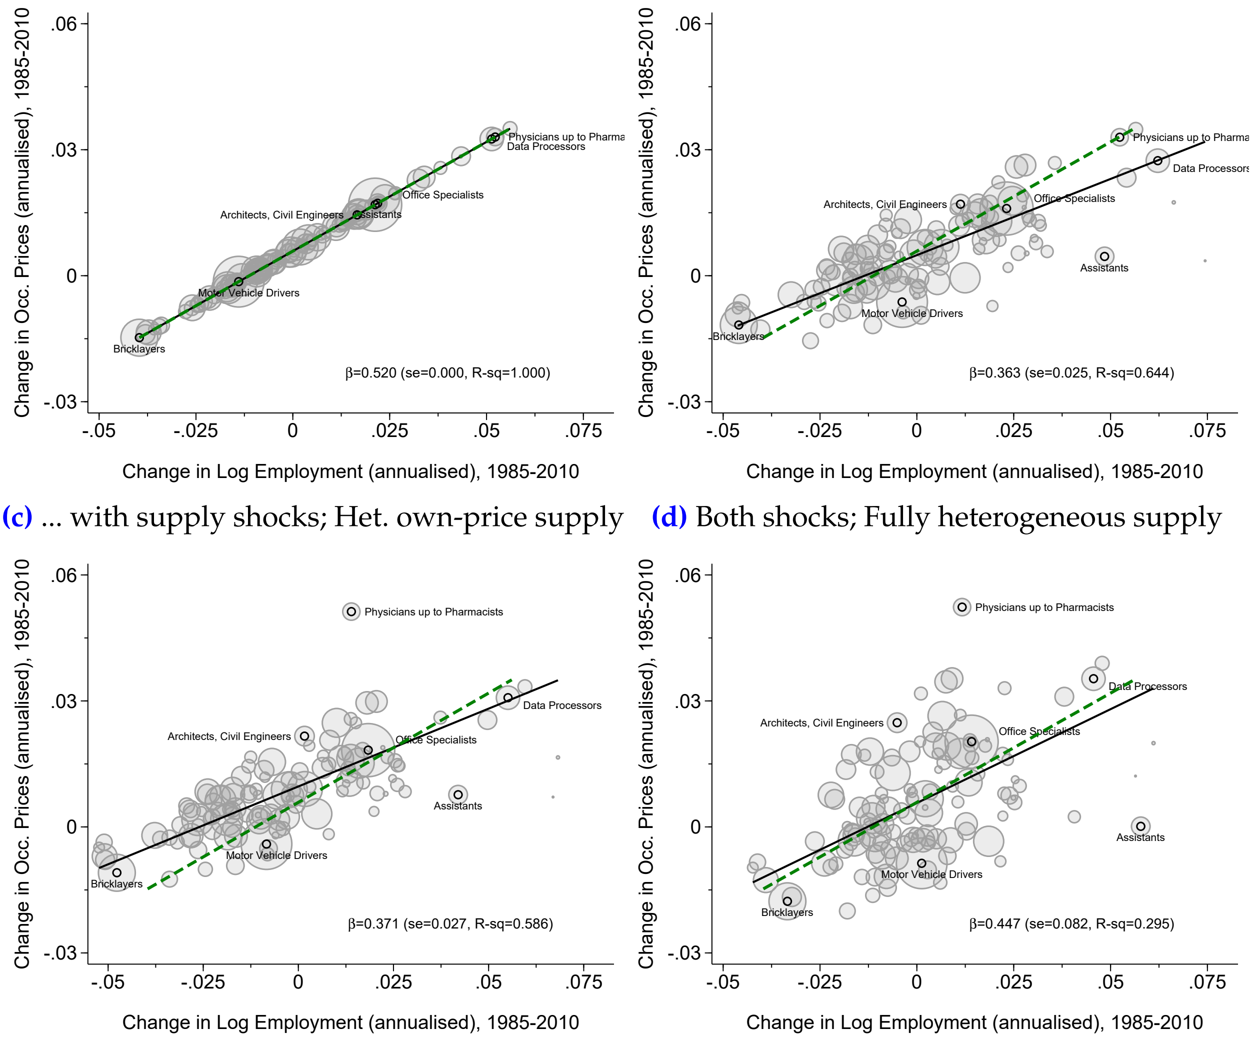

Notes: The figure shows occupational price and employment changes for different manipulations of the elasticity matrix D and $\Delta\mathbf{s}$, as described in Section 6.1. In Panel 4a, both supply shocks and heterogeneity in D are switched off (i.e. $\Delta s = \Delta s_{off} = 0$ and D_{hom}), leaving only demand shocks. Panel 4b first introduces the supply shocks (i.e. $\Delta s \neq 0$), then 4c adds own-elasticity heterogeneity (i.e. D_{own}). Finally, Panel 4d shows the full model (actual data) by including also heterogeneous cross-elasticities (i.e. full matrix D is used). For the exact description of the counterfactuals see Section 6.1 and Appendix F.2. The OLS with slope coefficients, standard errors, and R-squared is shown for each panel. For ease of comparison, the regression line in Panel 4a is repeated as green-dashed in all panels. Marker size indicates the baseline employment (in 1985) in each occupation.

[35] It is worth noting here that the points in this plot include average real price and employment growth, both of which are positive over the period. Accordingly, occupations with no relative demand shock are located slightly above and to the right of the origin.

wise and its slope reduces from 0.52 to 0.36, partly driven by positive demand *and* supply shocks in occupations such as Assistants. Still, the regression slope remains strongly positive, which is due to the larger dispersion of demand shocks than of shocks to supply.

The remaining two panels of Figure 4 show how the movements of occupational prices and employment are affected by labour supply heterogeneity. Panel c first introduces heterogeneity of occupations' own-price supply, but retains homogeneous spillovers (i.e. uses matrix D_{own} as discussed above). A geometric interpretation of the transition from panel b to c is that each occupational point is translated along its own demand curve and according to its own aggregate labour supply elasticity. Inelastic occupations move *counterclockwise* around the centre: in a Northwest direction for those with positive demand shocks, Southeast for those with negative demand shocks, and with no effective change for those with no shock to demand. Symmetrically, occupations that are more elastic than average move around the centre *clockwise*.

Panel c shows that the effect of allowing for this heterogeneity is, for the most part, small. This is consistent with the OLS regressions of Table 2. A strong exception is for Physicians and pharmacists, which is very own-price inelastic (see again Table 1) and experienced a large positive demand shock. This makes its implied price increase much higher, and its employment increase lower, compared to panel b (or compared to, say, Data processors, who exhibit an own-price elasticity of roughly average strength). Finally, panel d also includes heterogeneity in cross-occupation elasticities, and so reproduces the observed data. Compared to panel c, variation around the regression line increases substantially, such that the R-squared from a regression of price on employment reduces from 59% to 30%. As an illustration of this feature, displayed occupations such as Assistants and Motor vehicle drivers both move away from the regression line but in different directions. Motor vehicle drivers experienced a negative demand shift and, since demand shocks were also negative in close substitute occupations,[36] the impact was seen in strong wage declines, while drivers' employment share declined very little. On the other hand, Assistants experienced a positive demand shift, but the negative demand shocks in close substitute occupations made working in this occupation relatively even more attractive, serving to amplify the employment response.

Overall, the locus of points moves on average *counterclockwise* and the slope of the regression line increases from 0.37 to 0.45. These changes show the importance of allowing for heterogeneous spillovers to explain the data. In effect, the realised employment responses to demand shocks captured by the full matrix D are smaller, and wage responses

[36] Highly cross-elastic occupations to Motor vehicle driver include, among others, Railway engine drivers and Transportation equipment drivers as well as Stowers and furniture packers. For the latter, their employment share actually increased over time.

are larger, than those captured by matrix D_{own} or the fully homogeneous model. This is because shocks tend to be correlated among cross-elastic occupations: The experience of Motor vehicle drivers was more typical than the experience of Assistants. As seen in Figure 4d, the resulting lower effective labour supply responses also led to larger between-occupation wage inequality than in a model without cross-occupation effects.[37]

We finish this discussion by starting again at panel a and, from there, providing another assessment of the relative importance to dispersion in price and employment changes of supply shocks versus supply heterogeneity. Using changes in R-squared again as a metric, we see that the contributions are roughly equal. Panels a and b show that supply shocks cause a decline in the R-squared of 36 percentage points, while b and d show that supply heterogeneity accounts also for a decline of 35 points. Therefore, and consistent with Table 3, the relative impacts of supply heterogeneity and shocks are similar in explaining occupational changes. Moreover, we have discussed that, within this overall important contribution, different aspects of heterogeneity were important for explaining idiosyncratic outcomes of particular occupations.

7 Future Projections

An important question is how the labour market will be affected by technology shocks of the *future*. While our model is silent on what exactly the shocks will be, it can help make relevant predictions about their potential heterogeneous equilibrium impacts. We illustrate this by showing how projected enhanced automation shocks would affect employment, wages, and labour market flows from 2022 onward. We summarise the main analyses here; details on the data construction and further results are found in Appendix H.

We use the prediction sample already introduced in Section 4.1. This dataset has an updated and consistent occupation classification, which captures the trends toward more employment and job differentiation in areas such as health, education, and information technologies. Parallel to before, we construct the elasticity matrix for the new 126 occupations based on worker flows during 2012–2021.[38] We then supplement these data with expert assessments in 2022 about what share of tasks in each occupation could in

[37]The impact of heterogeneous spillovers is seen even more starkly in Appendix Figure F.1 where we introduce heterogeneous elasticities *before* introducing supply shocks. Without the background dispersion from these, the increase in the regression slope is highly obvious. We also display the impact of demand and supply shocks along the occupational wage distribution in Appendix Figure F.3–F.4. Among other things, these show that the lower effective labour supply responses led to even larger between-occupation wage inequality. For Figure 4, although changing the sequence with which we re-introduce model features makes them more or less salient graphically, it does not change their quantitative importance markedly.

[38]We now include Eastern Germans, women, and part-time workers in the sample. We keep our aggregate parameters $\theta = 4.8$ and $\sigma = 2.1$.

principle be replaced with current frontier technologies. We scale that share to have mean zero and the same standard deviation as the demand shocks backed-out over 1985–2010. That is, our measure of an occupation's predicted demand shock is based on its potential automatability relative to other occupations, cardinally interpreted and scaled to reflect the size of shocks in the past.

Parallel to before, panel a of Figure 5 shows these pure shocks to labour demand in the counterfactual model with fully homogenous labour supplies.[39] We see that Driving and sports instructors or Human medicine are among the occupations experiencing the most positive relative demand shocks while Industrial glass makers suffer the most negative shock. Some rather high-skilled occupations, including Accounting and Laboratory occupations in medicine, are also projected to endure relatively strong automation in the foreseeable future. Occupations such as IT consulting and sales or Machine building and operating locate within this range, where the former experiences a modestly positive and the latter a comparably negative demand shock.

Given the actual heterogeneity of labour supplies, the full model's projection is that equilibrium wage and employment changes will be substantially different from the pure ranking of shocks, however. Panel b of Figure 5 shows that occupations which are quite own-inelastic, such as Human medicine, experience relatively large wage changes whereas

Figure 5: Price and Employment Changes due to Predicted Automation Shocks

(a) ... with fully homogenous supply **(b)** ... with fully heterogeneous supply

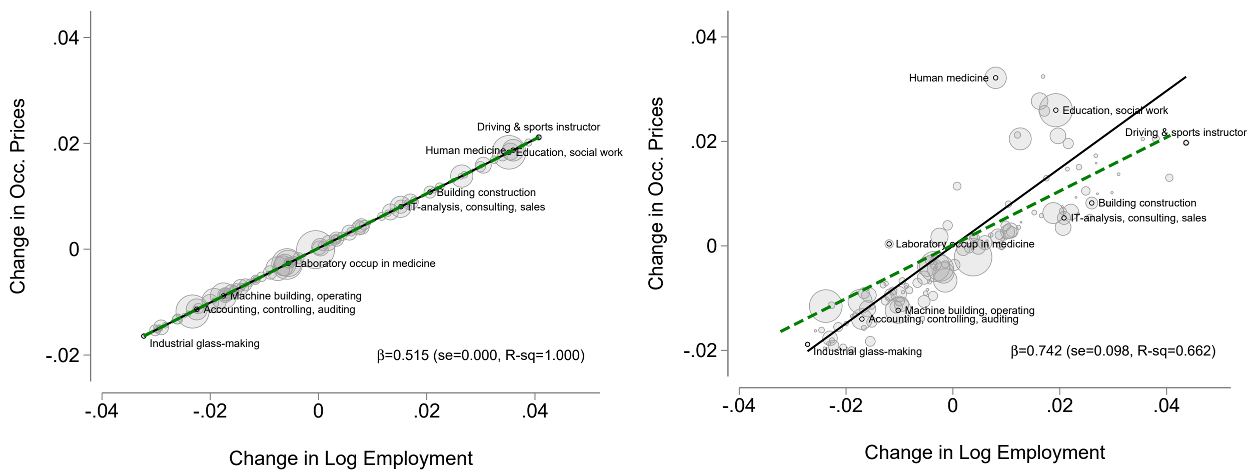

Notes: The figure shows predicted occupational price and employment changes given the elasticity matrix D and projected demand shocks in 2022. For details of how these objects were constructed, see the text. Panel a shows the (homogenous) impact of the shocks when heterogeneity in D is switched off (i.e. D_{hom}, constructed parallel to Section 6.1). Panel b shows the full predicted impact with heterogeneous occupational labour supplies (i.e. full matrix D is used).

[39] In terms of eq. (22)–(23), Figure 5a depicts outcomes when the elasticity matrix is set to D_{hom} and $\Delta \mathbf{s} = \mathbf{0}$. Of course, enhanced automation may constitute only a subset of all demand shocks $\Delta \mathbf{b}$ after 2022.

more own-elastic occupations, such as IT consulting and sales, see relatively strong employment but less wage growth. Importantly, and as before, cross-occupation spillovers play a large role. For example, while both Building construction and Machine operating are located around the average in terms of own-price elasticities, cross-occupation effects on employment are positive for the former, making it effectively more elastic, while they are negative for the latter. These cross effects can be most clearly seen in Laboratory medicine, whose employment strongly declines as a substantial number of its employment shifts toward substitutable Human medicine, raising equilibrium wages for those workers that remain in the occupation. Overall, the cross effects again steepen the slope of the regression line, thereby raising the dispersion of occupational wage changes while reducing the extent to which employment adjusts to automation demand shocks.

A key feature of these projected outcomes is the equilibrium price changes that result from the demand shocks and supply heterogeneity. Equation (13) in Remark 3 characterises these price changes as the interaction of elements of the spillover matrix V with the vector of demand shocks in all occupations $\Delta\mathbf{b}$. Table H.2 in the Appendix reports these pairwise 'price pressures' from nearby occupations on a selected set of occupations. It shows, for example, that demand increases in Human medicine have strong (positive) price pressures on Laboratory medicine while price pressures on IT consulting from its nearby occupations are much smaller.

With the price changes in hand, another key and policy-relevant feature is the bilateral flows of workers between occupation pairs that are projected in equilibrium. In particular, equation (15) in Remark 4 allows characterising the predicted net employment flow from any occupation k to any j as the product of their relative equilibrium price change with the according cross-elasticity.[40] Figure 6 illustrates this by displaying the main bilateral flows for IT consulting and sales. As seen before, this occupation is forecast to experience a moderately positive demand shock but projected to grow quite strongly in terms of employment. The growth occurs because IT consulting is able to draw in workers from a range of related substitutable occupations whose equilibrium price changes are smaller than its own. In particular, the flow chart in Figure 6 shows that workers from more management-oriented (e.g. Business organisation, Purchasing and sales) as well as more technical (Electrical engineering, Computer science) occupations can be attracted toward IT consulting once its relative wages rise. In contrast, occupations that experience even larger wage increases, especially in the health and education fields, are not able to draw many individuals away from IT consulting, since they are not sufficiently cross-elastic with it from a worker perspective. As a result, IT consulting is projected to grow by over

[40] Appendix H details how to compute these flows in terms of levels of employment shares. It also shows the formal decomposition of overall price changes into the underlying individual price pressures.

twenty percent as a share of total employment over the next ten years.

Figure 6: Projected Net In and Out Flows of the Occupation IT Consulting & Sales (%)

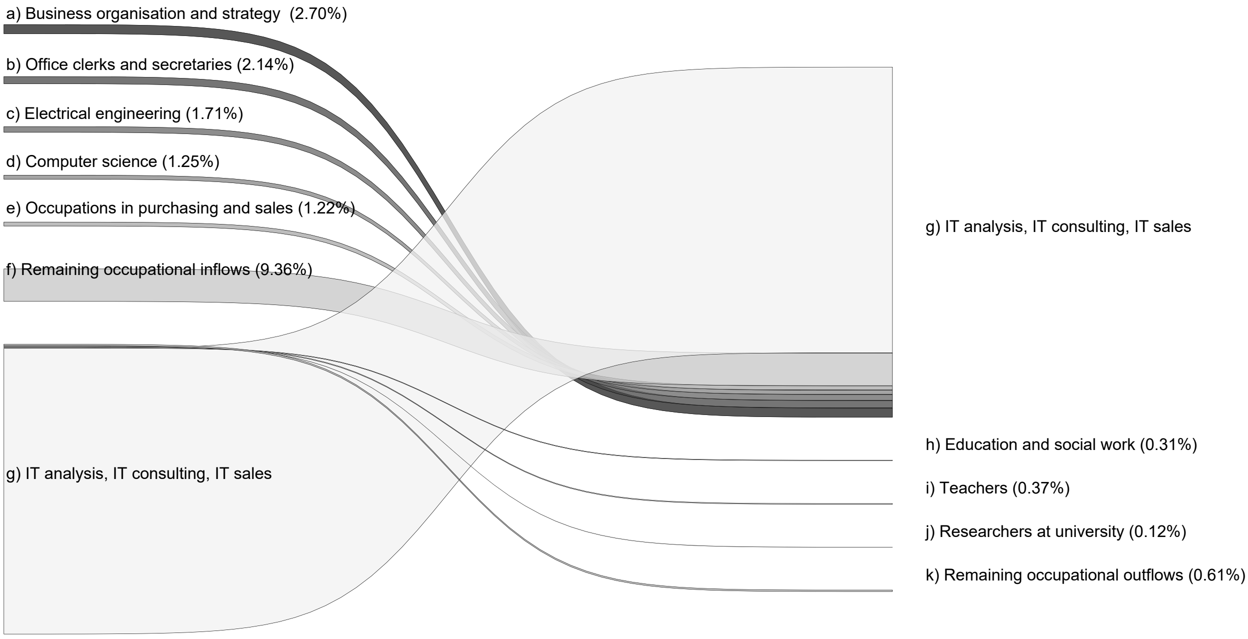

Notes: The figure depicts the projected flows over 10 years of IT consulting and sales with its top net sending and receiving occupations. The numbers shown are percentages of final (for inflows) and initial (for outflows) employment. IT consulting and sales makes up 0.95% of total employment initially and 1.15% after ten years, i.e. is projected to grow by over 20% as a share of its initial employment.

There exists a whole range of interesting occupations for flow analyses. In Appendix H, we discuss the occupation Machine building and operating, which turns out effectively low-elastic since highly substitutable occupations with it also endure negative shocks. As a result, Machine operators experience a substantial wage decline, which is so strong that it induces them to flow to a diverse set of more distant occupations (e.g., Business organisation, Drivers, and Building construction). We also discuss the flows of Laboratory occupations in medicine, which illustrate that the equilibrium pathways of employment adjustment emerge from the interplay between elasticities and relative price changes. In particular, Laboratory occupations in medicine have a substantially more negative demand shock than closely substitute Doctors' receptionists and assistants. Still, we observe no detectable flows from the former to the latter, since equilibrium wage changes turn out very similar. For further details and discussion on other occupations, such as Building construction, see Appendix H.

Overall, the analysis presented in this section displayed a couple of novel strengths for projecting the future of the labour market. First, by explicitly modelling the heterogeneous labour supplies, we show how varyingly occupational employment will adjust in response to shocks. This feature is directly complementary to recent work which focuses on predicting the demand shocks that will hit occupations in the context of rapid technological changes (e.g., Webb, 2020; Eloundou et al., 2023; Felten et al., 2023).[41] Second, we calculate how market prices change given the shocks and the labour supply. This modelling of price pressures is a fundamental departure from analyses that focus only on the heterogeneity of past flows between occupations, which is for example done in the context of job search advice (see also Belot et al., 2019; Altmann et al., 2023). As discussed in the case of Laboratory occupations in medicine, these flows and pathways for alternative employment opportunities may be substantially affected by equilibrium prices. Third, using a model that combines demand shocks with heterogeneity and market equilibrium can improve policy in a number of areas: It allows, for example, to better design job search advice and re-training programmes that are aligned with workers' existing employment experiences as well as with the actual occupational opportunities that will arise from demand and supply. It also allows for identifying precisely those occupational areas in which there is most need for active human capital policy.[42] Finally, any comprehensive analysis of the inequality implications of occupational shocks will need to take into account both the resulting employment and wage changes in equilibrium.

8 Conclusion

Shifts in the demand for occupations have led to significant changes in employment and wages across the developed world. However, a key aspect that remains relatively unexplored is how labour supply responses mediate these shifts. In this paper, we develop a tractable equilibrium model of the labour market for occupations that incorporates heterogeneous labour supply elasticities. We use this model to analyse the uneven effects of recent occupational demand and supply shifts, and to predict wage changes and employment flows in response to future automation-driven demand shocks.

To quantify these effects, we introduce a measure of occupation-specific labour supply elasticities, capturing how employment responds to wage changes across occupations. This includes 'own-price' and 'cross-price' elasticities, which reflect variations in substitutability across occupation pairs. We show how these elasticities can be derived from job flow moments and how they relate to key occupational characteristics, such as licensing requirements and task content. We embed this supply model into an equilibrium framework and implement it using administrative panel data from Germany. Our findings

[41]So-called 'employment projections' of agencies like the U.S. Bureau of Labor Statistics (2024) predict demand shocks more generally but typically do not model the labour supply side either.

[42]For example, and as seen above, the growth of dedicated IT occupations turns out not to be particularly constrained by labour supply. In contrast, providing the necessary new workers to the rising health sector, and facilitating the adjustments of workers in declining manufacturing jobs, may require special policy attention.

highlight two key insights: (i) occupational supply heterogeneity plays a crucial role in shaping labour market adjustments, and (ii) spillovers from correlated shocks among substitutable occupations are particularly important in determining whether occupations primarily adjust through employment shifts or wage changes.

We also apply our model alongside data from 'Job Futuromat' (2023), produced by the German Federal Employment Agency, to predict the occupational labour market effects of impending automation-driven demand shifts. Our analysis suggests that supply heterogeneities will likely amplify both skill shortages and demand declines. Specifically, we project upward wage pressure in the health and education sectors accompanied by a continued decline in manufacturing. By forecasting job mobility flows, our model provides detailed insights into occupational transitions, helping policymakers anticipate structural changes and potential labour market inequalities.

These findings underscore the need for targeted policies addressing occupation-specific frictions. Our model provides a quantitative tool for evaluating interventions that affect labour supply elasticities, such as those related to occupational licensing (Kleiner & Xu, 2024), educational content (Eckardt, 2024), or broader labour market policies (Autor et al., 2023). Future research could extend this framework to examine how occupational shifts affect individual careers, where heterogeneous substitutabilities are an important factor that is typically absent in the literature (e.g. Autor et al., 2014; Edin et al., 2023). Finally, our approach can be adapted to analyse frictions across other dimensions, such as geography or demographic disparities, helping to better understand inequalities and how they may be impacted by projected changes to work.

References

Acemoglu, D. & D. Autor (2011). Skills, Tasks, and Technologies: Implications for Employment and Earnings. In D. Card & O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4, Chapter 12, pp. 1043–1171. Elsevier. doi: 10.1016/S0169-7218(11)02410-5.

Acemoglu, D., D. Autor, J. Hazell, & P. Restrepo (2022). Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics 40*(S1), S293–S340. doi: 10.1086/718327.

Acemoglu, D. & P. Restrepo (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica 90*(5), 1973–2016. doi: 10.3982/ECTA19815.

Adão, R., M. Beraja, & N. Pandalai-Nayar (2024). Fast and Slow Technological Transitions. *Journal of Political Economy: Macroeconomics 2*(2). doi: 10.1086/730223.

Agarwal, N., A. Moehring, P. Rajpurkar, & T. Salz (2024). Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology. Available from: `https://economics.mit.edu/sites/default/files/2024-04/agarwal-et-al-diagnostic-ai.pdf`.

Alekseeva, L., J. Azar, M. Giné, S. Samila, & B. Taska (2021). The Demand for AI Skills in the Labor Market. *Labour Economics 71*, 102002. doi: 10.1016/j.labeco.2021.102002.

Altmann, S., A. Glenny, R. Mahlstedt, & A. Sebald (2023). The Direct and Indirect Effects of Online Job Search Advice. Available from: `https://drive.google.com/file/d/19l1OMYp49-VrXatLftRx1g5MILujEbAj/view`.

Antonczyk, D., B. Fitzenberger, & U. Leuschner (2009). Can a Task-Based Approach Explain the Recent Changes in the German Wage Structure? *Jahrbücher für Nationalökonomie und Statistik 229*(2-3), 214–238. doi: 10.1515/jbnst-2009-2-309.

Arnold, D. (2021). Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes. Available from: `https://darnold199.github.io/madraft.pdf`.

Artuç, E., S. Chaudhuri, & J. McLaren (2010). Trade Shocks and Labor Adjustment: A Structural Empirical Approach. *American Economic Review 100*(3), 1008–1045. doi: 10.1257/aer.100.3.1008.

Autor, D., L. Katz, & M. Kearney (2006). The Polarization of the US Labor Market. *American Economic Review: Papers and Proceedings 96*(2), 189–194. doi: 10.1257/000282806777212620.

Autor, D., L. Katz, & M. Kearney (2008). Trends in US Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics 90*(2), 300–323. doi: 10.1162/rest.90.2.300.

Autor, D. A., D. A. Mindell, & E. Reynolds (2023). *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines*. MIT Press. doi: 10.7551/mitpress/14109.001.0001.

Autor, D. H. & 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. doi: 10.1257/aer.103.5.1553.

Autor, D. H., D. Dorn, G. H. Hanson, & J. Song (2014). Trade adjustment: Worker-Level Evidence. *The Quarterly Journal of Economics 129*(4), 1799–1860. doi: 10.1093/qje/qju026.

Autor, D. H., F. Levy, & R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118(4)*, 1279–1333. doi: 10.1162/003355303322552801.

Bassier, I. (2024). Collective Bargaining and Spillovers in Local Labor Markets. Available from: https://drive.google.com/file/d/1PSo0lZ_uf51Ksen4hu9Ej-dGYWKzqlMM/view.

Belot, M., P. Kircher, & P. Muller (2019). Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice. *The Review of Economic Studies 86*(4), 1411–1447. doi: 10.1093/restud/rdy059.

Bender, S., H. Jürgen, R. Götz, & R. Helmut (1996). Die beschäftigtenstichprobe 1975-1990. Beiträge zur Arbeitsmarkt- und Berufsforschung 197 (Nuremberg, Germany: IAB).

Berger, D., K. Herkenhoff, & S. Mongey (2022). Labor Market Power. *American Economic Review 112*(4), 1147–1193. doi: 10.1257/aer.20191521.

Bhalotra, S., M. Fernández Sierra, & F. Wang (2023). The Distribution of the Gender Wage Gap: An Equilibrium Model. Available from: https://fanwangecon.github.io/assets/BhalotraFernandezWangMexicoFLFP.pdf.

Bocquet, L. (2024). The Network Origin of Slow Labor Reallocation. Available from: https://drive.google.com/file/d/1IDP_mC3uTv71RNWijWmfkrvvp477N0S3/view.

Böhm, M. J. (2020). The Price of Polarization: Estimating Task Prices under Routine-Biased Technical Change. *Quantitative Economics 11*(2), 761–799. doi: 10.3982/QE1031.

Böhm, M. J., H.-M. von Gaudecker, & F. Schran (2024). Occupation Growth, Skill Prices, and Wage Inequality. *Journal of Labor Economics 42*(1). doi: 10.1086/722084.

Borjas, G. J. (1980). The Relationship between Wages and Weekly Hours of Work: The Role of Division Bias. *The Journal of Human Resources 15*(3), 409–423. doi: 10.2307/145291.

Borusyak, K., R. Dix-Carneiro, & B. Kovak (2023). Understanding Migration Responses to Local Shocks. Available from: https://www.andrew.cmu.edu/user/bkovak/borusyak_dixcarneiro_kovak.pdf.

Brynjolfsson, E., D. Li, & L. Raymond (2025). Generative AI at Work. *The Quarterly Journal of Economics*, qjae044. doi: 10.1093/qje/qjae044.

Burstein, A., E. Morales, & J. Vogel (2019). Changes in Between-Group Inequality: Computers, Occupations, and International Trade. *American Economic Journal: Macroeconomics 11*(2), 348–400. doi: 10.1257/mac.20170291.

Caliendo, L., M. Dvorkin, & F. Parro (2019). Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica 87*(3), 741–835. doi: 10.3982/ECTA13758.

Card, D., A. R. Cardoso, J. Heining, & P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics 36*(S1), S13–S70. doi: 10.1086/694153.

Card, D., J. Heining, & P. Kline (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics 128*, 967–1015. doi: 10.1093/qje/qjt006.

Carlsson, M., J. Messina, & O. N. Skans (2016). Wage Adjustment and Productivity Shocks. *The Economic Journal 126*(595), 1739–1773. doi: 10.1111/ecoj.12214.

Cavaglia, C. & B. Etheridge (2020). Job Polarization and the Declining Quality of Knowlegde Workers: Evidence from the UK and Germany. *Labour Economics 66*, 101884. doi: 10.1016/j.labeco.2020.101884.

Cortes, G. & G. Gallipoli (2018). The Costs of Occupational Mobility: An Aggregate Analysis. *Journal of the European Economic Association 16*(2), 275–315. doi: 10.1093/jeea/jvx006.

Cortes, G. M. (2016). Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data. *Journal of Labor Economics 34*(1), 63–105. doi: 10.1086/682289.

Cortes, M., A. Lerche, U. Schönberg, & J. Tschopp (2024). Technological Change, Firm Heterogeneity and Wage Inequality. Available from: `https://www.dropbox.com/scl/fi/4gqonmxwt6em1agawdkj6/CLST_Mar2024.pdf?rlkey=xus6g121dsrmqigmtvpcnhqm7&dl=0`.

Cui, Z. K., M. Demirer, S. Jaffe, L. Musolff, S. Peng, & T. Salz (2025). The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. Available from: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4945566`.

Dauth, W. & J. Eppelsheimer (2020). Preparing the Sample of Integrated Labour Market Biographies (SIAB) for Scientific Analysis: a Guide. *Journal for Labour Market Research 54*, 10. doi: 10.1186/s12651-020-00275-9.

Dauth, W., S. Findeisen, J. Suedekum, & N. Woessner (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association 19*(6), 3104–3153. doi: 10.1093/jeea/jvab012.

Dustmann, C., J. Ludsteck, & U. Schönberg (2009). Revisiting the German Wage Structure. *The Quarterly Journal of Economics 124*(2), 843–881. doi: 10.1162/qjec.2009.124.2.843.

Eaton, J. & S. Kortum (2002). Technology, Geography, and Trade. *Econometrica 70*(5), 1741–1779. doi: 10.1111/1468-0262.00352.

Eckardt, D. (2024). Training Specificity and Occupational Mobility: Evidence from German Apprenticeships. Available from: `https://www.dropbox.com/scl/fi/0q0dxvgctllvpdmec3lqd/Eckardt_paper_Nov24.pdf?rlkey=0k7q6v7l8n0lahrz9r40frk4g&e=1&dl=0`.

Edin, P.-A., T. Evans, G. Graetz, S. Hernnäs, & G. Michaels (2023). Individual Consequences of Occupational Decline. *The Economic Journal 133*(654), 2178–2209. doi: 10.1093/ej/uead027.

Eloundou, T., S. Manning, P. Mishkin, & D. Rock (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. Technical report. doi: 10.48550/arXiv.2303.10130.

Felten, E. W., M. Raj, & R. Seamans (2023). Occupational Heterogeneity in Exposure to Generative AI. Technical report. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4414065.

Fitzenberger, B. (1999). Wages and Employment Across Skill Groups: An Analysis for West Germany. Springer. doi: 10.1007/978-3-642-58687-3.

Frodermann, C., A. Schmucker, S. Seth, & P. vom Berge (2021). Sample of Integrated Labour Market Biographies (SIAB) 1975-2019. Technical report, FDZ-Datenreport, 01/2021 (en), Nuremberg. doi: 10.5164/IAB.FDZD.2101.en.v1.

Gartner, H. (2005). The Imputation of Wages above the Contribution Limit with the German IAB Employment Sample. Technical report, FDZ Methodenreport. Available from: https://doku.iab.de/fdz/reporte/2005/MR_2.pdf.

Gathmann, C. & U. Schönberg (2010). How General is Human Capital? A Task-Based Approach. *Journal of Labor Economics 28*(1), 1–49. doi: 10.1086/649786.

German Federal Employment Agency (2023). Job Futuromat: Will Digital Technologies Change Your Job? Accessed March 07, 2024. Available from: https://job-futuromat.iab.de/en/.

Goos, M., A. Manning, & A. Salomons (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review 104*(8), 2509–2526. doi: 10.1257/aer.104.8.2509.

Graetz, G. & G. Michaels (2018). Robots at Work. *The Review of Economics and Statistics 100*(5), 753–768. doi: 10.1162/rest˙a˙00754.

Green, D. A., B. Sand, I. G. Snoddy, & J. Tschopp (2024). The Impact of Unions on Nonunion Wage Setting: Threats and Bargaining. Technical report. Available from: https://economics.ubc.ca/wp-content/uploads/sites/38/2021/07/Union_Spillovers_Manuscript.pdf.

Gregory, T., A. Salomons, & U. Zierahn (2022). Racing With or Against the Machine? Evidence on the Role of Trade in Europe. *Journal of the European Economic Association 20*(2), 869–906. doi: 10.1093/jeea/jvab040.

Hall, A., A. Siefer, & M. Tiemann (2012). BIBB/BAuA Employment Survey of the Working Population on Qualification and Working Conditions in Germany 2012. SUF 6.0; Research Data Center at BIBB (ed.); GESIS Cologne (data access); Bonn: Federal Institute for Vocational Education and Training. doi: 10.7803/501.12.1.1.60.

Hampole, M., D. Papanikolaou, L. D. Schmidt, & B. Seegmiller (2025). Artificial Intelligence and the Labor Market. Technical Report 33509, National Bureau of Economic Research. Available from: http://www.nber.org/papers/w33509.

Head, K. & T. Mayer (2013). Gravity Equations: Workhorse, Toolkit, and Cookbook. In *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier. Available from: https://sciencespo.hal.science/hal-00973067/document.

Hsieh, C.-T., E. Hurst, C. I. Jones, & P. J. Klenow (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica 87*(5), 1439–1474. doi: 10.3982/ECTA11427.

Humlum, A. (2021). Robot Adoption and Labor Market Dynamics. Accessed May 06, 2024. Available from: https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/61a4e7d3dfd0aa67849b9918/1638197204858/humlumJMP.pdf.

Hyman, B. G., B. K. Kovak, & A. Leive (2024). Wage Insurance for Displaced Workers. Available from: https://www.andrew.cmu.edu/user/bkovak/HKL_Wage_Insurance.pdf.

Jung, J. & J. Mercenier (2014). Routinization-Biased Technical Change and Globalization: Understanding Labor Market Polarization. *Economic Inquiry 52*(4), 1446–1465. doi: 10.1111/ecin.12108.

Kleiner, M. M. & M. Xu (2024). Occupational Licensing and Labor Market Fluidity. *Journal of Labor Economics*. doi: 10.1086/730120.

Lamadon, T., M. Mogstad, & B. Setzler (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review 112*(1), 169–212. doi: 10.1257/aer.20190790.

McFadden, D. L. (1973). Conditional Logit Analysis of Qualitative Choice Behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press: New York. Available from: https://eml.berkeley.edu/reprints/mcfadden/zarembka.pdf.

Morchio, I. & C. Moser (2024). The Gender Pay Gap: Micro Sources and Macro Consequences. Technical report, NBER WP 32408. Available from: https://www.nber.org/papers/w32408.

Paulus, W., B. Matthes, et al. (2013). The German Classification of Occupations 2010 – Structure, Coding and Conversion Table. *FDZ-Methodenreport 8*, 2013. Available from: https://doku.iab.de/fdz/reporte/2013/MR_08-13_EN.pdf.

Rohrbach-Schmidt, D. & M. Tiemann (2013). Changes in Workplace Tasks in Germany – Evaluating Skill and Task Measures. *Journal for Labour Market Research 46*, 215–237. doi: 10.1007/s12651-013-0140-3.

Sattinger, M. (1993). Assignment Models of the Distribution of Earnings. *Journal of Economic*

Literature 31(2), 831–880. Available from: https://www.jstor.org/stable/2728516.

Schubert, G., A. Stansbury, & B. Taska (2024). Employer Concentration and Outside Options. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3599454.

Sorkin, I. (2018). Ranking Firms Using Revealed Preference. *The Quarterly Journal of Economics 133*(3), 1331–1393. doi: 10.1093/qje/qjy001.

Spitz-Oener, A. (2006). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics 24*(2), 235–270. doi: 10.1086/499972.

Steiner, V. & K. Wagner (1998). Has Earnings Inequality in Germany Changed in the 1980s? *Zeitschrift für Wirtschafts- und Sozialwissenschaften 118*(1), 29–59. doi: 10.3790/schm.118.1.29.

Traiberman, S. (2019). Occupations and Import Competition: Evidence from Denmark. *American Economic Review 109*(12), 4260–4301. doi: 10.1257/aer.20161925.

U.S. Bureau of Labor Statistics (2024). Employment Projections Method Overview. Accessed February 02, 2025. Available from: https://www.bls.gov/emp/documentation/projections-methods.htm#occupational_employment.

Vicari, B. (2014). Degree of Standardised Certification of Occupations. Technical report, FDZ Methodenreport. Available from: https://doku.iab.de/fdz/reporte/2014/MR_04-14_EN.pdf.

Webb, M. (2020). The Impact of Artificial Intelligence on the Labor Market. Available from: https://www.michaelwebb.co/webb_ai.pdf.

For Online Publication (Appendix for):

The Impact of Labour Demand Shocks when Occupational Labour Supplies are Heterogeneous

Michael J. Böhm, Ben Etheridge & Aitor Irastorza-Fadrique

A Formal Results on the Elasticity Matrix

This section further develops the model introduced in Section 2 of the paper, integrating additional aspects for a more comprehensive analysis. We begin by presenting formal derivations of the main remarks and a deeper exploration of their underlying intuition. We then derive another formal result on the full vector of price changes.

A.1 Derivation of Remark 1 (Elasticities and Job Flows)

We start by formally deriving key Remark 1. To simplify notation in the following, we define 'choice index' $\lambda(\mathbf{p}) \equiv \frac{1}{\sum_{k=1}^{N}\exp(\theta p_k + a_{ik})}$, where $\mathbf{p}$ represents the vector of log prices. The fraction of individuals working in sector j as a function of log prices, denoted by $E_j(\mathbf{p})$, can then be expressed as:

$$E_j(\mathbf{p}) = \sum_i \tau_i \lambda(\mathbf{p}) \exp\left(\theta p_j + a_{ij}\right)$$

Recall that our interest centres on (own- and cross-occupation) price elasticities, the response of employment in occupation j to occupation k's log price change. Using the accounting identity presented in equation (3), we formally write this as:

$$\frac{\partial E_j(\mathbf{p})}{\partial p_k} = \sum_i \tau_i \left(\lambda(\mathbf{p}) \frac{\partial \exp\left(\theta p_j + a_{ij}\right)}{\partial p_k} + \frac{\partial \lambda(\mathbf{p})}{\partial p_k} \exp\left(\theta p_j + a_{ij}\right) \right)$$

Computing the second element in the brackets, $\frac{\partial \lambda(\mathbf{p})}{\partial p_k}$, gives:

$$\begin{aligned}
\frac{\partial \lambda(\mathbf{p})}{\partial p_k} &= -\frac{\theta \exp\left(\theta p_k + a_{ik}\right)}{\left(\sum_s \exp\left(\theta p_s + a_{is}\right)\right)^2} \\
&= -\theta \frac{1}{\sum_s \exp\left(\theta p_s + a_{is}\right)} \frac{\exp\left(\theta p_k + a_{ik}\right)}{\sum_s \exp\left(\theta p_s + a_{is}\right)} \\
&= -\theta \lambda(\mathbf{p}) \pi_{ik}(\mathbf{p})
\end{aligned}$$

By combining these results, we derive the following expression:

$$\frac{\partial E_j(\mathbf{p})}{\partial p_k} = \begin{cases} \sum_i \tau_i \theta \left(\pi_{ij}(\mathbf{p})\left(1 - \pi_{ij}(\mathbf{p})\right)\right) & \text{if } j = k \\ -\sum_i \tau_i \theta \left(\pi_{ij}(\mathbf{p})\, \pi_{ik}(\mathbf{p})\right) & \text{otherwise} \end{cases} \tag{24}$$

Finally, writing $e_j \equiv \ln E_j(\mathbf{p})$, we obtain:

$$\begin{aligned}\frac{\partial e_j(\mathbf{p})}{\partial p_k} &= \frac{1}{E_j(\mathbf{p})}\frac{\partial E_j(\mathbf{p})}{\partial p_k} \\ &= \theta\begin{cases}\frac{\sum_i \tau_i\left(\pi_{ij}(\mathbf{p})\left(1-\pi_{ij}(\mathbf{p})\right)\right)}{\sum_i \tau_i \pi_{ij}(\mathbf{p})} & \text{if } j=k \\ \frac{-\sum_i \tau_i\left(\pi_{ij}(\mathbf{p})\pi_{ik}(\mathbf{p})\right)}{\sum_i \tau_i \pi_{ij}(\mathbf{p})} & \text{otherwise}\end{cases}\end{aligned}$$

These are equations (4)–(5) in Section 2. It shows that the short-term partial derivative of occupation j's log employment share with respect to k's log price can be computed using (baseline) transition probabilities, and a pecuniary parameter θ. We next discuss alternative formulations of the elasticities in terms of moments of job flows.

A.2 Discussion of Remark 2 (Occupational Substitutabilities)

A.2.1 Derivation

We have described the off-diagonal elements of the elasticity matrix D as:

$$d_{jk} = -\frac{1}{\tau_j}\sum_i \tau_i \pi_{ij}\pi_{ik}$$

where π_{ij}, π_{ik} are elements of the transition matrix and τ_i is the ith element of the associated stationary vector. To interpret this further, consider the weighted covariance between columns of the normalised transition matrix:

$$\begin{aligned}Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) &\equiv \sum_i \tau_i\left(\tilde{\pi}_{ij} - \mathbb{E}_\tau\tilde{\pi}_{.,j}\right)\left(\tilde{\pi}_{ik} - \mathbb{E}_\tau\tilde{\pi}_{.,k}\right) \\ &= \sum_i \tau_i\left(\tilde{\pi}_{ij}-1\right)\left(\tilde{\pi}_{ik}-1\right)\end{aligned}$$

where

$$\tilde{\pi}_{iq} \equiv \frac{\pi_{iq}}{\tau_q}$$

and the second line follows from the first because $\sum_i \tau_i\tilde{\pi}_{iq} = \frac{1}{\tau_q}\sum_i \tau_i\pi_{iq} = \frac{\tau_q}{\tau_q} = 1$.

Expanding this further:

$$\begin{aligned}Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) &= \sum_i \tau_i\left(\tilde{\pi}_{ij}-1\right)\left(\tilde{\pi}_{ik}-1\right) \\ &= \sum_i \tau_i\tilde{\pi}_{ij}\tilde{\pi}_{ik} - \sum_i \tau_i\tilde{\pi}_{ij} - \sum_i \tau_i\tilde{\pi}_{ik} + \sum_i \tau_i\end{aligned}$$

$$= \frac{1}{\tau_j \tau_k} \sum_i \tau_i \pi_{ij} \pi_{ik} - 1 - 1 + 1$$
$$= -\frac{1}{\tau_k} d_{jk} - 1$$

Rearranging gives equation (6). Then summing over $k \neq j$, and remembering that $1 - \tau_j = \sum_{k \neq j} \tau_k$ and $d_{jj} = -\sum_{k \neq j} d_{jk}$, gives equation (7).

Finally, we derive an alternative version of (7) in terms of occupations' job-flow dispersion. Start with the on-diagonal elements of the elasticity matrix D:

$$d_{jj} = \frac{1}{\tau_j} \sum_i \tau_i \pi_{ij} \left(1 - \pi_{ij}\right)$$

Similar to the above, we can express this in terms of the weighted variance of normalised transition probabilities:

$$\begin{aligned}
d_{jj} &= \frac{1}{\tau_j} \sum_i \tau_i \pi_{ij} - \frac{1}{\tau_j} \sum_i \tau_i \pi_{ij}^2 \\
&= \quad 1 \quad - \frac{1}{\tau_j} \sum_i \tau_i \pi_{ij}^2 \\
&= 1 - \frac{1}{\tau_j} \left(Var_\tau \left(\pi_{.,j}\right) + \left(\mathbb{E}_\tau \pi_{.j}\right)^2 \right) \\
&= 1 - \frac{1}{\tau_j} \left(Var_\tau \left(\pi_{.,j}\right) + \tau_j^2 \right) \\
&= 1 - \tau_j \left(1 + \frac{1}{\tau_j^2} Var_\tau \left(\pi_{.,j}\right) \right) \\
&= 1 - \tau_j \left(1 + Var_\tau \left(\tilde{\pi}_{.,j}\right)\right) \qquad (25)
\end{aligned}$$

Rearranging gives $d_{jj} = -\tau_j Var_\tau \left(\tilde{\pi}_{.,j}\right) + 1 - \tau_j$.

A.2.2 Choice of Normalisation and Invariance to the Occupation Classification

Now we turn to a discussion of our choices of normalisations.

We first consider $Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) = \sum_i \tau_i \left(\tilde{\pi}_{ij} - \mathbb{E}_\tau \tilde{\pi}_{.,j}\right) \left(\tilde{\pi}_{ik} - \mathbb{E}_\tau \tilde{\pi}_{.,k}\right)$. Because $\mathbb{E}_\tau \tilde{\pi}_{.,j} = \mathbb{E}_\tau \tilde{\pi}_{.,k} = 1$, we argue this term is invariant to occupation size. To show this empirically, we examine the distribution of this term for occupational classifications at various levels of coarseness. In particular, Table A.1 reports the median across occupations for three levels of aggregation: 4 main groups as described below in Appendix D, 10 occupation groups corresponding to one-digit categories of the 1988 *Klassifikation der Berufe*, and the 120 occupations considered in the analysis (see Table E.5 for the full list).

We now consider the variance terms. We can also write d_{jj} as follows

$$
\begin{aligned}
d_{jj} &= -\sum_{k\neq j} d_{jk} \\
&= \sum_{k\neq j} \tau_k \left(1 + Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)\right) \\
&= \sum_{k\neq j} \tau_k + \sum_{k\neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) \\
&= 1 - \tau_j + \sum_{k\neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) \qquad (26)
\end{aligned}
$$

Equating equations (25) and (26), we see that

$$
\begin{aligned}
Var_\tau\left(\tilde{\pi}_{.,j}\right) &= -\frac{1}{\tau_j}\sum_{k\neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) \\
\implies \tau_j Var_\tau\left(\tilde{\pi}_{.,j}\right) &= -\sum_{k\neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)
\end{aligned}
$$

These expressions show two things. First, because $Var_\tau\left(\tilde{\pi}_{.,j}\right)$ is necessarily greater than zero, then $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ is below zero on average.[43] Second, if $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ is of order $\mathcal{O}(1)$, then $\frac{1}{\tau_j}\sum_{k\neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ is of order $\mathcal{O}(N)$. In contrast, $\tau_j Var_\tau\left(\tilde{\pi}_{.,j}\right) = -\sum_{k\neq j} \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ is a weighted average of the covariance terms, and so is of order $\mathcal{O}(1)$. To show this empirically, Table A.1 also reports the median value across occupations for both measures of the variance, again for the three levels of occupational aggregation.

Table A.1: Median Values of Model Components Across Occupation Pairs

# Occs	$Cov(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k})$	$Var(\tilde{\pi}_{.,j})$	$\tau_j Var(\tilde{\pi}_{.,j})$
4	-0.76	2.20	0.54
10	-0.75	5.48	0.58
120	-0.78	126.91	0.57

Notes: Variances and covariance computed across sending occupations, given destination occupations j and k. Table then shows median values across these destination occupations. The occupations in the aggregation to four broad groups are (1) managers, professionals, and technicians, (2) sales and office workers, (3) production workers, operators, and craftsmen, and (4) workers in services and care occupations. In the ten broad groups, they are 1-digit level occupations as in, e.g, Acemoglu & Autor (2011); Böhm et al. (2024). For further details on occupations and their aggregations see Section D.1.

[43] This also shows that $\sum_k \tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) = 0$.

A.3 Remark 5: Vector of Price Changes

This section develops a further result on the aggregation of Remark 2 for individual own- and cross-price elasticities. This formalises the effects of the full vector of price changes and provides a rigorous interpretation of overall employment changes studied in eq. (8) in terms of distributions of worker flows.

Remark 5 (Vector of Price Changes) *Matrix D can be expressed as follows*

$$D = I - W - W \circ C \tag{27}$$

where I is the identity matrix, W is the matrix of stationary employment shares with j,k-th element τ_k, $\circ$ is the element-by-element product, and C is the symmetric matrix with j,k-th element $c_{jk} = Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$, which captures the 'occupational substitutability' between occupations j and k.

Accordingly, following a vector of price changes $\Delta\mathbf{p}$, then the change in the employment share in occupation j is given by

$$\begin{aligned} \Delta e_j &\approx \theta \mathbf{d_j} \Delta \mathbf{p} \\ &= \theta \left(\underbrace{\Delta p_j - \Delta \mathbb{E}_\tau \mathbf{p}}_{\substack{\text{real price} \\ \text{change}}} + \underbrace{Cov_\tau\left(c_{.,j}, \Delta p_j - \Delta \mathbf{p}\right)}_{\substack{\text{occupational} \\ \text{substitutability}}} \right) \qquad (28) \\ &= \theta \left(\underbrace{\left(1 - \tau_j - \tau_j c_{jj}\right) \Delta p_j}_{\substack{\text{own-occupation} \\ \text{effect}}} + \underbrace{\sum_{k \neq j} \left(-\tau_k - \tau_k c_{jk}\right) \Delta p_k}_{\substack{\text{total cross-} \\ \text{occupation effect}}} \right) \qquad (29) \end{aligned}$$

where $\mathbb{E}_\tau \mathbf{p}$ is the (weighted) average of prices across occupations and we drop a time subscript for ease of notation. Similarly, $Cov_\tau\left(c_{.,j}, \Delta p_j - \Delta \mathbf{p}\right)$ captures the (weighted) covariance between the j-th column of C, $c_{.,j}$, and the vector of relative price changes $\Delta p_j - \Delta \mathbf{p}$.

Remark 5 complements the interpretations contained in Remark 2. In the formulation in equation (28), the effect of a vector of price changes on a given occupation consists of two components. First is the direct effect of real price changes in that occupation itself, net of the change in the economy-wide price (wage) index. This term aggregates the 'direct' and 'price index' terms contained in equations (6) and (7). Second is the total effect of occupational substitutabilities: Employment growth is larger if price growth is higher relative to more similar occupations. In fact, empirically, price changes are positively correlated across similar occupations, and so this last component tends to attenuate the direct effect of price changes. To see this, consider, for example, wage growth in occupations high in analytical tasks. Price growth in these occupations has been highest relative to routine and manual occupations, which saw the largest declines, but which are also dissimilar in terms of occupational flows. Therefore, for these analytical occupations, this last term is likely negative, offsetting the positive effect from the first two terms.

Equation (29) then builds on this formulation by relating it back to equation (9), which forms the basis of our empirical application. Equation (29) therefore expresses the effect of a vector of price changes in terms of two components which we can easily take to data, and which can be interpreted in terms of the joint distribution of these price changes with steady-state job flows.

Derivation: The expression

$$D = I - W - W \circ C$$

follows directly from Remark 2. The diagonal element c_{jj} of C is $Var_\tau\left(\tilde{\pi}_{.,j}\right)$.

We therefore have that

$$\begin{aligned}
\Delta e_j &= \theta \mathbf{d_j} \Delta \mathbf{p} \\
&= \theta \sum_k \left(i_{jk} - \tau_k - \tau_k c_{jk}\right) \Delta p_k \\
&= \theta \left(\sum_k i_{jk} \Delta p_k - \sum_k \tau_k \Delta p_k - \sum_k \tau_k c_{jk} \Delta p_k \right) \\
&= \theta \left(\Delta p_j - \Delta \mathbb{E}_\tau p - \sum_k \tau_k c_{jk} \left(\Delta p_k - \Delta p_j\right) \right) \\
&= \theta \left(\Delta p_j - \Delta \mathbb{E}_\tau p + Cov_\tau \left(c_{.,j}, \Delta p_j - \Delta p_.\right)\right)
\end{aligned}$$

as given in the text. The fourth line follows from the third because $\sum_k \tau_j c_{jk} = 0 \implies \sum_k \tau_j c_{jk} \Delta p_j = 0$. The final line follows from the fourth because similarly $\mathbb{E}_\tau c_{.,j} = 0$ and column vector $c_{.,j} = c_{j,.}$ because C is symmetric.

B Generalisations of the Labour Supply Model

B.1 Overlapping Generations Model

This section extends our theoretical framework to incorporate an overlapping-generations (OLG) structure, building upon the static model developed in Sections 2 and 3. The extended framework introduces three key elements: time-varying idiosyncratic occupation-amenity draws, aggregate risk in occupational wages, and endogenous human capital accumulation. The final addition generates within-occupation wage heterogeneity. While we present a two-period model for analytical tractability, the framework can be readily generalised to higher-frequency time intervals. The main insight of the analysis is that, even with this richer structure, the effect of wage shocks on occupational employment can still be characterised as interpretable functions of base occupational flows.

The model relates to that in, for example, Artuç et al. (2010) and Caliendo et al. (2019). A key difference is that, rather than focusing on transition dynamics following discrete policy changes (such as trade liberalisation), we incorporate persistent uncertainty in the wage process. This modification makes our framework more suitable for analysing longer-term trends in occupational demand. Compared to the benchmark model, the OLG structure yields three distinct advantages: 1) A detailed accounting framework that decomposes employment changes into early-career and late-career components; 2) A precise characterisation of labour market entry elasticities, capturing how new cohorts respond to changing occupational returns; and 3) An analysis of both short-run and longer-run elasticities.

B.1.1 Model Set-Up

OLG Structure

There are two working periods, early and late career, indexed by $g = 1$ and $g = 2$ (age is denoted by g because 'a' is used in earlier notation). There are N occupations. The decision problem at $g = 1$ is one of labour-market entry and captures early-career occupational choices given fixed education. The decision at $g = 2$ captures career choice from, say, age 40 onwards.

Let E_t, a vector of length N, capture employment observed at time t, such that

$$E_t = \tilde{E}_{t,1} + \tilde{E}_{t,2}$$

where $\tilde{E}_{t,g}$ is employment at time t for those aged g.

Wage Process

We model wages as following a persistent process. Let $\mathbf{p}_{t+1}$ be a vector of length N of occupational wages at time t. As is empirically plausible, we model $\mathbf{p}_{t+1}$ as following a random walk:

$$\mathbf{p}_{t+1} = \mathbf{p}_t + \boldsymbol{\eta}_{t+1}$$

where $\mathbf{p}_t$ is the vector of log wages at time t and $\boldsymbol{\eta}_t$ is a multivariate Gaussian following a $N(\kappa \mathbf{1}_N, \Psi)$ distribution, where κ is a scalar, $\mathbf{1}_N$ a vector of ones, and we henceforth denote $\kappa \mathbf{1}_N$ by $\bar{\boldsymbol{\eta}}$. We do not model the production side here but imagine that wages arise from a production function similar to that specified in Section 3, together with some underlying process for demand shocks. Even so, some structure on the production function is imposed by our wage process. For example, that wage shocks affect each cohort identically implies that younger and older workers are perfect substitutes within occupations.

Life-cycle Model

We now consider the life-cycle problem. As in the benchmark model, we allow for a continuum of workers, indexed by $\omega \in \Omega$. Each worker receives an occupation-specific match preference draw at age g, $\epsilon_g^j(\omega)$, which follows a standardised extreme value I (Gumbel) distribution. It is important to note that here we adapt the notation from the main body of the text by denoting occupations with superscripts. To be specific about distributions, $\epsilon_g^j(\omega)$ is i.i.d distributed with location parameter $-\gamma$ and scale 1, where γ is the Euler-Mascheroni constant. This ensures that $\epsilon_g^j(\omega)$ has mean 0 and variance $\frac{\pi^2}{6}$.

Problem at Career stage 2

We consider generic time t and first consider those in late career. If originating in occupation j at the beginning of age $g = 2$, then utility from choosing occupation k is given by

$$u_{t,2}^{jk}(\omega) = \theta \left(p_t^k + \underbrace{s_1(\omega) + h_2^{jk}}_{\substack{\text{age-2} \\ \text{human capital}}} \right) + b_2^k + \epsilon_2^k(\omega) \tag{30}$$

where p_t^k is the log price (wage) per unit of real output for working in occupation k at time t. Next, $s_1(\omega)$ reflects individual-specific skills, which may be inherited from the previous period (period 1), h_2^{jk} are real switching cost terms, which here we allow to differ over career stage for generality, and we interpret as changes to, or destruction of, human capital. For ease of interpretation with the wage data, we assume that $h^{ii} = c$, a constant, for all i, which reflects constant skill accumulation across occupations.[44] b_2^k then captures differential amenities across all occupations.[45] As in the benchmark model, θ gives the wage elasticity. It can also be thought of as scaling the importance of the idiosyncratic shocks. This formulation also implies that utility is measured in amenity units.

Define the maximised value for individual ω at career-stage 2 as follows:

$$v_{t,2}^j(\mathbf{p}_t, \omega) = \max_k \left(\theta \left(p_t^k + s_1(\omega) + h_2^{jk} \right) + b_2^k + \epsilon_2^k(\omega) \right)$$

Then $v_{t,2}^j(\mathbf{p}_t, \omega)$ itself follows a Gumbel distribution with scale 1, location parameter $\ln \sum_k \exp\left(\theta \left(p_t^k + s_1(\omega) + h_2^{jk} \right) + b_2^k \right) - \gamma$, and $\mathbf{p}_t$ is the vector of period-2 wages.

Average utility, given wages, can therefore be defined as

$$\tilde{V}_{t,2}^j(\mathbf{p}_t, s) \equiv \mathbb{E}_\epsilon \left[v_{t,2}^j(\mathbf{p}_t, \omega) \right] = \theta s + \ln \sum_k \exp\left(\theta \left(p_t^k + h_2^{jk} \right) + b_2^k \right)$$

which is a function of all the prices. Useful for later will be defining $V_{t,2}^j(\mathbf{p}_t)$ by

$$V_{t,2}^j(\mathbf{p}_t) \equiv \tilde{V}_{t,2}^j(\mathbf{p}_t, s) - \theta s$$

which is the expected value at normalised/'zero' skill level $s = 0$.

Note that the marginal effect of a change in wages on expected utility is given by:

$$\begin{aligned} \frac{\partial \tilde{V}_{t,2}^j}{\partial p_t^k} = \frac{\partial V_{t,2}^j}{\partial p_t^k} &= \frac{\partial}{\partial p_t^k} \ln \sum_l \exp\left(\theta \left(p_t^l + h_2^{jl} \right) + b_2^l \right) \\ &= \theta \pi_{t,2}^{jk} \end{aligned} \tag{31}$$

where $\pi_{t,2}^{jk}(\mathbf{p}_t)$, depends on wages but not human capital, and captures the probability that k is the most attractive occupation when beginning period-t and age-2 in occupation j. Alternatively, in our application, $\pi_{t,2}^{jk}$ is the period-t gross flow (or transition probability) from j to k among late-career individuals. Pursuing this discussion further, expression (31)

[44] Our approach to identifying underlying price movements separately from human capital accumulation, discussed in Section 4, allows for heterogeneous skill development across occupations. In the current appendix, we prefer to keep the exposition deliberately more straightforward.

[45] Including the b_2^k term is important because it allows the combined switching cost $\theta h_2^{jk} + b_2^k$ to be unrestricted. It is worth remarking we could define the switching cost relative to the counterfactual of staying, by $\theta\left(h_2^{jk} - h_2^{jj}\right) + b_2^k - b_2^j$. Similarly, this is unrestricted.

is natural given that V_2^j can also be defined as:

$$V_{t,2}^j(\mathbf{p}_t) = \sum_l \pi_2^{jl}\left(\theta\left(p_t^l + h_2^{jl}\right) + b_2^l + \mathbb{E}\left(\epsilon_2^l | l \text{ is chosen}\right)\right)$$

and, it turns out, marginal changes in the selection terms at the end of this expression cancel.

Problem at Career stage 1

The first career stage captures those entering the labour market from schooling, split by schooling-track or skill group ς. Expected utility from choosing occupation j is given by

$$u_{t,1}^{\varsigma,j}(\omega) = \phi\left(p_t^j + s_0(\omega) + h_1^{\varsigma,j}\right) + b_1^j + \epsilon_1^j(\omega) + \beta \underbrace{\mathbb{E}_{\mathbf{p}_{t+1}} \tilde{V}_{2,t+1}^j(\mathbf{p}_{t+1}, s_1)}_{\substack{\text{Continuation}\\ \text{utility from}\\ \text{occ j}}} \tag{32}$$

$$s_1 = s_0 + h_1^{\varsigma,j}$$

where β is a discount factor such that $0 < \beta < 1$, and ϕ captures the elasticity of labour entry, which may differ from the elasticity of labour switches in later career. This structure allows for substitutabilities on the supply side for closely related occupations, through the $h_1^{\varsigma,j}$ terms. These reflect differential human capital accumulation effects of entering occupations across different underlying skill groups.

Utility can be more simply written as

$$u_{t,1}^{\varsigma,j}(\omega) = \phi\left(p_t^j + \left(1+\tilde{\beta}\right)\left(s_0(\omega) + h_1^{\varsigma,j}\right)\right) + b_1^j + \epsilon_1^j(\omega) + \beta \mathbb{E}_{\mathbf{p}_{t+1}} V_{2,t+1}^j(\mathbf{p}_{t+1})$$

where $\tilde{\beta} = \beta\frac{\theta}{\phi}$. We therefore write the age-1, period-t problem as

$$\begin{aligned} v_{t,1}^{\varsigma}(\mathbf{p}_t, s) &\equiv \max_j u_{t,1}^{\varsigma,j}(\omega) \\ &= \max_j \left(\phi\left(p_t^j + \left(1+\tilde{\beta}\right)\left(s + h_1^{\varsigma,j}\right)\right) + b_1^j + \epsilon_1^j(\omega) + \beta U_{t+1,2}^j(\mathbf{p}_t)\right) \end{aligned}$$

where $U_{t+1,2}^j()$ is given by:

$$U_{t+1,2}^j(\mathbf{p}_t) \equiv \mathbb{E}_{\eta_{t+1}}\left(V_{t+1,2}^j(\mathbf{p}_t + \boldsymbol{\eta}_{t+1})\right) \tag{33}$$

$v_{t,1}^{\varsigma}()$ again follows a Gumbel distribution, here with location parameter

$$\phi\left(1+\tilde{\beta}\right)s+\ln\sum_k \exp\left(\phi p_t^k+\left(1+\tilde{\beta}\right)h_1^{\varsigma,k}+b_1^k+\beta U_{t+1,2}^k\left(\mathbf{p}_t\right)\right)-\gamma.$$

B.1.2 Summarising the Decision Problem and Aggregating Employment

Define $a_1^{\varsigma,j}\equiv\phi\left(1+\tilde{\beta}\right)h_1^{\varsigma,j}+b_1^j$ and $a_2^{jk}\equiv\theta h_2^{j,k}+b_2^k$, merging human capital and amenity effects of occupational choice. Given equations (30), (32) and (33), we can characterise the decision problem through the choice probabilities as follows:

$$\pi_{t,1}^{\varsigma,j}\left(\mathbf{p}_t\right)=\frac{\exp\left(\phi p_t^j+a_1^{\varsigma j}+\beta U_{t+1,2}^j\left(\mathbf{p}_t\right)\right)}{\sum_l \exp\left(\phi p_t^l+a_1^{\varsigma l}+\beta U_{t+1,2}^l\left(\mathbf{p}_t\right)\right)} \tag{34}$$

$$\pi_{t,2}^{ij}\left(\mathbf{p}_t\right)=\frac{\exp\left(\theta p_t^j+a_2^{ij}\right)}{\sum_l \exp\left(\theta p_t^l+a_2^{il}\right)}$$

where $\mathbf{p}_t$ is the only relevant (exogenous) state variable. Aggregate employment can then be characterised as:

$$\tilde{E}_{t,1}^j\left(\mathbf{p}_t\right)=\sum_\varsigma \tau^\varsigma\pi_{t,1}^{\varsigma,j}\left(\mathbf{p}_t\right)$$

$$\tilde{E}_{t,2}^j\left(\mathbf{p}_t,\mathbf{p}_{t-1}\right)=\sum_i \tilde{E}_{t-1,1}^i\left(\mathbf{p}_{t-1}\right)\pi_{t,2}^{ij}\left(\mathbf{p}_t\right) \tag{35}$$

where $\tilde{E}_{t-1,1}^i$ is given/fixed at time t but is a function of $t-1$ prices. τ^ς is the share of skill group ς which here we assume is fixed always. Finally, note that (35) implies that, in general, aggregate employment is a function of a full history of prices.

B.1.3 Model Timing

In the coarse modelling of this appendix, we consider time period $t-1$ as corresponding approximately to the years $1975-85$, t corresponding to $1985-1995/2000$ and $t+1$ corresponding to $1995/2000-2010$. Similarly, $g=1$ corresponds to ages $25-40$ and $g=2$ corresponds to ages $40-59$. This way the evolution of careers corresponds roughly with the evolution of calendar time. We leave a more precise alignment of model periods with data for future work.

We focus attention on a shock revealed at the beginning of time t and consider both one-period elasticities and two-period elasticities. Being more detailed, we imagine that the new wages for year 2000 are revealed to workers from 1985. $\tilde{E}_{t,1}^j\left(\mathbf{p}_t\right)$ for example, then captures early-career employment in year 2000, at wages observed by the econometrician in 2000 but revealed to workers from year 1985, with e.g. $\pi_{t,1}^{\varsigma,j}\left(\mathbf{p}_t\right)$ capturing the choice probabilities for labour-market entrants in this period, and, by analogy, the job flows.

B.1.4 The Effect of Period-t Shocks on end-of Period t Allocations

We now turn to our first main focus, the short-run effect of wage shocks. As in the benchmark model, our approach focuses on calculating gradients and elasticities of employment shares.

Effect on Late Career

Using (35) and abbreviating notation somewhat we have:

$$\begin{aligned}\frac{\partial \tilde{E}^{j}_{t,2}}{\partial p^{k}_{t}} &= \sum_{i} \tilde{E}_{t-1,1} \frac{\partial \pi^{ij}_{t,2}}{\partial p^{k}_{t}} \\ &= \begin{cases} -\theta \sum_i \tilde{E}_{t-1,1} \pi^{ij}_{t,2} \pi^{ik}_{t,2} & \text{if } j \neq k \\ \theta \sum_i \tilde{E}_{t-1,1} \pi^{ij}_{t,2} \left(1 - \pi^{ij}_{t,2}\right) & \text{if } j = k \end{cases}\end{aligned} \tag{36}$$

Effect on Early Career

Now consider early career, and first consider gradients of flows: $\frac{\partial \pi^{\varsigma,j}_{t,1}}{\partial p^{k}_{t}}$. For convenience, we drop the subscript on wages. Using (34) for $j = k$, this then implies

$$\frac{\partial \pi^{\varsigma,j}_{t,1}}{\partial p^{j}} = \pi^{\varsigma,j}_{t,1} \left(\phi \left(1 - \pi^{\varsigma,j}_{t,1}\right) + \beta \left(\frac{\partial U^{j}_{t+1,2}(\mathbf{p})}{\partial p^{j}} - \sum_{l} \pi^{\varsigma,l}_{t,1} \frac{\partial U^{l}_{t+1,2}(\mathbf{p})}{\partial p^{j}} \right) \right) \tag{37}$$

where $\beta \left(\frac{\partial U^{j}_{t+1,2}(\mathbf{p})}{\partial p^{j}} - \sum_{l} \pi^{\varsigma,l}_{t,1} \frac{\partial U^{l}_{t+1,2}(\mathbf{p})}{\partial p^{j}} \right)$ is the extra term coming from the life-cycle model.

For $j \neq k$, we have

$$\frac{\partial \pi^{\varsigma,j}_{t,1}}{\partial p^{k}} = \pi^{\varsigma,j}_{t,1} \left(-\phi \pi^{\varsigma,k}_{t,1} + \beta \left(\frac{\partial U^{j}_{t+1,2}(\mathbf{p})}{\partial p^{k}} - \sum_{l} \pi^{\varsigma,l}_{t,1} \frac{\partial U^{l}_{t+1,2}(\mathbf{p})}{\partial p^{k}} \right) \right) \tag{38}$$

Using a second-order Taylor-series expansion around the mean of $\boldsymbol{\eta}$, we obtain:

$$\frac{\partial U^{j}_{t+1,2}(\mathbf{p})}{\partial p^{k}} \approx \theta \left(\pi^{jk}_{t+1,2} + \mathbb{E}_{\boldsymbol{\eta}} \boldsymbol{\eta}' \frac{\partial^2 \left(\pi^{jk}_{2} (\mathbf{p} + \boldsymbol{\eta}) \right)}{\partial \boldsymbol{\eta}^2}_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}} \boldsymbol{\eta} \right) \tag{39}$$

For now, and for convenience, we consider shocks to wages as 'MIT' shocks and impose $\mathbb{E}_{\boldsymbol{\eta}} \boldsymbol{\eta}' \frac{\partial^2 \left(\pi^{jk}_{2} (\mathbf{p} + \boldsymbol{\eta}) \right)}{\partial \boldsymbol{\eta}^2}_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}} \boldsymbol{\eta} = 0$, which implies that $\frac{\partial U^{j}_{t+1,2}(\mathbf{p})}{\partial p^{k}} \approx \theta \pi^{jk}_{t+1,2}$. We provide a fuller characterisation of the effect of wage risk at the end of this Section.

Putting this together we have, for $j = j$:

$$\frac{\partial \pi_{t,1}^{\varsigma,j}}{\partial p^j} = \pi_{t,1}^{\varsigma,j}\left(\phi\left(1-\pi_{t,1}^{\varsigma,j}\right)+\beta\theta\left(\pi_{t+1,2}^{jj}-\sum_l \pi_{t,1}^{\varsigma,l}\pi_{t+1,2}^{lj}\right)\right)$$

and for $j \neq k$:

$$\frac{\partial \pi_{t,1}^{\varsigma,j}}{\partial p^k} = \pi_{t,1}^{\varsigma,j}\left(-\phi\pi_{t,1}^{\varsigma,k}+\beta\theta\left(\pi_{t+1,2}^{jk}-\sum_l \pi_{t,1}^{\varsigma,l}\pi_{t+1,2}^{lk}\right)\right)$$

Under the continued assumption of MIT shocks, this implies that

$$\begin{aligned}\frac{\partial \tilde{E}_{t,1}^{j}}{\partial p_t^k} &= \sum_{\varsigma}\tau^{\varsigma}\frac{\partial \pi_{t,1}^{\varsigma,j}}{\partial p_t^k}\\ &= \begin{cases}\sum_{\varsigma}\tau^{\varsigma}\pi_{t,1}^{\varsigma,j}\left(-\phi\pi_{t,1}^{\varsigma,k}+\beta\theta\left(\pi_{t+1,2}^{jk}-\sum_l \pi_{t,1}^{\varsigma,l}\pi_{t+1,2}^{lk}\right)\right) & \text{if } j \neq k\\ \sum_{\varsigma}\tau^{\varsigma}\pi_{t,1}^{\varsigma,j}\left(\phi\left(1-\pi_{t,1}^{\varsigma,j}\right)+\beta\theta\left(\pi_{t+1,2}^{jj}-\sum_l \pi_{t,1}^{\varsigma,l}\pi_{t+1,2}^{lj}\right)\right) & \text{if } j = k\end{cases}\end{aligned} \tag{40}$$

Elasticities and the Effect on Aggregate Employment Shares

The gradient of aggregate employment is given by:

$$\frac{\partial E_t^j}{\partial p_t^k} = \frac{\partial \tilde{E}_{t,1}}{\partial p_t^k} + \frac{\partial \tilde{E}_{t,2}}{\partial p_t^k}$$

and in terms of elasticities

$$\frac{\partial \ln E_t^j}{\partial p_t^k} = \frac{\tilde{E}_{t,1}^j}{E_t^j}\frac{\partial \ln \tilde{E}_{t,1}^j}{\partial p_t^k} + \frac{\tilde{E}_{t,2}^j}{E_t^j}\frac{\partial \ln \tilde{E}_{t,2}^j}{\partial p_t^k} \tag{41}$$

B.1.5 Empirical Implementation and Relationship to Benchmark Results

We can use the results derived so far to provide an empirical implementation analogous to our main empirical exercises. For example, the right hand side of equation (41) can ultimately be expressed in terms of employment shares, estimable transition probabilities and parameters ϕ, θ and β. One could implement a two-stage estimation strategy for these aggregate supply elasticities. The first stage would be to identify θ using the discretised version of equation (36), employing historical task intensities as instruments for price changes. The second stage would involve estimating ϕ and β by incorporating our first-stage estimate $\hat{\theta}$ into equation (40), utilising the same instrumental variables approach.

An important theoretical consideration is the relationship between this extended model's parameters and those in our benchmark specification. A key insight emerges: when $\beta = 0$ (no forward-looking behavior) and $\phi = \theta$ (identical elasticities across career stages), then (40) collapses to our benchmark gradient. Intuitively, therefore, our benchmark estimates are limited to the extent that we miss long-term forward-looking behaviour and the supply elasticities differ for labour-market entrants and for mature workers. On the second point, again we refer to empirical results in Section 5 showing that the within-career aggregate elasticity is estimated to be similar to the aggregate elasticity of total employment.

B.1.6 Extension 1: Longer-Term Elasticities

We also want to consider medium/long-run elasticities. Here we provide a sketch. In terms of modelling, we therefore assess the effect of period-t shocks on end-of period $t+1$ allocations. For simplicity, assume no shock in $t+1$, i.e. $\boldsymbol{\eta}_{t+1} = 0$. (We could also imagine that the same shocks occur twice, i.e. we have $\boldsymbol{\eta}_{t+1} = \boldsymbol{\eta}_t$.) We can now use the following aggregation:

$$
\begin{aligned}
\tilde{E}^j_{t+1,1}\left(\mathbf{p}_t + \boldsymbol{\eta}_{t+1}\right) &= \sum_{\varsigma} \tau^{\varsigma} \pi^{\varsigma,j}_{t+1,1}\left(\mathbf{p}_t + \boldsymbol{\eta}_{t+1}\right) \\
\implies \tilde{E}^j_{t+1,1}\left(\mathbf{p}_t\right) &= \sum_{\varsigma} \tau^{\varsigma} \pi^{\varsigma,j}_{t+1,1}\left(\mathbf{p}_t\right) \\
\tilde{E}^j_{t+1,2}\left(\mathbf{p}_t + \boldsymbol{\eta}_{t+1}, \mathbf{p}_t\right) &= \sum_{i} \tilde{E}^i_{t,1}\left(\mathbf{p}_t\right) \pi^{ij}_{t+1,2}\left(\mathbf{p}_t + \boldsymbol{\eta}_{t+1}\right) \\
\implies \tilde{E}^j_{t+1,2}\left(\mathbf{p}_t, \mathbf{p}_t\right) &= \sum_{i} \tilde{E}^i_{t,1}\left(\mathbf{p}_t\right) \pi^{ij}_{t+1,2}\left(\mathbf{p}_t\right)
\end{aligned}
$$

Very briefly, as before, gradients for the young are characterised by

$$
\frac{\partial \tilde{E}^j_{t+1,1}}{\partial p^k_t} = \sum_{\varsigma} \tau^{\varsigma} \frac{\partial \pi^{\varsigma,j}_{t+1,1}}{\partial p^k_t}
$$

with little change. But now we have

$$
\begin{aligned}
\frac{\partial \tilde{E}^j_{t+1,2}}{\partial p^k_t} &= \frac{\partial}{\partial p^k_t} \sum_{i} \tilde{E}^i_{t,1}\left(\mathbf{p}_t\right) \pi^{ij}_{t+1,2}\left(\mathbf{p}_t\right) \\
&= \sum_{i} \frac{\partial}{\partial p^k_t} \tilde{E}^i_{t,1} \pi^{ij}_{t+1,2} + \sum_{i} \tilde{E}^i_{t,1} \frac{\partial}{\partial p^k_t} \pi^{ij}_{t+1,2}
\end{aligned}
$$

i.e. the long-run elasticity for the older workers also includes follow-on effects from changing early-career employment. This is now a general function of flows, plus *both* ϕ and θ.

We leave an exploration of these longer-term elasticities for future work.

B.1.7 Extension 2: Characterising the Ex-ante Effect of Aggregate Risk

Equations (37), (38), and (39) show that the effect of wage risk on entry flows is captured by the term $\mathbb{E}_{\eta}\boldsymbol{\eta}'\frac{\partial^2\left(\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})\right)}{\partial\boldsymbol{\eta}^2}\Big|_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}}\boldsymbol{\eta}$ where $\frac{\partial^2\left(\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})\right)}{\partial\boldsymbol{\eta}^2}\Big|_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}}$ is the Hessian of $\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})$ with respect to the vector $\boldsymbol{\eta}$, evaluated at its mean. We can improve on the modelling above by computing this matrix as follows. As is familiar, the first-derivatives of $\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})$ can be written in vector format as:

$$\frac{\partial\left(\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})\right)}{\partial\boldsymbol{\eta}}\Bigg|_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}} = \pi_2^{jk}\begin{pmatrix} -\pi_2^{j1} \\ \dots \\ 1-\pi_2^{jk} \\ \\ -\pi_2^{jN} \end{pmatrix}$$

The matrix of second derivatives is then:

$$\frac{\partial^2\left(\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})\right)}{\partial\boldsymbol{\eta}^2}\Bigg|_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}} = \pi_2^{jk}\begin{pmatrix} \pi_2^{j1}\left(2\pi_2^{j1}-1\right) & \dots & \dots & 2\pi_2^{j1}\pi_2^{jl} & \dots \\ \dots & & & & \\ & & \left(\pi_2^{jk}-1\right)\left(2\pi_2^{jk}-1\right) & & \\ & & & & \\ \dots & & & & \end{pmatrix}$$

We can allow for a general structure on the covariance matrix of $\boldsymbol{\eta}$, but for illustrative purposes assume it follows a distribution $N\left(0,\sigma^2 I_N\right)$ where I_N is the identity matrix of size N. Then

$$\begin{aligned} \mathbb{E}_{\eta}\boldsymbol{\eta}'\frac{\partial^2\left(\pi_2^{jk}(\mathbf{p}+\boldsymbol{\eta})\right)}{\partial\boldsymbol{\eta}^2}\Bigg|_{\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}}\boldsymbol{\eta} &= \sigma^2\pi_2^{jk}\left[\sum_l \pi_2^{jl}\left(2\pi_2^{jl}-1\right)+1-2\pi_2^{jk}\right] \\ \implies \frac{\partial U_2^j(\mathbf{p})}{\partial p^k} &\approx \theta\pi_2^{jk}\left(1+\sigma^2\left[\sum_l \pi_2^{jl}\left(2\pi_2^{jl}-1\right)+1-2\pi_2^{jk}\right]\right) \end{aligned} \tag{42}$$

Absent income risk ($\sigma^2=0$), the interpretation of equation (42) is straightforward and follows that of equation (31). The gradient of expected utility in occupation j w.r.t a change in wages in occupation k is measured by the expected flows from j to k. This captures the probability that the transition from j to k is expected to be chosen.

The effect of risk on the gradient with respect to wages, given by equation (42), depends on whether risk makes the future transition more or less attractive. This effect is non-linear and rich, even with the simple structure on the covariance matrix of $\boldsymbol{\eta}$. At the extremes, when $\pi_2^{jk} = 1$ ($\pi_2^{jk} = 0$), then workers always (never) choose to transition to k, and wage risk has no first-order effect on its attractiveness. For values of π_2^{jk} that are high, but less than 1, then wage risk reduces the chance that the transition from j to k will be the most attractive and so unambiguously attenuates the gradient $\frac{\partial U_2^j(\mathbf{p})}{\partial p^k}$. However, when π_2^{jk} is low but above 0 then, for some configurations of moving costs, wage fluctuations can *increase* the probability of this transition and so increase the gradient $\frac{\partial U_2^j(\mathbf{p})}{\partial p^k}$. An interesting knife-edge case is when flows from j are split evenly between k and a single other occupation l, (e.g. $\pi_2^{jl} = \pi_2^{jk} = \frac{1}{2}$ when only two occupations attract positive flows). Then, when wage risk is equal across k and l, its introduction does not change which occupation is expected to be the most attractive and so again its effect on the gradient $\frac{\partial U_2^j(\mathbf{p})}{\partial p^k}$ is zero.

Although not shown, wage risk itself typically raises the *level* of expected utility. This is because of the linearity of the within-period felicity function and because of the option value of choosing between multiple occupations.

B.2 Aggregation of Subgroups

In this section, we consider the accuracy of our elasticity estimates when we imagine distinct subgroups with different preferences and flows. We thus extend the framework to allow for varying mobility rates during the career across population subgroups.

B.2.1 Model Set-Up

As standard, preferences are given by

$$u_{ij}(\omega) = \theta^s p_j + a_{ij}^s + \epsilon_j(\omega)$$

where now importantly we allow for switching costs, a_{ij}^s, to differ by subgroup s. This includes amenities but might also include constant group-specific occupational wage premia (i.e., wage premia unaffected by shocks).

B.2.2 Elasticities

If end-of-period employment $E_j = \sum_s E_j^s$, then an accounting identity is

$$\frac{\partial \ln E_j(\mathbf{p})}{\partial p_k} = \sum_s \frac{E_j^s}{E_j} \frac{\partial \ln E_j^s(\mathbf{p})}{\partial p_k}$$

Let $\lambda_j^s = \frac{E_j^s}{E_j}$, so that, when evaluated at **p**, we can write

$$\frac{\partial \ln E_j(\mathbf{p})}{\partial p_k} = \sum_s \lambda_j^s \frac{\partial \ln E_j^s(\mathbf{p})}{\partial p_k}$$

where λ_j^s captures the fraction of occupation j coming from group s and

$$\sum_s \lambda_j^s = 1$$

To proceed, we will impose that $\theta^s = \theta$ is common across subgroups.[46] With this in hand, we can write

$$\begin{aligned}
\frac{\partial \ln E_j(\mathbf{p})}{\partial p_k} &= \sum_s \lambda_j^s \theta d_{jk}^s \\
&= \theta \sum_s \lambda_j^s d_{jk}^s \\
&= \theta \begin{cases} -\sum_s \lambda_j^s \sum_i \frac{\tau_i^s \pi_{ij}^s \pi_{ik}^s}{\sum_i \tau_i^s \pi_{ij}^s} & \text{if } k \neq j \\ \sum_s \lambda_j^s \sum_i \frac{\tau_i^s \pi_{ij}^s \left(1-\pi_{ij}^s\right)}{\sum_i \tau_i^s \pi_{ij}^s} & \text{if } k = j \end{cases} \\
&= \theta \begin{cases} -\sum_s \lambda_j^s \sum_i \frac{\tau_i^s \pi_{ij}^s \pi_{ik}^s}{\tau_j^s} & \text{if } k \neq j \\ \sum_s \lambda_j^s \sum_i \frac{\tau_i^s \pi_{ij}^s \left(1-\pi_{ij}^s\right)}{\tau_j^s} & \text{if } k = j \end{cases}
\end{aligned}$$

We calculate d_{jk}^s as standard using steady-state baseline shares τ_j^s which here are group-specific. This allows us to progress to the final line and replace $\sum_i \tau_i^s \pi_{ij}^s$ with τ_j^s.

Now denote our benchmark calculation of the elasticity as $\theta \tilde{d}_{jk}$, as given in Remark 1 and eq (5). In what follows, and in the most relevant cases of subgroups defined either by age or skill-level (education), we show that $\tilde{d}_{jk}$ and $\sum_s \lambda_j^s d_{jk}^s$ yield very similar aggregate elasticities. For age, this is because relative elasticities are highly correlated across subgroups. For skill groups, occupations are quite segmented by education such that aggregate elasticities mostly reflect either high- or low-skilled workers' mobility in each occupation.

[46] Mobility rates can still vary between subgroups due to variation in transition costs (e.g. larger difference between a_{ij}^s and a_{ii}^s among $s = older$).

B.2.3 When elasticities are highly correlated among subgroups

We start by considering different subgroups by age. We categorise people into younger (i.e. those aged 25–39, denoted by $s = y$) and older (i.e. those aged 40–59, denoted by $s = o$). In the computation of $\sum_s \lambda_j^s d_{jk}^s$, we use the average employment during the baseline period 1975–1984 for $\lambda_j^s = \frac{E_j^s}{E_j}$.

We find that d_{jk}^s are highly correlated between age groups (correlation of around 0.8 for both diagonal and non-diagonal elements) and with $\tilde{d}_{jk}$ (around 0.9) in the data. As Figure B.1a shows, levels are lower among older workers, i.e. $d_{jk}^o < d_{jk}^y$, consistent with lower mobility and flexibility as well as higher switching costs (i.e. less time to get the benefits of changing occupations) for older workers.[47] For the sake of argument, we make this extreme and suppose $d_{jk}^o \eta = d_{jk}^y \; \forall j, k$ where $\eta > 1$. We get

$$\sum_s \lambda_j^s d_{jk}^s = d_{jk}^y \left[\lambda_j^y + (1 - \lambda_j^y) \eta \right]$$

which would only deviate from a common elasticity, $corr(d_{jk}^y, \sum_s \lambda_j^s d_{jk}^s) << 1$, if age composition varies a lot between occupations j.

We illustrate that $\tilde{d}_{jk} \approx \sum_s \lambda_j^s d_{jk}^s$ for $s = (y, o)$ in Figure B.1b. Here, we report this relationship for the diagonal elements of matrix D, since these provide the clearest interpretation. Displaying the off-diagonal elements would yield a similar picture. Off-diagonals' skewness, however, would lead to a denser clustering of many points around zero, making the visualisation somewhat cluttered and harder to interpret.

B.2.4 When subgroups work in different occupations

We next consider different subgroups by education. We classify those without an *Abitur* (i.e. German qualification that certifies a person has completed the highest track of secondary school and is ready for university) to be low-skilled ($s = l$), and those with an *Abitur* to be high-skilled ($s = h$). In our data, d_{jk}^s are not too correlated between education levels (correlation is 0.45 for the diagonal elements and 0.35 for the non-diagonal elements). It may be natural, almost mechanical, that workers with different skills and training populate different occupations.

[47] The positive relationship between own-price elasticities for younger and older workers, shown with a quadratic fit in Figure B.1a, would appear even stronger if analysed using rank correlation, emphasising the consistent ordering of elasticities across age groups.

Figure B.1: Elasticities by Age Groups

(a) Younger (aged 25-39) vs Older (aged 40-59) **(b)** $\tilde{d}_{jk}$ vs $\sum_s \lambda_j^s d_{jk}^s$ by Age

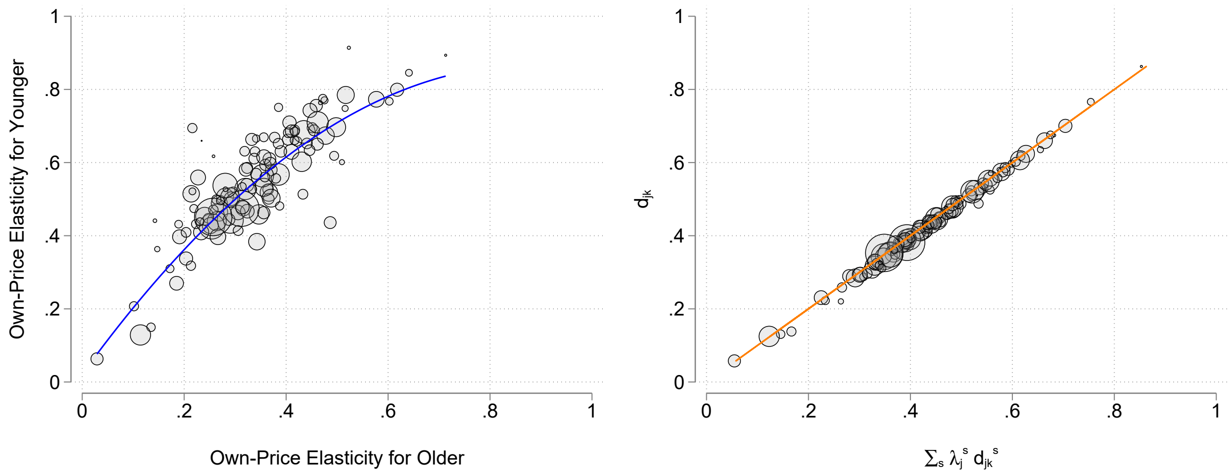

Notes: Figure B.1a shows the relationship (with a quadratic fit) between own-price elasticities for younger (i.e. those aged 25-39) and older (i.e. those aged 40-59) workers. Figure B.1b shows the relationship between $\tilde{d}_{jk}$ and $\sum_s \lambda_j^s d_{jk}^s$ for $s = (y, o)$ also for the diagonal elements of Matrix D (i.e. own-price elasticities). Marker size indicates the baseline employment share (in 1985) in each occupation.

For the sake of argument, suppose that there are two education groups ($s \in \{l, h\}$) and the market is completely segmented:

$$\pi_{ij}^s = \begin{cases} \pi_{ij} & \text{if } s(i) = s(j) \\ 0 & \text{if } s(i) \neq s(j) \end{cases} \quad \tau_i^s = \begin{cases} \tau_i / \lambda^s & \text{if } s(i) = s \\ 0 & \text{if } s(i) \neq s \end{cases} \quad \lambda_j^s = \begin{cases} 1 & \text{if } s(j) = s \\ 0 & \text{if } s(j) \neq s \end{cases}$$

where λ^s is the overall population share of skill group s. We get for $k \neq j$

$$\begin{aligned} -\frac{\partial \ln E_j(\mathbf{p})}{\partial p_k} &= \theta \sum_{s \in \{l,h\}} \lambda_j^s \sum_i \frac{\tau_i^s \pi_{ij}^s \pi_{ik}^s}{\tau_j^s} \\ &= \theta 1\{s(j) = s(k) = l\} \sum_{i, s(i) = l} \frac{\tau_i^l \pi_{ij} \pi_{ik}}{\tau_j^l} + \theta 1\{s(j) = s(k) = h\} \sum_{i, s(i) = h} \frac{\tau_i^h \pi_{ij} \pi_{ik}}{\tau_j^h} \\ &= \theta \sum_{i, s(i) = l} \frac{\tau_i \pi_{ij} \pi_{ik}}{\tau_j} + \theta \sum_{i, s(i) = h} \frac{\tau_i \pi_{ij} \pi_{ik}}{\tau_j} \\ &= \theta \sum_i \frac{\tau_i \pi_{ij} \pi_{ik}}{\tau_j} = \theta \tilde{d}_{jk} \end{aligned} \tag{43}$$

where $1\{s(j) = s(k) = l\}$ is an indicator variable for flows between low-skilled occupations and we used the fact that with segmented occupations $\frac{\tau_i^s}{\tau_j^s} = \frac{\tau_i}{\tau_j}$ if $s(i) = s(j)$.

Equation (43) illustrates the fact that with different relative flow rates among subgroups, and accordingly potentially quite differing d_{jk}^s across s (as Figure B.2a depicts), we may still have that $\tilde{d}_{jk}$ are close to the correct $\sum_s \lambda_j^s d_{jk}^s$. In fact, with fully segmented markets, exactly $\tilde{d}_{jk} = \sum_s \lambda_j^s d_{jk}^s$.

We illustrate that $\tilde{d}_{jk} \approx \sum_s \lambda_j^s d_{jk}^s$ for $s = (l, h)$ in Figure B.2b. This is consistent with the fact that occupations are quite segmented by education group. As before, we report the relationship for the diagonal elements of the elasticity matrix D.

Figure B.2: Elasticities by Education Groups

(a) Low-Skilled vs High-Skilled **(b)** $\tilde{d}_{jk}$ vs $\sum_s \lambda_j^s d_{jk}^s$ by Education

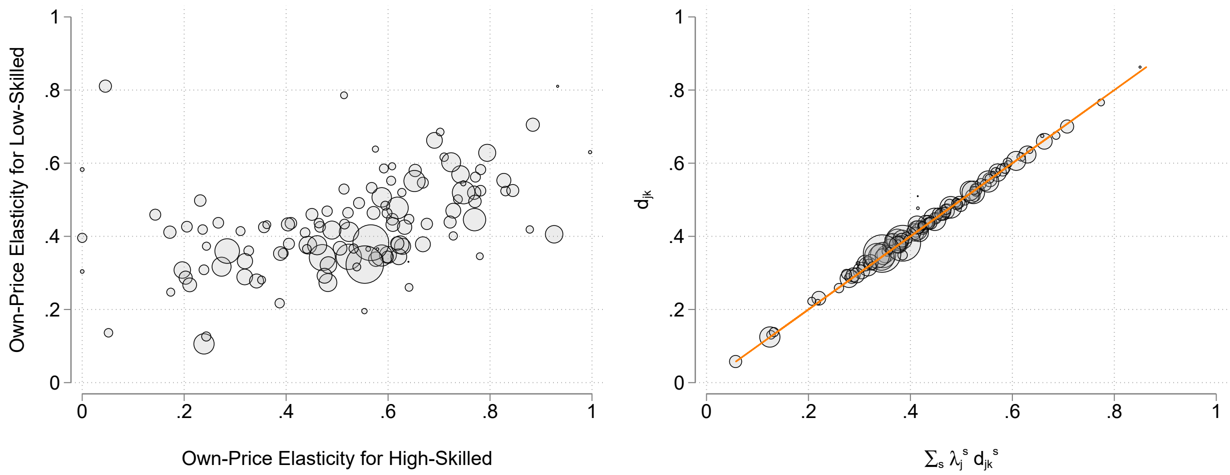

Notes: Figure B.2a shows the relationship between own-price elasticities for low-skill (i.e. those without an *Abitur*) and high-skill (i.e. those with an *Abitur*). Figure B.2b shows the relationship between $\tilde{d}_{jk}$ and $\sum_s \lambda_j^s d_{jk}^s$ for $s = (l, h)$ also for the diagonal elements of Matrix D (i.e. own-price elasticities). Marker size indicates the baseline employment share (in 1985) in each occupation.

B.3 Accounting for Non-Employment Transitions

A driver of heterogeneity in occupational growth that we omit in the main analysis is the extensive margin of employment. This may be important, for instance, in the context of the secular decline of German unemployment in the 2000s.

In line with eq. (1), we interpret indirect utility in M different non-employment states $m \in \{N+1, \ldots, N+M\}$ as containing pecuniary payoffs, transition costs, and idiosyncratic components. While pecuniary payoffs p_m are unobserved, the empirical framework can be extended to control switches to and from different non-employment states.

We start by computing a new elasticity matrix that includes all transitions to and from non-employment states. We can then extend eq. (16) to $N + M$ occupations, with M referring to different non-employment sectors:

$$\Delta e_j \approx \theta \sum_{k=1}^{N+M} d_{jk} \Delta p_k = \theta \sum_{k=1}^{N} d_{jk} \Delta p_k + \sum_{m=N+1}^{N+M} (\theta \Delta p_m) d_{jm} \tag{44}$$

The first summation on the right-hand side represents our standard (own- and cross-occupation) effects, while in the second summation, we explicitly group factors $\theta\Delta p_m$ together. This is to indicate that here we treat d_{jm} as control variables for occupation j's elasticity with respect to non-employment state m. The $\theta\Delta p_m$ coefficient on the respective control represents the combination of pecuniary preferences and changes in non-employment 'prices'.

In what follows, we show the results from these estimations accounting for unemployment transitions (i.e. $M = 1$ non-employment sectors). The source of data regarding unemployment is the Benefit Recipient History (*Leistungsempfängerhistorik* - LeH) of the IAB. The LeH, with data available from 1975, covers periods during which individuals receive earnings replacement benefits from the Federal Employment Agency. This includes unemployment benefits, unemployment assistance as well as maintenance allowances. Column (2) of Table E.5 reports the resulting own-price elasticities accounting for unemployment transitions. Overall, these tend to be slightly larger than d_{jk} in our baseline matrix D, reflecting a relevant amount of transitions in many occupations with unemployment, but they are also clearly correlated (with a correlation bigger than 0.9). As a result, the figure which splits occupations by own-price elasticity accounting for unemployment (not reported) is almost identical to Figure 2b in the main text. Table B.1 then reports the estimation results for eq. (44). The R-squared is higher than in the main text as more of the heterogeneity in employment growth can be explained when allowing for occupations' different elasticities with respect to unemployment. Importantly, the estimated role of own- and cross-occupation effects turn out similar to the main results.

Table B.1: Accounting for Unemployment Transitions.
Determinants of Employment Changes: Own- and Cross-Effects (OLS–IV)

		Dependent Variable: Δe_j		
		Unrestricted Model	Restricted Model	
		(1)	(2)	(3)
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	3.77		
		(0.75)		
heterogeneous own effect:	$(d_{jj} - \bar{d}_{diag})\Delta p_j$	3.04	4.26	6.03
		(0.81)	(0.69)	(1.74)
total cross effect:	$\sum_{k\neq j} d_{jk}\Delta p_k$	2.73		
		(1.12)		
R-squared		0.467	0.451	-
Number of occupations		120	120	120
Estimation method		OLS	OLS	IV
Non-employment controls		Yes	Yes	Yes
F-statistic 1st Stage		-	-	23

Notes: Specifications as in the main text Section 5 other than that regressions now control for occupations' elasticities with unemployment (d_{jm}). The regressors in column (2) are the full $\sum_k d_{jk}\Delta p_k = \mathbf{d}_j\Delta\mathbf{p}$ together with $d_{j,N+1}$ where $k = N+1$ is the unemployment sector. In column (3), these are instrumented by $\mathbf{d}_j V\mathbf{r}$ while controlling for $d_{j,N+1}$ (see eq. (19)).

C Labour Demand and Equilibrium

This section provides detailed derivations behind the discussion in Section 3. In what follows, we present the main features of the demand and supply sides, characterise equilibrium, and show an extension with monopsonistic employers.

C.1 Baseline Competitive Model

We consider an economy-wide constant elasticity of substitution (CES) production technology

$$Y = A\left(\sum_i \beta_i E_i^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \text{ s.t. } \sum \beta_i = 1$$

where i is for occupation, E for employment, β_i are the factor intensities of different occupation inputs and $\sigma > 0$ is the elasticity of substitution across occupations.

The first order conditions yield, for all i,

$$\beta_i E_i^{\frac{-1}{\sigma}} A\left(\sum_i \beta_i E_i^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}-1} = \mathfrak{p}_i$$

where $\mathfrak{p}_i$ is the wage (the price of labour) level for occupation i and the price of the output good is normalised to 1.

To begin, consider demands relative to occupation N:

$$\begin{aligned}\tilde{E}_i &\equiv \ln \tfrac{E_i}{E_N} &= \ln\left(\frac{\beta_i}{\beta_N}\frac{\mathfrak{p}_N}{\mathfrak{p}_i}\right)^{\sigma} \\ & &= \ln\left(\beta_{-i}\frac{\beta_i}{1-\beta_i}\frac{1}{\tilde{\mathfrak{p}}_i}\right)^{\sigma}\end{aligned}$$

where $\tilde{\mathfrak{p}}_i \equiv \frac{\mathfrak{p}_i}{\mathfrak{p}_N}$ and $\beta_{-i} \equiv \frac{1-\beta_i}{\beta_N} = \frac{\sum_{j\neq i}\beta_j}{\beta_N}$. In what follows, we will consider incremental changes to $\ln\frac{\beta_j}{1-\beta_j}$ with proportionate off-setting changes to β_k for $k \neq j$.

It is worth noting that $\frac{d\ln\frac{\beta_i}{\beta_N}}{d\ln\frac{\beta_i}{1-\beta_i}} = \frac{d\ln\beta_{-i}\frac{\beta_i}{1-\beta_i}}{d\ln\frac{\beta_i}{1-\beta_i}} = 1$. On the other hand, $\frac{d\ln\frac{\beta_i}{\beta_N}}{d\ln\frac{\beta_j}{1-\beta_j}} = 0$ because proportional changes to β_i and β_N are equal and offsetting. In more compact notation, we can therefore write

$$\tilde{E}_i^d\left(\tilde{p}_i\left(\mathbf{b},\mathbf{s}\right),\tilde{\beta}_i\right) = \ln\left(\tilde{\beta}_i\frac{1}{\tilde{\mathfrak{p}}_i}\right)^{\sigma} \tag{45}$$

where $\tilde{p}_i$ is the log of $\tilde{\mathfrak{p}}_i$, $\mathbf{b}$ is the $(N-1)$ vector of relative productivities (i.e. demand shifters $\left(\ln \frac{\beta_j}{1-\beta_j}\right)$), $\mathbf{s}$ is a vector of supply shifters that do not directly affect demand, and $\tilde{\beta}_i = \frac{\beta_i}{\beta_N}$. Note that relative demand for employment in occupation i depends on the relative price in that occupation *only*.

In fact, we are interested in log employment shares $e_i = \ln \frac{E_i}{\sum_j E_j} = \ln \frac{E_i}{\bar{E}}$. In this case, demands depend on productivities and prices of other occupations. We are interested in perturbations around the steady state, so in keeping with the rest of the paper, we denote steady-state share of occupation i by τ_i. This gives a demand curve $e_i^d\left(\langle \tilde{p}\left(\mathbf{b}, \mathbf{s}\right)\rangle, \mathbf{b}\right)$, which is a function of all prices and demand shifters.

To calculate derivatives, first note that, around the steady state:

$$\frac{\partial e_j^d}{\partial p_i}|_{p_{k \neq i}} = -\frac{\tau_i}{1-\tau_i}\frac{\partial e_i^d}{\partial p_i}|_{p_{k \neq i}}$$

i.e. given a change to p_i, and holding fixed all other prices (made explicit by the notation $|_{p_{k\neq i}}$), then adding up ensures this identity, because all other occupations are equally proportionately offset.[48] Therefore, we have that:

$$\begin{aligned}
\frac{\partial e_i^d}{\partial p_i} &= \frac{\partial \ln \frac{E_i}{E_N}}{\partial p_i} + \frac{\partial \ln \frac{E_N}{\bar{E}}}{\partial p_i} \\
&= \frac{\partial \ln \frac{E_i}{E_N}}{\partial p_i} + \frac{\partial e_N^d}{\partial p_i} \\
&= -\sigma - \frac{\tau_i}{1-\tau_i}\frac{\partial e_i^d}{\partial p_i} \\
\implies \frac{\partial e_i^d}{\partial p_i} &= -\left(1-\tau_i\right)\sigma
\end{aligned}$$

This also implies that for $j \neq i$:

$$\begin{aligned}
\frac{\partial e_i^d}{\partial p_j} &= -\frac{\tau_j}{1-\tau_j}\frac{\partial e_j^d}{\partial p_j} \\
&= \tau_j \sigma
\end{aligned}$$

Together, these are result (11) in the main text. A similar logic implies that $\frac{\partial e_i^d}{\partial \ln \frac{\beta_j}{1-\beta_j}}$ follows a similar structure.

[48] Note that adding up requires $\sum_k \frac{\partial e_k^d}{\partial p_i} E_k = 0$, which implies $\frac{\partial e_i^d}{\partial p_i} E_i + \sum_{k \neq i} \frac{\partial e_k^d}{\partial p_i} E_k = 0$. Noting that a property of CES demands given by eq. (45) are that $\frac{\partial e_k^d}{\partial p_i} = \frac{\partial e_l^d}{\partial p_i} \equiv \frac{\partial e_{-i}^d}{\partial p_i}$ for $k, l \neq i$, then we have that $\frac{\partial e_i^d}{\partial p_i} E_i + \frac{\partial e_{-i}^d}{\partial p_i} \sum_{k \neq i} E_k = 0 \implies \frac{\partial e_i^d}{\partial p_i} e_i + \frac{\partial e_{-i}^d}{\partial p_i} \sum_{k \neq i} e_k = 0$. Rearranging and using τ_i give the result.

We therefore have a demand function $e_i^d\left(\langle p\left(\mathbf{b},\mathbf{s}\right)\rangle,\mathbf{b}\right)$ with partial derivatives for prices given by elements of the matrix $\sigma\left(W-I\right)$, with rank $N-1$, where I is the identity matrix, and W is the matrix of employment shares, as defined in Appendix A.3. The matrix of derivatives with respect to demand shifters is given by $\sigma\left(I-W\right)$, equally of rank $N-1$.

C.1.1 Labour Supply

As extensively discussed in the main text, we have that $\frac{\partial e_j^s}{\partial p_k}=\theta d_{jk}$. The matrix of supply derivatives is therefore given by θD, similarly of rank $N-1$.

We have some flexibility in defining the effect of supply shifters, as long as they satisfy adding up, i.e. that $\sum_i \frac{\partial e_i^s}{\partial s_j}\tau_i=0$. We can satisfy this by letting $\frac{\partial e_i^s}{\partial s_j}\equiv-\tau_j$ for $i\neq j$ and $\frac{\partial e_j^s}{\partial s_j}\equiv 1-\tau_j$. Then $\sum_i \frac{\partial e_i^s}{\partial s_j}\tau_i=\left(1-\tau_j\right)\tau_j-\sum_{i\neq j}\tau_j\tau_i=\tau_j\left(1-\tau_j-\sum_{i\neq j}\tau_i\right)=0$. The matrix of derivatives with respect to supply shifters is therefore given by $I-W$.

C.1.2 Equilibrium Characterisation

Similarly to before, we can write

$$e_i\left(\mathbf{b},\mathbf{s}\right)=e_i^s\left(\langle p\left(\mathbf{b},\mathbf{s}\right)\rangle,\mathbf{s}\right)=e_i^d\left(\langle p\left(\mathbf{b},\mathbf{s}\right)\rangle,\mathbf{b}\right) \tag{46}$$

where both supply and demand curves depend on the full system of prices.

In what follows, for ease of exposition, it is useful to define the following matrices for gradients of equilibrium quantities $\{E_j\}$ and prices $\{p_j\}$.

Notation	Typical element
Ξ	$\frac{de_i}{d\left(\ln\frac{\beta_j}{1-\beta_j}\right)}$
Γ	$\frac{de_i}{ds_j}$
V	$\frac{dp_i}{d\left(\ln\frac{\beta_j}{1-\beta_j}\right)}$
S	$\frac{dp_i}{ds_j}$

Solving for Price Gradients using $e_i^s\left(\right)=e_i^d\left(\right)$

Differentiating $e_i^s\left(\right)=e_i^d\left(\right)$ from eq. (46) with respect to $\ln\frac{\beta_j}{1-\beta_j}$ we obtain:

$$\sum_k\frac{\partial e_i^s}{\partial p_k}\frac{\partial p_k}{\partial\left(\ln\frac{\beta_j}{1-\beta_j}\right)}=\sum_k\frac{\partial e_i^d}{\partial p_k}\frac{\partial p_k}{\partial\left(\ln\frac{\beta_j}{1-\beta_j}\right)}+\frac{\partial e_i^d}{\partial\ln\frac{\beta_j}{1-\beta_j}} \tag{47}$$

Expressing this in matrix notation gives

$$
\begin{aligned}
\theta D V &= \sigma\left(W-I\right)V+\sigma\left(I-W\right) \\
\implies \left(\theta D+\sigma\left(I-W\right)\right)V &= \sigma\left(I-W\right) \qquad (48)
\end{aligned}
$$

where V is a matrix with i, jth element $\frac{\partial p_i}{\partial\left(\ln\frac{\beta_j}{1-\beta_j}\right)}$ that we wish to solve.

At this point, we notice that $\left(\theta D+\sigma\left(I-W\right)\right)$ has rank $N-1$. However, we can also notice that $\left(I-W\right)$ is the de-meaning operator, such that for vector x, then $\left(I-W\right)x = x-1_N\sum_i \tau_i x_i$, where 1_N is a column vector of ones. Therefore, we can solve eq. (48) as long as we make the appropriate normalisation. Specifically, we define price gradients such that $\sum_i \tau_i \frac{\partial p_i}{\partial\left(\ln\frac{\beta_j}{1-\beta_j}\right)} = 0$, i.e. the weighted price gradient is 0.

Recall that this normalisation is without loss of generality because the model is invariant to additive shifts in prices. In this case, we can solve for V as

$$
V = \left(\frac{\theta}{\sigma}D+I\right)^{-1}\left(I-W\right) \qquad (49)
$$

which in fact guarantees the normalisation by construction.

Next, we consider gradients with respect to supply shifters. Differentiating with respect to s_j we obtain:

$$
\begin{aligned}
\sum_k \frac{\partial e_i^s}{\partial p_k}\frac{\partial p_k}{\partial s_j}+\frac{\partial e_i^s}{\partial s_j} &= \sum_k \frac{\partial e_i^d}{\partial p_k}\frac{\partial p_k}{\partial s_j} \\
\implies \theta D S + I - W &= \sigma\left(W-I\right)S \\
\implies \left(\theta D+\sigma\left(I-W\right)\right)S &= -\left(I-W\right)
\end{aligned}
$$

Similarly to above, we can solve for S using a normalisation of price gradients with respect to a supply shock. That is, setting $\sum_i \tau_i \frac{\partial p_i}{\partial s_j} = 0$ and again without loss of generality, we obtain:

$$
\begin{aligned}
S &= -\left(\theta D+\sigma I\right)^{-1}\left(I-W\right) \qquad (50) \\
&= -\frac{1}{\sigma}V
\end{aligned}
$$

Solving for Quantity Gradients using $e_i\left(\right) = e_i^d\left(\right)$ and $e_i\left(\right) = e_i^s\left(\right)$

Differentiating the identity $e_i\left(\mathbf{b}, \mathbf{s}\right) = e_i^d\left(\left\langle p\left(\mathbf{b}, \mathbf{s}\right)\right\rangle, \mathbf{b}\right)$ w.r.t. s_j we get

$$\begin{aligned} \frac{\partial e_i}{\partial s_j} &= \sum_k \frac{\partial e_i^d}{\partial p_k} \frac{\partial p_k}{\partial s_j} \\ \implies \Gamma &= -\sigma\left(I - W\right) S = -\sigma S = V \end{aligned}$$

and then differentiating the identity $e_i\left(\mathbf{b}, \mathbf{s}\right) = e_i^s\left(\left\langle p\left(\mathbf{b}, \mathbf{s}\right)\right\rangle, \mathbf{b}\right)$ w.r.t. $\ln \frac{\beta_j}{1-\beta_j}$ we get

$$\frac{de_i}{d\left(\ln \frac{\beta_j}{1-\beta_j}\right)} = \sum_k \frac{\partial e_i^s}{\partial p_k} \frac{\partial p_k}{\partial\left(\ln \frac{\beta_j}{1-\beta_j}\right)}$$

which provides the matrix equation

$$\Xi = \theta D V$$

C.1.3 Observed Changes

Let $\Delta\mathbf{e}$ be the vector of observed changes in labour shares, with ith element, Δe_i. Similarly, let $\Delta\mathbf{b}$ be the vector of productivity (or demand) shifts, $\Delta\mathbf{s}$ the vector of supply shifts, and $\Delta\mathbf{p}$ be the change in prices. Then we have that

$$\begin{aligned} \Delta\mathbf{p} &\approx V\Delta\mathbf{b} + S\Delta\mathbf{s} \\ &= V\Delta\mathbf{b} - \frac{1}{\sigma} V\Delta\mathbf{s} \end{aligned} \tag{51}$$

and

$$\begin{aligned} \Delta\mathbf{e} &\approx \Xi\Delta\mathbf{b} + \Gamma\Delta\mathbf{s} \\ &= \theta D V\Delta\mathbf{b} - \sigma S\Delta\mathbf{s} \\ &= \theta D V\Delta\mathbf{b} + V\Delta\mathbf{s} \end{aligned} \tag{52}$$

These expressions, corresponding to eq. (13) and eq. (14) in the main text, describe changes to labour shares and prices in terms of demand and supply shocks, price elasticities, and model parameters θ and σ.

Finally, expressions (51) and (52) also inform the regression framework. From eq. (51), note that $\theta D\Delta\mathbf{p} = \theta D V\Delta\mathbf{b} + \theta D S\Delta\mathbf{s}$. Using this to substitute $\Delta\mathbf{b}$ out of eq. (52) yields:

$$\implies \Delta\mathbf{e} \approx \theta D\Delta\mathbf{p} - \theta D S\Delta\mathbf{s} - \sigma S\Delta\mathbf{s}$$

$$
\begin{aligned}
&= \theta D \Delta \mathbf{p} - (\theta D + \sigma) S \Delta \mathbf{s} \\
&= \theta D \Delta \mathbf{p} - (-(I - W)) \Delta \mathbf{s} \\
&= \theta D \Delta \mathbf{p} + \Delta \mathbf{s} \qquad (53)
\end{aligned}
$$

where the last line follows from the penultimate line because the vector of supply shocks is defined to be suitably normalised.

C.2 Alternative: Equilibrium with Monopsonistic Employers

The benchmark model presented above is characterised by competitive behaviour of labour supply and demand, which seems plausible when studying occupations. In this section, we show that our main results do not change in versions of the model where employers behave imperfectly competitively. The underlying intuition is that time-constant markdowns (or markups), though differing across occupations, do not appear in relative log (percentage) changes of demand or supply over time.

C.2.1 Setting

Aggregate production continues to be according to a CES function under perfect competition and where the price of the final good is normalised to 1.

$$
Y = A \left(\sum_j \beta_j Y_j^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \text{ s.t. } \sum \beta_j = 1, \sigma > 1 \qquad (54)
$$

This leads to the price of occupation j's output being equal to its marginal product:

$$
\chi_j = \beta_j Y_j^{-\frac{1}{\sigma}} A \left(\sum_j \beta_j Y_j^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}-1} \qquad (55)
$$

At the level of the intermediaries, we can impose more structure on the model to allow for different market structures. Suppose there is a unit measure of atomistic employers indexed by l that are homogenous, from a labour supply and demand perspective, within each occupation j. In the benchmark competitive model, we assume that workers receive identical amenity draws from each of the firms, who therefore have no market power.

To introduce market power, suppose instead that workers obtain job offers from a *single* firm within each occupation. For concreteness, we might imagine that the economy is regionally segregated, such that there is a single firm for each of the N occupations within each worker's district, and workers must work in their own district. In this case, employers know that they face an upward-sloping supply curve, as workers' only outside options are to choose an alternative occupation. As is standard, we assume employers cannot discriminate between workers (here according to their origin occupation). Then labour supply elasticities within occupations remain the same as in the main text. That is, the labour supply elasticity to each employer in j is a constant $\varepsilon^S_{lj} = \varepsilon^S_j = \theta d_{jj}$, which can also be interpreted as the weighted average cross-elasticity facing occupation j (see discussion in Appendix A.2).

We also continue to assume a linear production function in each occupation such that occupational output is equal to occupational employment i.e. $Y_j = E_j = \int E_{lj} \, dF(l)$. Labour demand elasticities also become homogeneous as $\varepsilon^D_{lj} = \varepsilon^D_j$.[49] In this case, analysing individual firm behaviour is equivalent to analysing behaviour at the occupation level, which we do for convenience from now on.

C.2.2 Equilibria

Profits for each occupation's representative employer become

$$\pi_j = \chi_j E_j - \mathfrak{p}_j E_j \tag{56}$$

which are maximised with respect to E_j and where $\mathfrak{p}_j$ are wages in levels. Before examining monopsony, it is useful to return to the case of perfectly competitive markets.

Competitive behaviour in product and factor markets

With their decision, the employer influences neither prices and we get as first-order condition (FOC):

$$\chi_j = \mathfrak{p}_j = \beta_j E_j^{-\frac{1}{\sigma}} A \left(\sum_j \beta_j E_j^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}-1}$$

which is exactly the case of the main text and Appendix C.1.

We now consider the setting of interest.

[49] If occupation j is sufficiently small, approximately $\varepsilon^D_j \approx -\sigma$; otherwise, there will be an attenuating size-adjustment term. In Appendix C.1, we have shown that the elasticity of occupation j's employment share with respect to its (relative) price is exactly $-(1-\tau_j)\sigma$. What matters in the following is simply that ε^D_j, and for that matter ε^S_j, are not affected by (marginal) changes to occupational productivities or prices.

Monopsony in the labour market

The employer takes into account how its demand affects wages, such that the FOC to the profit function (56) becomes:

$$\chi_j = \mathfrak{p}_j + \frac{\partial \mathfrak{p}_j}{\partial E_j} E_j = \mathfrak{p}_j \left(\frac{\varepsilon_j^S + 1}{\varepsilon_j^S} \right) = \frac{\mathfrak{p}_j}{\mu_j^S},$$

where $\mu_j^S < 1$ is the wage markdown, relative to labour's value marginal product, in occupation j. Other than that, we can still take condition (55) to get log labour demand demand relative to occupation N:

$$\tilde{E}_j^d = \ln \left(\tilde{\beta}_j \frac{\tilde{\mu}_j^S}{\tilde{\mathfrak{p}}_j} \right)^{\sigma} \tag{57}$$

which is equation (32) in Appendix C.1 with just an additional term reflecting relative markdowns $\tilde{\mu}_j^S = \frac{\mu_j^S}{\mu_N^S}$.

From now on, the analysis of price and labour demand shocks gives exactly the same results as in Appendix C.1. This is because equation (57) has the same structure as before, markdowns are constant, and we showed that derivatives of resulting employment shares e_j^d with respect to p_i or $\ln \left(\frac{\beta_i}{1-\beta_i} \right)$ take a simple form. In short, the changes in employment shares do not depend on $\tilde{\mu}_j^S$.[50] We therefore obtain the same derivates, especially the spillover matrix V, and analysis of the equilibrium effects of shocks as before.

Concluding discussion – Imperfect competition in both markets

As a final point, we could alter the set-up up to allow for monopolist behaviour by intermediate producers in the product market, such that occupational employers also take into account the effect of their decisions on output prices. Again, the equilibrium response to shocks remains unchanged. Now the FOC for the intermediate producer becomes

$$\chi_j + \frac{\partial \chi_j}{\partial E_j} E_j = \chi_j \left(\frac{\varepsilon_j^D - 1}{\varepsilon_j^D} \right) = \frac{\mathfrak{p}_j}{\mu_j^S} \quad \text{or} \quad \frac{\chi_j}{\mathfrak{p}_j} = \frac{\mu_j^D}{\mu_j^S},$$

where $\mu_j^D = \frac{\varepsilon_j^D}{\varepsilon_j^D - 1} > 1$ is j's markup over marginal cost and we have the well-known double-marginalisation of output- relative to input price of a monopolist-monopsonist.

[50] We could study effects on demanded employment shares, e_j^d, of changes in markdowns (i.e. occupations' labour supply elasticities). These derivatives would also have the form $\sigma(W - I)$.

Still, equation (57) retains the same structure, $\tilde{E}_j^d = \ln\left(\tilde{\beta}_j \frac{1}{\tilde{\mathfrak{p}}_j} \frac{\tilde{\mu}_j^S}{\tilde{\mu}_j^D}\right)^\sigma$, and the equilibrium effects of changes in demand or prices are unaffected by these markups/markdowns in level terms.

D Data Appendix

This section presents a detailed presentation of the data. First, we discuss the SIAB data and outline the procedures for sample selection and wage imputation. We then review the data on tasks and occupational characteristics.

D.1 The SIAB Data

We use the Sample of Integrated Labour Market Biographies (*Stichprobe der Integrierten Arbeitsmarktbiographien*, Frodermann et al., 2021) for our analyses.[51] The SIAB is a 2% sample of the population of the Integrated Employment Biographies (IEB) provided by the Institute for Employment Research (*Institut für Arbeitsmarkt- und Berufsforschung* – IAB). It includes employees covered by social security, marginal part-time workers (after 1999), unemployment benefit recipients, individuals who are officially registered as job-seeking, and individuals who are participating in programs of active labour market policies. It is possible to track the employment status of a person exact to the day. The source of data regarding employment is the Employee History (*Beschäftigtenhistorik* - BeH) of the IAB. The BeH covers all white- and blue-collar workers as well as apprentices as long as they are not exempt from social security contributions. It excludes civil servants, self-employed people, regular students, and individuals performing military service.

The SIAB data contains an individual's full employment history, including a consistent-over-time occupational classifier (between 1975-2010 and then again from 2012 onwards), the corresponding nominal daily wage, and socio-demographic variables such as age, gender, or level of education. Data are available in a spell structure, making it possible to observe the same person at several employers within a year. In a few cases, these spells overlap when workers have multiple employment contracts at a time. We transform the spell structure into a yearly panel by identifying the longest spell within a given year and deleting all the remaining spells (following Böhm et al., 2024).

[51] Access to the data is subject to signing a contract with the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB).

D.1.1 Sample Selection and Variable Description

We consider two different samples in this paper: the historical sample (1975-2010, Sections 4-6) and the prediction sample (from 2012 onwards, Section 7).[52] To work with a homogeneous sample throughout, the historical sample is restricted to Western German full-time male workers aged 25–59. Since the level and structure of wages differ substantially between former East and West Germany, we drop workers who were ever employed in Eastern Germany. Our focus on full-time jobs is driven by the absence of data on hours worked. The restriction on men is primarily due to the old occupation classification, which was devised with male employment in mind (Paulus et al., 2013). Excluding younger workers, we ensure the vast majority conclude their formal education by the time they enter the sample. Besides, we stop relatively early (at age 59) because early retirement programs were common in Germany, particularly in the late 1970s and the 1980s. The prediction sample relaxes these restrictions and includes both men and women, Eastern and Western Germans, as well as those working part-time and full-time.

We exclude workers with wages below the limit for which social security contributions have to be paid, mainly workers in marginal jobs (also known as mini-jobs). These jobs were not subject to social security taxation before 1999. After the first reform in 1999, the tax-free wage threshold was fixed during the period 1999 to 2003 at 325 euros per month. In 2003, the range of exempted earnings was expanded up to 400 euros, effective until 2012. The minimum threshold for mini-jobbers increased in 2013 from 400 to 450 euros per month. Approximately 10% of observations are affected by this restriction. We drop wage spells of workers whose last spell is in apprenticeship training as the first wage after apprenticeship is often a mixture between new wage and apprenticeship wage (this only affects 0.48% of the sample). We also drop all spells of workers who are always foreign workers (less than 5% of observations).[53] Finally, workers without information on their occupation or wages are dropped from the analysis.

Occupation classification. For the historical sample analysis, we use the 120 three-digit occupations from the SIAB's Scientific Use File as our main units of analysis. These occupations are consistently coded during the period 1975–2010 (from the KldB 1988 classification system) and listed in Table E.5. In Appendix Table A.1, we also consider occupations at the 1-digit level and aggregate them into four broad groups following the literature (Acemoglu & Autor, 2011; Böhm et al., 2024). These are (1) managers, professionals, and technicians (Mgr-Prof-Tech), (2) sales and office workers (Sales-Office), (3) production

[52] Recall that occupations in the SIAB are consistently coded, with no removals or additions, during the period 1975-2010. After this, a structural break occurs and a new occupational classification is introduced, which is fully consistent again from 2012 onwards.

[53] Workers who are classified as German at some point but foreign at another are not dropped.

workers, operators, and craftsmen (Prod-Oper-Crafts), and (4) workers in services and care occupations (Serv-Care). For the prediction exercise, we use the 126 three-digit occupations based on the new classification of occupations introduced in 2010 (the KldB 2010 classification system). This classification captures the trends towards more employment and job differentiation in areas such as health, education, and information technologies.

Wages. The available wage variable is the employee's gross daily nominal wage in euros. It is calculated from the fixed-period wages reported by the employer and the duration of the original notification period in calendar days. Despite being accurately measured as the employer can be punished for incorrect reporting, two major drawbacks are of special relevance to our analysis. First, due to a cap on social security contributions, wages are right-censored. As is common in administrative data sources, earnings above the upper earnings limit for statutory pension insurance are only reported up to this limit. The upper earnings limit for statutory pension insurance differs from year to year as well as between Eastern and Western Germany, where the decisive factor is the location of the establishment. Second, the income components being subject to social security tax were extended in 1984. Before that, one-time payments such as bonuses were not included in the daily wage benefit measure. We discuss how we deal with these two issues below. Finally, to ensure comparability across years, wages are deflated by the Consumer Price Index reported in the Federal Statistical Office of Germany, with 2010 as the base year.

D.1.2 Imputation of Right-Censored Wages

The SIAB data is based on process data used to calculate retirement pensions and unemployment insurance benefits, implying the wage information is top-coded and only relevant up to the Social Security contribution ceiling. While this feature only affects approximately 8.5% of observations on average across years in our historical sample (aged 25–59, full-time, excluding marginal workers), the proportion of censored observations differs across subgroups. By gender, top-coded wages amount to roughly 11% for men and 3.3% for women. Differences are also substantial by education groups. Whereas only 1.1% of the spells of individuals who enter the labour market without post-secondary education are affected by top-coding, the share of right-censored wages increases to 5.2%, 9.4%, and 30.8% for those who completed vocational education and training, an *Abitur*, and a university degree, respectively. The share of top-coded wages also increases over the life cycle. While censoring only affects less than 2% of observations for those aged 25-29, the fraction of top-coded wages rises to more than 11% for those older than 40.

To impute top-coded wages, we follow Dustmann et al. (2009) and Card et al. (2013).[54] We first define age-education cells based on seven age groups (with 5-year intervals; 25–29; 30–34; 35–39; 40–44; 45–49; 50–54; 55–59) and four education groups (as described above). Within each of these cells (and thereby allowing a different variance for each education and age group), we estimate Tobit wage equations separately by year, gender, and East-West Germany. We predict the upper tail of the wage distribution including controls for age (quadratic), tenure (quadratic), a part-time dummy, as well as interactions between age (quadratic) and the different education groups. To control for worker fixed effects, we construct the mean of an individual's log wage in other years, the fraction of censored wages in other years, and a dummy variable if the person was only observed once in her life.[55] We use the predicted values $X'\hat{\beta}$ from the Tobit regressions together with the estimated standard deviation $\hat{\sigma}$ to impute the censored log wages y^c as follows:

$$y^c = X'\hat{\beta} + \hat{\sigma}\Phi^{-1}[k + u(1-k)]$$

where Φ is the standard normal density function, u is a random draw from a uniform distribution ranging between zero and one, $k = \Phi[(c - X'\hat{\beta})/\hat{\sigma}]$ and c is the censoring point, which differs by year and East–West Germany. See Gartner (2005) for details.[56] In a very few cases ($< 0.001\%$), imputed wages are exceedingly high. As a minor adjustment, we limit imputed wages to ten times the 99th percentile of the latent wage distribution.

D.1.3 The Structural Wage Break 1983/1984

The income components being subject to the social security tax were extended in Germany in 1984 (for further details, see Bender et al. (1996) and Steiner & Wagner (1998)). Before 1984, one-time payments, such as bonuses, were not included in the daily wage benefit measure. Starting in 1984, these variable parts of the wage were included. We follow Fitzenberger (1999) and Dustmann et al. (2009) and deal with this structural break by correcting wages prior to 1984 upwards. The correction is based on the idea that higher quantiles appear to be more affected by the structural break than lower quantiles, as higher percentiles are likely to receive higher bonuses. To this end, we estimate locally weighted regressions, separately for men and women, of the wage ratio between 1982 and 1983 (i.e. before the break), and between 1983 and 1984 (i.e. after the break) on the wage

[54] To ensure that all censored wages are covered in the imputation procedure, we mark all observations with wages four euro below the assessment ceiling as in Dauth & Eppelsheimer (2020).

[55] For those observed only once, the mean wage and mean censoring indicator are set to sample means.

[56] Dustmann et al. (2009) consider different imputation methods, such as restricting the variance to be the same across all education and age groups, or assuming the upper tail of the wage distribution follows a Pareto distribution. They conclude that the imputation method that assumes that the error term is normally distributed with a different variance by age and education works better than the other imputation methods. This method is also chosen in more recent papers such as Cortes et al. (2024) and Böhm et al. (2024).

percentiles in 1983 and 1984, respectively. The correction factor is then computed as the difference between the predicted, smoothed values from the two wage ratio regressions. In a way similar to that of Dustmann et al. (2009), to account for differential overall wage growth between the periods from 1982 to 1983 and from 1983 to 1984, we subtract from the correction factor the smoothed value of the wage ratio in 1983, averaged between the second and fortieth quantiles. Finally, wages prior to 1984 are corrected by multiplying them by 1 plus the correction factor. After this, some wages are corrected above the censoring limit. Dustmann et al. (2009) reset these wages back to the censoring limit and impute them in the same way they imputed wages that were above the limit anyway. Here we follow Böhm et al. (2024) and do not reset wages back to the censoring limit if they were corrected above the limit but leave them at their break corrected values.

D.2 Data on Tasks and Occupational Characteristics

We use the Qualifications and Career Surveys (QCS, Hall et al., 2012), conducted by the Federal Institute for Vocational Education and Training (BiBB), to obtain information on tasks performed in occupations. The QCS, which have been previously used, e.g. by Spitz-Oener (2006); Antonczyk et al. (2009); Gathmann & Schönberg (2010), are representative cross-sectional surveys with 20,000–35,000 individuals in each wave who respond about the tasks required in their occupations. These include, for example, how often they repair objects, how often they perform fraction calculus, or how often they have to persuade co-workers. We classify questions as representing either analytical, interactive, routine, or manual tasks and assign a value of 0, 1/3, or 1, depending on whether the answer is 'never', 'sometimes', or 'frequently'. We pool the QCS waves in 1979 and 1985/1986 to compute task intensities across occupations by averaging over all the responses. We use this information to study how task intensity relates to our price elasticity measures, and instrument demand changes across occupations over the period 1985-2010.

Task distance. To measure the distance between occupations in the task space (reflecting the degree of dissimilarity in the mix of tasks), we follow Cortes & Gallipoli (2018) and use the angular separation (correlation) of the observable vectors x_j and x_k:[57]

$$\text{AngSep}_{jk} = \frac{\sum_{a=1}^{A} \left(x_{aj} \cdot x_{ak}\right)}{\left[\sum_{a=1}^{A} \left(x_{aj}\right)^2 \cdot \sum_{a=1}^{A} \left(x_{ak}\right)^2\right]^{\frac{1}{2}}} \tag{58}$$

[57]The angular separation is the cosine angle between the occupations' vectors in the task space.

where x_{aj} is the intensity of task dimension a in occupation j and A is the total number of dimensions being considered (analytical, routine, and manual). We transform this to a distance measure $dist_{jk}$ that is increasing in dissimilarity:

$$dist_{jk} = \frac{1}{2}(1 - AngSep_{jk})$$

The measure varies between zero and one; it will be closer to zero the more two occupations overlap in their skill requirements. The mean task distance between occupations in our data is 0.5, with a standard deviation of 0.29 (see Table E.1). The most distant possible move is between an economic and social scientist and a carpenter. Examples of pairs of occupations with low distance measures are between a sheet metal worker and a tile setter, or between a glass processor and a plastic processor.

Occupational licensing. To obtain measures of occupational licensing, we use the indicators for *standardised certification* requirements and degree of *regulation* developed by Vicari (2014). These indicators are based on BERUFENET, the online career information portal provided by the German Federal Employment Agency – a rich job title database similar to the US O*NET. They are calculated by categorising very narrow occupations (8-digit) based on the presence or absence, under federal or state law, of standardised training certificates required for professional activities. This is done in three steps. First, each 8-digit occupation is assigned a value of 0 or 1 based on whether the access to the occupational activity is linked to standardised credentials. Second, each occupation is merged with the feature 'regulation', i.e. whether legal and administrative regulations exist for an occupation and whether a specific qualification is necessary to practice it. Finally, the indicator 'standardised certification' uses both pieces of information about the standardisation of the credentials and regulation. These 0-1 values are finally aggregated at the 3-digit occupational classification (i.e. the 120 occupations used in our historical analysis), weighted by the number of individuals employed in each occupation. Intuitively, the degree of regulation indicates whether legal and administrative regulations exist which bind the access to and practice of the occupation, including the necessity of holding a specific title as proof of competence. The occupational certification further includes whether access to exercising the professional activity is linked to a standardised training credential. These indicators are constructed as a metric value between 0 and 1, with the indicator increasing in the degree of certification and regulation.

E Further Analysis in the Historical Period

This section presents further analysis to complement the one in Sections 4-5.

E.1 Descriptive Statistics in the Historical Sample

Table E.1 shows summary statistics for the 120 occupations. In the top panel, we see that variation of employment growth in the cross-section of occupations is substantial, with 10th percentile occupations shrinking at 1.8 log points annually (averaged over the period 1985–2010) and 90th percentile occupations growing at 2.4 log points, respectively. When weighting by initial size, the negative average employment growth partly stems from the fact that formerly large manufacturing- and craft-related occupations have shrunk over time.[58] Second, annualised occupational price growth, as given by our preferred measure (wage growth of stayers in the occupation), is positive at 0.59 log points, again with considerable variation around this average (-0.96 and +2.17 log points for occupations at the 10th and 90th percentile, respectively). Only slightly less variation is found for our alternative measure of occupational prices à la Cortes (2016).

The middle panel of Table E.1 shows, among others, the distribution of occupational certification and regulation (coded between 0 and 1) and the shares of workers with university degrees. The bottom panel shows task intensities (analytical, routine, manual) across the 120 occupations. Consistent with earlier work (Gathmann & Schönberg, 2010), there exists substantial variation. For example, the median occupation is more than twice as routine-intensive as the occupation at the lowest decile. Task distance is normalised between zero and one, and best interpreted as a ranked ordinal variable (see its construction in the previous section). Still, the table reports e.g. distance at the 10th percentile (i.e. occupations using relatively similar task sets) and at the 90th percentile (occupations using rather different task sets).

Table E.2 displays summary statistics for annualised employment and occupational price changes separately by each five-year sub-period from 1985 to 2010. We see substantial variation over time: e.g. average wage and employment growth was substantially faster in the pre-unification years 1985–1990 and turned negative in the economically sluggish early 2000s.

Table E.3 presents summary statistics for the transition probability matrix, Π, and the elasticity matrix, D. Diagonal elements (i.e. probabilities for staying and own-price elasticities) are on average substantially larger than off-diagonal elements (for switching

[58]The results of our main analyses do not substantively differ whether we weight occupations by their initial size or not.

occupations and cross-price elasticities). However, dispersions of off-diagonal elements are higher relative to their means and skewness is clearly substantial in these variables. As discussed in the main text, cross-elasticities at the top of the distribution are as high as some of the own-elasticities, but thereafter fall off very rapidly in size. For example, the 99th percentile cross-elasticity (0.04 in Table E.3) is already somewhat lower than the minimum own-elasticity (0.07 in Table 1).

The persistence of elasticity components across time is also shown in Table E.3. In particular, the matrix of elasticities is constructed for different five-year periods (1975–1980,…, 2000–2005, 2005–2010), and then the relation of the respective own-elasticities and cross-elasticities (the matrix elements) is separately studied across those periods. Autocorrelations turn out high, in the range of 0.75–0.90 even for the long time distances between the early and late periods. This is consistent with the high autocorrelation of occupational task contents reported in Gathmann & Schönberg (2010) and with the findings when estimating our model pooled in these five-year sub-periods.

The second column of Table E.3 shows that occupational cross-price elasticities are strongly skewed and with high kurtosis. We decompose the log of the cross-price elasticities using the expression in Remark 2 as follows:

$$\ln\left(-d_{jk}\right) = \ln\left(\tau_k\right) + \ln\left(Cov_\tau(\tilde{\pi}_{\cdot,j}, \tilde{\pi}_{\cdot,k}) + 1\right)$$

That is, the variance of log differences in cross-price elasticities can be decomposed into variances of log differences in sector sizes and occupational substitutabilities (plus one, to make them all positive). Panel B of Table E.4 shows that, in fact, most of the dispersion of $\ln\left(-d_{jk}\right)$, and hence the skewness in levels of d_{jk}, is driven by the dispersion of $\ln\left(Cov_\tau(\tilde{\pi}_{\cdot,j}, \tilde{\pi}_{\cdot,k}) + 1\right)$, while the dispersion in log occupational sizes contributes less than 30%. Although not shown in Table E.4, but can be easily inferred, the covariance of log occupational size with the substitutability term is negligible. Such covariance terms are often important in models of matching between worker and employer types, generating skewed wage distributions (e.g. Sattinger, 1993). Here, this interaction does not matter and cross-elasticities largely inherit their distribution from the occupational substitutabilities.

Own-elasticities are distributed approximately normally in *levels*. A formal test fails to reject normality based on the skewness and kurtosis reported in Table E.3 above.[59] Although we do not explore the reason for this feature rigorously here, we conjecture it is because own-price elasticities comprise the sum of many apparently independently

[59] p-value on the skewness test is 0.40, with a p-value on the kurtosis test of 0.13.

distributed terms, as Remark 2 indicates. In line with this feature, and with the first expression of Remark 2, we decompose this elasticity as:

$$d_{jj} = \underbrace{\sum_{k \neq j} \tau_k Cov_\tau \left(\tilde{\pi}_{\cdot,j}, \tilde{\pi}_{\cdot,k} \right)}_{\substack{\text{aggregate} \\ \text{substitutability}}} + \underbrace{1}_{\text{direct}} - \underbrace{\tau_j}_{\substack{\text{price} \\ \text{index}}}$$

As Panel A of Table E.4 shows, and consistent with the discussion in the main text, the variation in aggregate substitutabilities is by far the dominant component of the variance of own-price elasticities. Compared to this, the variation in occupation sizes and its covariance with aggregate substitutability are minuscule.

We show how own-price elasticities d_{jj} relate to several occupational characteristics in panel (a) of Figure E.1. These include the share of workers with university degrees, workers' mean age, occupational certification and regulation as well as analytical, routine, and manual task intensities. Panel (b) of Figure E.1 plots cross-price elasticities against occupational task distance.

Finally, Table E.5 offers the full list of the 120 occupations ranked by their respective own-price elasticities, together with their employment sizes in 1985 and 2010. The table also reports own-price elasticities when the model includes transitions to and from unemployment, as discussed earlier in Appendix B.3.

Table E.1: Summary Statistics for the 120 Occupations.

	Mean	Weighted Mean	Std.Dev.	p10	p50	p90	Observ.
Annualised Employment and Occupational Price Changes (1985–2010)							
Log Employment	0.107	−0.123	1.921	−1.843	−0.065	2.369	120
Prices: Stayers' Wage Growth	0.586	0.516	1.354	−0.959	0.408	2.168	120
Prices: à la Cortes (2016)	1.102	1.065	0.953	−0.009	0.949	2.308	120
Other Occupational Characteristics							
Initial Employment Size in 1985 (%)	0.833	1.763	0.883	0.213	0.543	1.639	120
Employment Size in 2010 (%)	0.833	1.789	1.030	0.193	0.501	1.738	120
Occupational Certification	0.712	0.751	0.258	0.290	0.810	0.970	120
Occupational Regulation	0.103	0.079	0.228	0	0	0.380	120
Share of University Degree (%)	13.51	11.69	23.23	0.57	1.83	46.32	120
Mean Workers' Age	40.55	40.92	1.68	38.59	40.46	42.35	120
Task Intensity and Distance							
Analytical	0.069	0.064	0.075	0.010	0.039	0.181	120
Manual	0.095	0.089	0.071	0.016	0.075	0.186	120
Routine	0.151	0.153	0.079	0.062	0.131	0.271	120
Task Distance	0.499	0.497	0.296	0.061	0.541	0.870	14280
Proxy for demand shocks **r**	0.177	0.178	0.149	−0.037	0.217	0.326	120

Notes: The table presents summary statistics for annualised employment and occupational price changes during 1985–2010, occupational characteristics (e.g. the share of workers with university degrees by occupation), and task content information (i.e. analytical, manual, routine, and task distance). The last row presents the summary statistics for our proxy of demand shocks **r**, defined as Routine + Manual - Analytical (see Section 5.2, especially eq. (18)). The weighted mean is weighted by each occupation's employment share in 1985.

Table E.2: Summary Statistics. Annualised Employment and Occupational Price Changes by Sub-Periods

	Mean	Weighted Mean	Std.Dev.	p10	p50	p90	Autocorr. with 5-yr lag
Panel A. 1985–1990							
Δe (Log empl. change)	2.59	2.28	2.57	−0.15	2.32	5.72	-
Δp (Stayers' Wages)	2.10	2.08	1.44	0.40	1.85	4.07	-
Δp (à la Cortes, 2016)	2.38	2.38	1.19	0.99	2.23	4.08	-
Panel B. 1990–1995							
Δe	0.05	0.13	2.51	−3.13	−0.25	3.62	0.56
Δp (Stayers' Wages)	0.17	0.11	1.36	−1.33	−0.04	1.97	0.84
Δp (à la Cortes, 2016)	0.58	0.50	1.09	−0.71	0.33	2.11	0.75
Panel C. 1995–2000							
Δe	−0.19	−0.24	2.67	−2.87	−0.46	2.71	0.46
Δp (Stayers' Wages)	0.48	0.52	1.79	−1.57	0.25	2.56	0.83
Δp (à la Cortes, 2016)	0.75	0.82	1.51	−0.97	0.56	2.50	0.75
Panel D. 2000–2005							
Δe	−1.64	−1.43	2.27	−4.49	−1.46	1.35	0.71
Δp (Stayers' Wages)	−0.24	−0.17	1.32	−1.90	−0.24	1.51	0.84
Δp (à la Cortes, 2016)	0.09	0.12	1.07	−1.15	0.01	1.54	0.82
Panel E. 2005–2010							
Δe	−0.27	−0.04	2.18	−3.07	−0.31	2.07	0.59
Δp (Stayers' Wages)	0.42	0.61	1.38	−1.14	0.12	2.17	0.77
Δp (à la Cortes, 2016)	0.57	0.76	1.25	−0.88	0.22	2.25	0.82

Notes: The table presents summary statistics for annualised employment (Δe) and occupational price changes (Δp) for different 5-year periods (i.e. 1985-1990, 1990-1995, 1995-2000, 2000-2005, and 2005-2010). The last column refers to the autocorrelation between that period and the preceding 5-year period, e.g. the autocorrelation of employment changes between 1990-1995 relative to employment changes in 1985-1990.

Table E.3: Summary Statistics. Elasticity Matrix and Transition Probability Matrix

	Elasticity Matrix D		**Transition Probability Matrix Π**		**Matrix V**		**Matrix DV**	
	Own-Price Elasticity (d_{jj})	Cross-Price Elasticity ($-d_{jk} \times 100$)	Diagonal Elements (π_{jj})	Off-Diagonal Elements ($\pi_{jk} \times 100$)	Diagonal Elements (v_{jj})	Off-Diagonal Elements ($v_{jk} \times 100$)	Diagonal Elements (dv_{jj})	Off-Diagonal Elements ($dv_{jk} \times 100$)
Mean	0.434	0.364	0.746	0.214	0.508	-0.427	0.210	-0.177
Std. Dev.	0.128	0.939	0.090	0.660	0.079	0.916	0.035	0.319
Variance	0.016	0.882	0.008	0.436	0.006	0.839	0.001	0.102
Skewness	0.177	14.672	-0.722	17.449	1.203	-1.482	-1.090	-7.490
Kurtosis	3.634	493.494	4.393	585.670	5.881	54.605	5.565	113.507
p10	0.294	0.007	0.627	0.000	0.418	-1.096	0.174	-0.397
p50	0.430	0.111	0.754	0.046	0.501	-0.226	0.215	-0.082
p90	0.604	0.867	0.839	0.516	0.587	0.016	0.249	-0.017
p99	0.796	4.021	0.931	2.585	0.746	1.109	0.279	-0.001
Average autocorr. 5-year	0.875	0.875	0.867	0.804	0.895	0.931	0.893	0.918
Autocorrelation 30-year	0.749	0.736	0.762	0.615	0.745	0.823	0.744	0.806
Correlation with matrix D	-	-	-0.995	-0.944	-0.961	-0.224	0.968	0.946
Correlation with matrix V	-0.961	-0.224	0.934	0.246	-	-	-0.989	0.015
Number of Observations	120	14,280	120	14,280	120	14,280	120	14,280

Notes: The table presents summary statistics for the elasticity matrix D (Remark 1), the transition probability matrix Π, matrix V as in equation (49) and DV, where we use the equilibrium solution for $\frac{\theta}{\sigma}$. The average (5-year period) autocorrelation is computed by averaging autocorrelations of reported variables between 1980-1985 and 1975-1980, 1985–1990 and 1980–1985, 1990–1995 and 1985–1990, 1995–2000 and 1990–1995, and so on. The 30-year autocorrelation refers to the autocorrelation between the latest period 2005–2010 and the earliest period 1975–1980.

Table E.4: Summary Statistics: Elasticity Components

	Mean	Std. Dev.	Variance	Min	p10	p50	p90	Max	Skewness
Panel A. Elasticity components in Remark 2									
Aggregate Substitutability, $\sum_{k \neq j} \tau_k Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$	−0.558	0.126	0.0161	−0.918	−0.692	−0.565	−0.389	−0.202	0.131
Stationary Employment Size, τ	0.008	0.012	0.0001	0.001	0.002	0.004	0.017	0.090	4.360
Occupational Substitutability, $Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$	−0.429	2.418	5.848	−0.999	−0.979	−0.779	0.203	129.849	36.912
Panel B. Log components of Cross-Price Elasticity (d_{jk})									
(a) Stationary Employment Size		1.072	1.150						−0.488
(b) Occupational Substitutability + 1		1.582	2.504						−0.259
Cross-Price Elasticity: (a) + (b)		1.917	3.676						−0.479

Notes: Panel A of the table presents summary statistics for the elasticity components of the own-price and cross-price elasticity (as discussed in Remark 2). Panel B of the table decomposes the log of the cross-price elasticity into variances of log occupational size and log occupational substitutability (plus one, to make them all positive). The number of observations is 120 for own-price elasticity (corresponding to each occupation) and 14280 for cross-price elasticity and its components (corresponding to each occupation-pair).

Figure E.1: Own-Price and Cross-Price Elasticity: Comparison with External Metrics

(a) Own-Price Elasticity and Occupational Characteristics

(b) Cross-Price Elasticity and Task Distance

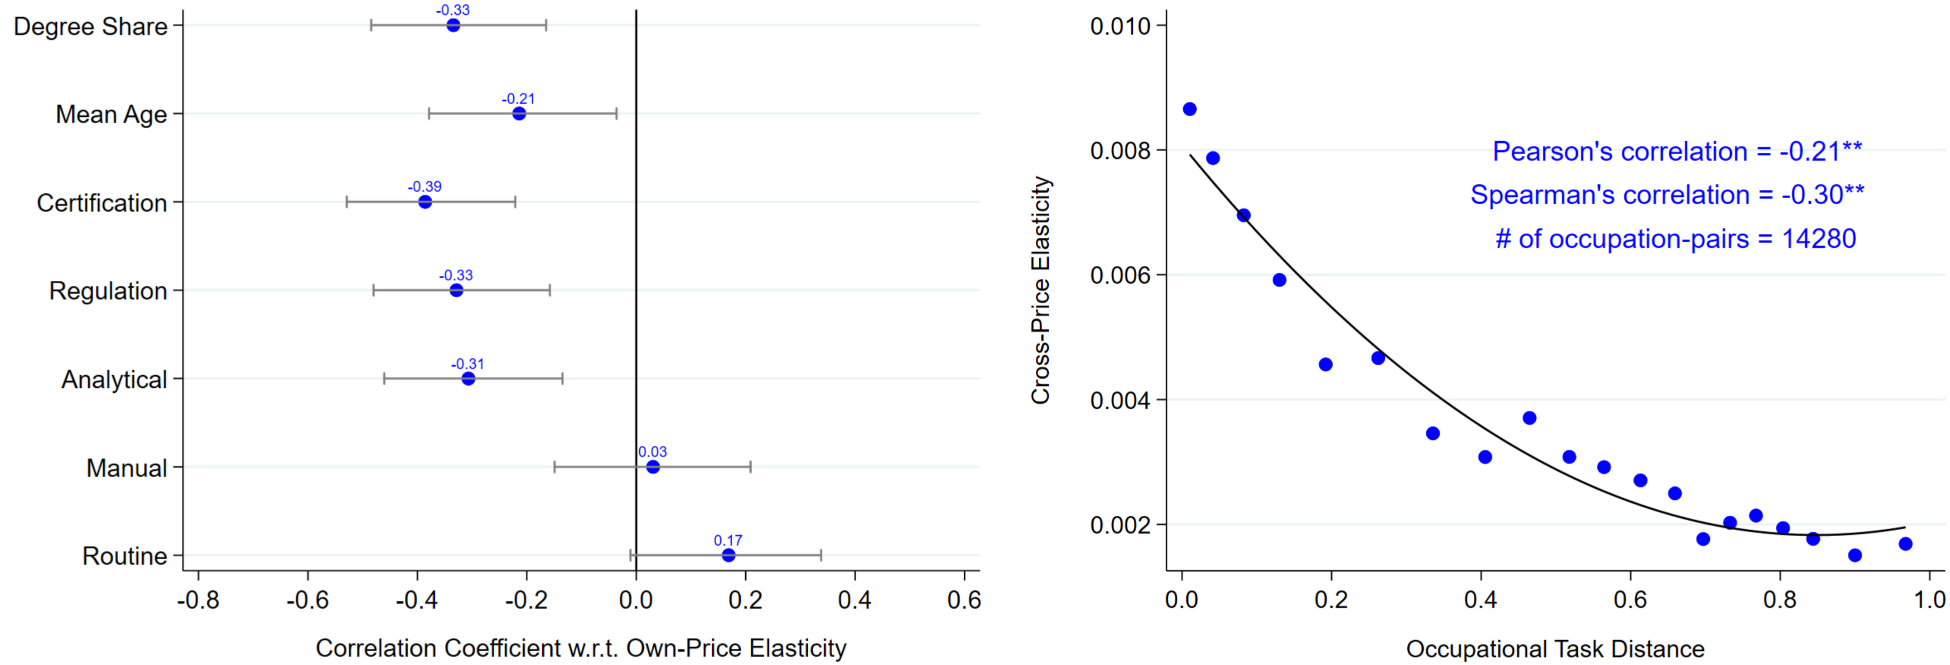

Notes: Panel (a) reports how own-price elasticity, namely d_{jj}, correlates with skill requirements across 120 occupations. Occupational certification and regulations come from Vicari (2014). Task content (analytical, manual, and routine) are measured using BiBB, see Appendix D.2. Correlations weighted by initial employment in each occupation. Panel (b) shows the relationship (with a quadratic fit) between cross-price elasticity, namely $-d_{jk}$, and occupational task distance measured as in Cortes & Gallipoli (2018).

Table E.5. All 120 Occupations Ranked by Diagonal Elements d_{jj}, and their Employment Size

	Own-Price Elasticity		% Share of Employment	
Occupations (based on German KldB 1988 Classification)	d_{jj}	d_{jj}^{NE}	1985	2010
Physicians up to Pharmacists	0.07	0.09	0.65	0.81
Bank specialists up to building society specialists	0.13	0.14	1.79	1.98
Nurses, midwives	0.16	0.19	0.37	0.67
Dental technicians up to doll makers, model makers, taxidermists	0.18	0.22	0.32	0.24
Non-medical practitioners up to masseurs, physiotherapists and related occupations	0.19	0.29	0.13	0.22
Journalists up to librarians, archivists, museum specialists	0.20	0.26	0.28	0.35
Hairdressers up to other body care occupations	0.23	0.40	0.06	0.06
Architects, civil engineers	0.23	0.26	0.83	0.69
Soldiers, border guards, police officers up to judicial enforcers	0.27	0.27	0.38	0.51
Musicians up to scenery/sign painters	0.28	0.42	0.29	0.31
Foremen, master mechanics	0.29	0.31	1.39	0.75
Health insurance specialists (not social security) up to life, property insurance specialists	0.29	0.33	0.85	0.89
Chemical laboratory assistants up to photo laboratory assistants	0.30	0.32	0.26	0.25
Doormen, caretakers up to domestic and non-domestic servants	0.30	0.34	0.97	0.97
Type setters, compositors up to printers (flat, gravure)	0.31	0.33	0.75	0.36
Gardeners, garden workers up to forest workers, forest cultivators	0.31	0.38	1.18	1.15
Social workers, care workers up to religious care helpers	0.31	0.34	0.42	0.68
Carpenters	0.32	0.36	1.57	1.17
Tile setters up to screed, terrazzo layers	0.33	0.43	0.42	0.30
Nursing assistants	0.33	0.42	0.20	0.33
Mechanical, motor engineers	0.33	0.35	1.07	1.21
Electrical fitters, mechanics	0.33	0.38	2.78	2.76
Chemists, chemical engineers up to physicists, physics engineers, mathematicians	0.33	0.36	0.35	0.34
Bricklayers up to concrete workers	0.34	0.43	2.95	1.20
Home wardens, social work teachers	0.34	0.45	0.28	0.46
Music teachers, n.e.c up to other teachers	0.34	0.38	0.27	0.32
Electrical engineers	0.34	0.35	1.00	1.18
Entrepreneurs, managing directors, divisional managers	0.34	0.37	2.63	2.11
Data processing specialists	0.35	0.39	1.18	3.46
Members of Parliament, Ministers, elected officials up to association leaders, officials	0.36	0.43	0.33	0.48
Measurement technicians up to remaining manufacturing technicians	0.36	0.39	0.81	0.48
Painters, lacquerers (construction)	0.36	0.41	1.11	0.91
Office specialists	0.36	0.38	6.10	8.15
Dietary assistants, pharmaceutical assistants up to medical laboratory assistants	0.36	0.56	0.03	0.05
Chemical plant operatives	0.36	0.39	1.25	0.97
Navigating ships officers up to air transport occupations	0.37	0.42	0.39	0.28
Paper, cellulose makers up to other paper products makers	0.37	0.42	0.53	0.50
Artistic and audio, video occupations up to performers, professional sportsmen, auxiliary artistic occup	0.37	0.47	0.27	0.25
Motor vehicle drivers	0.38	0.42	5.57	5.39
Toolmakers up to precious metal smiths	0.38	0.40	1.13	0.80
Cost accountants, valuers up to accountants	0.38	0.42	0.82	0.51
Railway engine drivers up to street attendants	0.39	0.41	0.77	0.61
Bakery goods makers up to confectioners (pastry)	0.39	0.42	0.41	0.41
Other technicians	0.39	0.41	1.96	2.43
Commercial agents, travellers up to mobile traders	0.39	0.43	1.58	1.10
Miners up to shaped brick/concrete block makers	0.40	0.46	1.33	0.47
Roofers	0.40	0.48	0.37	0.40
Survey engineers up to other engineers	0.40	0.44	0.75	1.82
Plumbers	0.40	0.42	1.35	1.23

Table E.5—continued

Occupations (based on German KlDB 1988 Classification)	Own-Price Elasticity		% Share of Employment	
	d_{jj}	d_{jj}^{NE}	1985	2010
Technical draughtspersons	0.40	0.43	0.60	0.48
Biological specialists up to physical and mathematical specialists	0.40	0.45	0.30	0.20
Mechanical engineering technicians	0.41	0.41	0.91	0.82
Butchers up to fish processing operatives	0.41	0.45	0.65	0.47
Turners	0.41	0.41	0.97	0.73
Generator machinists up to construction machine attendants	0.42	0.45	1.42	0.73
Goods examiners, sorters, n.e.c	0.42	0.48	0.90	0.58
Ceramics workers up to glass processors, glass fishers	0.42	0.46	0.40	0.22
Agricultural machinery repairers up to precision mechanics	0.42	0.48	0.53	0.54
Machine attendants, machinists' helpers up to machine setters (no further specification)	0.43	0.46	0.58	0.51
Stucco workers, plasterers, rough casters up to insulators, proofers	0.43	0.48	0.53	0.32
Metal grinders up to other metal-cutting occupations	0.43	0.47	0.50	0.35
Cooks up to ready-to-serve meals, fruit, vegetable preservers, preparers	0.43	0.53	0.62	1.05
Spinners, fibre preparers up to skin processing operatives	0.43	0.48	0.56	0.19
Motor vehicle repairers	0.43	0.45	1.63	1.65
Goods painters, lacquerers up to ceramics/glass painters	0.44	0.47	0.50	0.37
Chemical laboratory workers up to vulcanisers	0.44	0.49	0.41	0.30
Cutters up to textile finishers	0.44	0.47	0.24	0.08
Cashiers	0.44	0.40	0.10	0.07
Street cleaners, refuse disposers up to machinery, container cleaners and related occupations	0.44	0.48	0.63	0.72
Drillers up to borers	0.44	0.48	0.59	0.41
Iron, metal producers, melters up to semi-finished product fettlers and other mould casting occupations	0.45	0.52	0.96	0.60
Electrical engineering technicians up to building technicians	0.45	0.45	1.39	1.47
Wine coopers up to sugar, sweets, ice-cream makers	0.45	0.51	0.46	0.37
Room equippers up to other wood and sports equipment makers	0.45	0.47	0.39	0.27
Plant fitters, maintenance fitters up to steel structure fitters, metal shipbuilders	0.45	0.48	2.18	1.36
Carpenters up to scaffolders	0.46	0.58	0.63	0.49
Postmasters up to telephonists	0.46	0.54	0.30	0.36
Forwarding business dealers	0.46	0.51	0.42	0.47
Engine fitters	0.47	0.50	2.04	1.43
Farmers up to animal keepers and related occupations	0.47	0.48	0.49	0.42
Welders, oxy-acetylene cutters	0.47	0.50	0.72	0.51
Telecommunications mechanics, craftsmen up to radio, sound equipment mechanics	0.47	0.47	0.82	0.45
Steel smiths up to pipe, tubing fitters	0.47	0.54	0.58	0.34
Wood preparers up to basket and wicker products makers	0.48	0.52	0.48	0.26
Office auxiliary workers	0.49	0.52	0.34	0.31
Sheet metal workers	0.49	0.57	0.40	0.36
Wholesale and retail trade buyers, buyers	0.51	0.55	1.65	1.88
Factory guards, detectives up to watchmen, custodians	0.51	0.55	0.67	0.67
Special printers, screeners up to printer's assistants	0.51	0.54	0.35	0.21
Sheet metal pressers, drawers, stampers up to other metal moulders (non-cutting deformation)	0.51	0.53	0.53	0.32
Paviours up to road makers	0.52	0.58	0.49	0.32
Tourism specialists up to cash collectors, cashiers, ticket sellers, inspectors	0.53	0.58	0.49	0.65
Tracklayers up to other civil engineering workers	0.53	0.58	0.78	0.32
Metal polishers up to metal bonders and other metal connectors	0.53	0.57	0.44	0.28
Management consultants, organisers up to chartered accountants, tax advisers	0.53	0.51	0.41	1.29
Transportation equipment drivers	0.53	0.57	0.52	0.45
Warehouse managers, warehousemen	0.54	0.57	2.21	1.58
Housekeeping managers up to employees by household cheque procedure	0.54	0.69	0.05	0.08
University teachers, lecturers at higher technical schools up to technical, vocational, factory instructors	0.54	0.61	0.38	0.50

Table E.5—continued

Occupations (based on German KlDB 1988 Classification)	Own-Price Elasticity		% Share of Employment	
	d_{jj}	d_{jj}^{NE}	1985	2010
Economic and social scientists, statisticians up to scientists	0.56	0.66	0.35	0.57
Stowers, furniture packers up to stores/transport workers	0.56	0.61	1.95	2.91
Stenographers, shorthand-typists, typists up to data typists	0.56	0.59	0.11	0.12
Other mechanics up to watch-, clockmakers	0.56	0.61	0.45	0.79
Electrical appliance fitters	0.57	0.59	0.43	0.60
Plastics processors	0.57	0.62	0.67	0.86
Packagers, goods receivers, despatchers	0.57	0.60	0.86	0.92
Locksmiths, not specified up to sheet metal, plastics fitters	0.59	0.61	1.32	1.54
Salespersons	0.60	0.64	1.57	2.06
Laundry workers, pressers up to textile cleaners, dyers, and dry cleaners	0.60	0.65	0.06	0.06
Building labourer, general up to other building labourers, building assistants	0.61	0.71	1.26	0.97
Electrical appliance, electrical parts assemblers	0.62	0.64	0.22	0.20
Other assemblers	0.63	0.65	0.31	0.81
Household cleaners up to glass, building cleaners	0.63	0.75	0.26	0.41
Publishing house dealers, booksellers up to service-station attendants	0.63	0.67	0.17	0.13
Restaurant, inn, bar keepers, hotel proprietors, catering trade dealers up to waiters, stewards	0.64	0.65	0.35	0.58
Metalworkers (no further specification)	0.67	0.69	1.07	1.38
Assistants (no further specification)	0.71	0.74	0.75	3.00
Ancillary hospitality workers	0.74	0.81	0.21	0.12
Medical receptionists	0.80	0.46	0.01	0.02
Nursery teachers, child nurses	0.80	0.75	0.02	0.09

Notes: The table provides diagonal elements of the elasticity matrix D, not accounting (column (1), our baseline specification) and accounting for unemployment transitions (column (2), an extension of our model discussed in Appendix B.3). Columns (3)–(4) report the occupation's percentage share of employment in 1985 and 2010, respectively. The IAB translates *Uebrige Gaestebetreuer* as 'Other attending on guests'. We instead translate as 'ancillary hospitality workers'.

E.2 Occupational Changes over 1985–2010

Figure 2b in the main text splits occupations at the median of d_{jj} and draws two separate regression lines. Figure E.2 below alternatively splits occupations into d_{jj} quartiles. The resulting four regression lines are visibly ranked by predicted labour supply elasticity, with the lowest d_{jj} quartile (in blue colour) exhibiting the steepest relation of employment vs prices, the highest d_{jj} quartile (in red colour) exhibiting the flattest relationship, and the middle quartiles (in green and orange) ranked in between.

Figure E.2: Occupational Price and Employment Changes by Own-Price Elasticity d_{jj} Quartiles

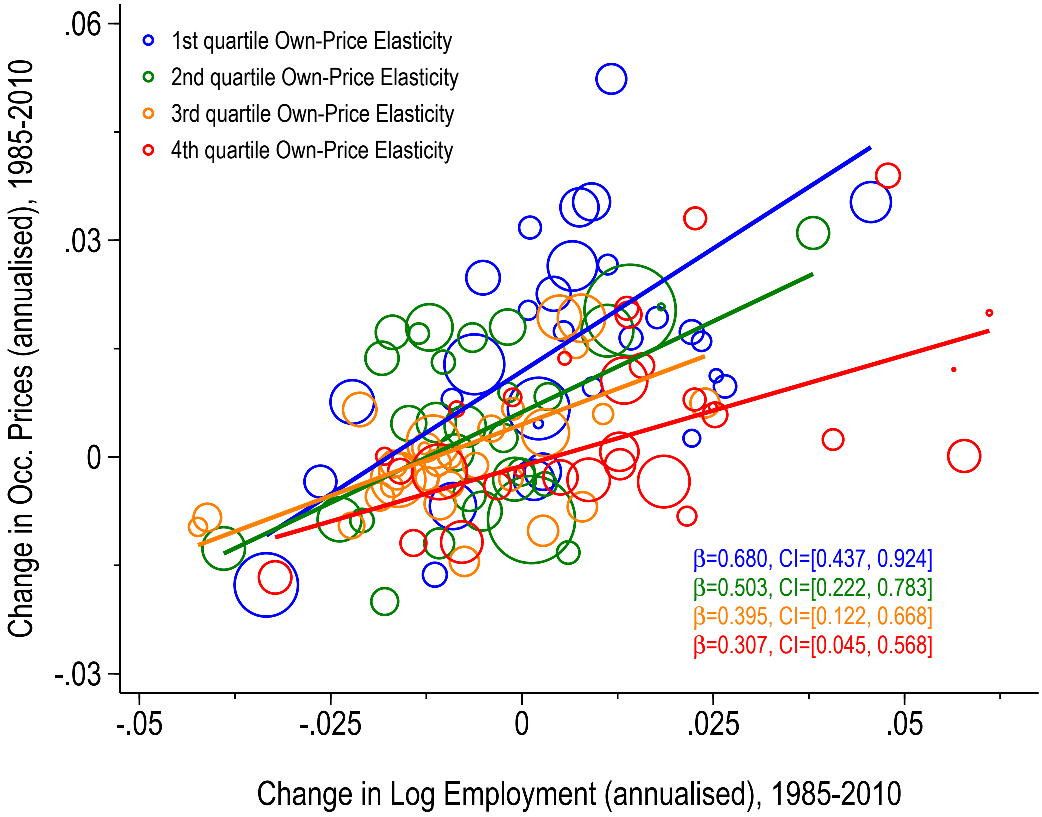

Notes: The figure shows the lines from an occupation-size weighted regression of price change on employment change, split by occupations in the lowest (blue), second (green), third (orange), and highest (red) quartile of own-price elasticity d_{jj}. β refers to the slope coefficient, and CI stands for the 95% confidence interval. Marker size indicates the baseline employment (in 1985) in each occupation.

Table E.6 considers the case in which own-occupation effects are not further split into a fixed relationship and the additional effect of the heterogeneity in elasticities d_{jj}. That is, it directly implements an unrestricted and a restricted version of eq. (9). Note that the coefficient in column (2) of the table is negative because of omitted variable bias (OVB). In the short regression only on cross-effects of column (2), $\Delta e_j = \theta_2 \sum_{k \neq j} d_{jk} \Delta p_k + \varepsilon_j$, this leads to an OVB for θ_2 of $\theta_1 \frac{Cov(d_{jj}\Delta p_j, \sum_{k\neq j} d_{jk}\Delta p_k)}{Var(\sum_{k\neq j} d_{jk}\Delta p_k)}$. Considering these covariances as taken over random draws of price changes, for given j, k and d_{jk}, the numerator in this expression can be rewritten as $d_{jj} \sum_{k\neq j} d_{jk} Cov(\Delta p_j, \Delta p_k)$. Since d_{jk} is large negative for highly substitutable occupations, and close to zero for occupations that are further apart, and because prices for substitutable occupations tended to move in the same direction, then $d_{jj} \sum_{k\neq j} d_{jk} Cov(\Delta p_j, \Delta p_k) \ll 0$, which signs the OVB.

Table E.6: Determinants of Employment Changes: Own- and Cross-Effects (OLS)

		Dependent Variable: Δe_j			
		(1)	(2)	(3)	(4)
own effect:	$d_{jj}\Delta p_j$	1.81		4.10	
		(0.32)		(0.88)	4.15
total cross effect:	$\sum_{k\neq j} d_{jk}\Delta p_k$		−2.14	4.03	(0.70)
			(0.59)	(1.29)	
R-squared		0.310	0.163	0.394	0.394
Number of occupations		120	120	120	120

Notes: The table presents the unweighted estimates from different versions of eq. (9). Regressor in column (4) is $\sum_k d_{jk}\Delta p_k$, i.e. corresponding to the full model and as in the main text. All regressions include a constant. Observations weighted by occupation j's initial employment size. Period 1985–2010. Standard errors in parentheses; all coefficients shown are significant at the 1% level.

E.3 Identification of the Labour Supply Parameter θ

E.3.1 Estimation Strategy from the Model Equilibrium

Equation (53) is our basic regression equation (eq. (15) in the main text), extending eq. (8) to include supply shocks. The logic of requiring the IV is that, given that $\Delta \mathbf{s}$ is not observed, then an OLS regression of Δe_j on $\mathbf{d_j}\Delta\mathbf{p}$ will not work, because $\mathbf{d_j}\Delta\mathbf{p}$ is correlated with these shocks.

Suppose we have a variable, which we denote r_j, that is correlated with demand shifters $\Delta b_j \equiv \ln \frac{\beta_j}{1-\beta_j}$ but not with supply shifters Δs_j. In matrix notation:

$$\Delta \mathbf{b} = \kappa 1_N + \lambda \mathbf{r} + \bar{\eta}$$

where κ and λ are scalars, 1_N is a vector of ones and $\bar{\eta}$ is a vector of shocks.

Then, from eq. (51):

$$\begin{aligned}
& \Delta \mathbf{p} \approx V \Delta \mathbf{b} + S \Delta \mathbf{s} \\
\implies & \Delta \mathbf{p} \approx \lambda V \mathbf{r} + \bar{\epsilon} + S \Delta \mathbf{s} \\
\implies & D \Delta \mathbf{p} \approx \lambda D V \mathbf{r} + D \bar{\epsilon} + D S \Delta \mathbf{s} \\
& = \lambda D \left(\frac{\theta}{\sigma} D + I \right)^{-1} (I - W) \, \mathbf{r} + D \bar{\epsilon} + D S \Delta \mathbf{s} \\
& = \lambda D \left(\frac{\theta}{\sigma} D + I \right)^{-1} \tilde{\mathbf{r}} + D \bar{\epsilon} + D S \Delta \mathbf{s}
\end{aligned}$$

where the second line follows from the first because, if v_{ij} is the i, jth element of V, then $\sum_j v_{ij} = 0$. Vector $\tilde{\mathbf{r}}$ is the employment-share-weighted-demeaned version of $\mathbf{r}$ and finally, $\bar{\epsilon} \equiv V \bar{\eta}$. This is relationship (17) in the main text.

In terms of regressing Δe_j on the vector of price changes, this implies that an appropriate instrument for $\mathbf{d}_j \Delta \mathbf{p}$ is $\mathbf{d}_j V \mathbf{r}$.

Implementing this model requires having some information on the demand elasticity σ. We choose a calibration based on estimates from the literature. Based on a range of $\sigma \in [1.81, 2.10]$ from Burstein et al. (2019) and initial information on the potential value of θ from Table 2, we calibrate $\frac{\theta}{\sigma} = 2.3$ as a benchmark throughout the paper. Table E.7 shows the robustness of our results to different values of $\frac{\theta}{\sigma}$.

Similar to Figure 3a in the main text, we display the relationship between r_j and Δe_j in Figure E.3 below. As expected, the regression line is slightly steeper among more elastic occupations, which react to a demand shock relatively more in terms of employment.

Table E.7: Robustness to Different Values of $\frac{\theta}{\sigma}$

	$\frac{\theta}{\sigma} = 0.001$	$\frac{\theta}{\sigma} = 0.1$	$\frac{\theta}{\sigma} = 1$	$\frac{\theta}{\sigma} = 1.5$	$\frac{\theta}{\sigma} = 2$	$\frac{\theta}{\sigma} = 2.3$	$\frac{\theta}{\sigma} = 2.5$	$\frac{\theta}{\sigma} = 3$	$\frac{\theta}{\sigma} = 4$
IV estimate for θ	5.20	5.19	4.95	4.87	4.81	4.78	4.76	4.72	4.66
Implied σ	5200	519	4.95	3.25	2.41	2.08	1.90	1.57	1.17

Notes: The table shows the robustness of our IV estimate to different values of $\frac{\theta}{\sigma}$. The second row reports the implied σ. The case highlighted in blue ($\frac{\theta}{\sigma} = 2.3$) is the benchmark used throughout the paper.

Figure E.3: IV Reduced-Form for Employment

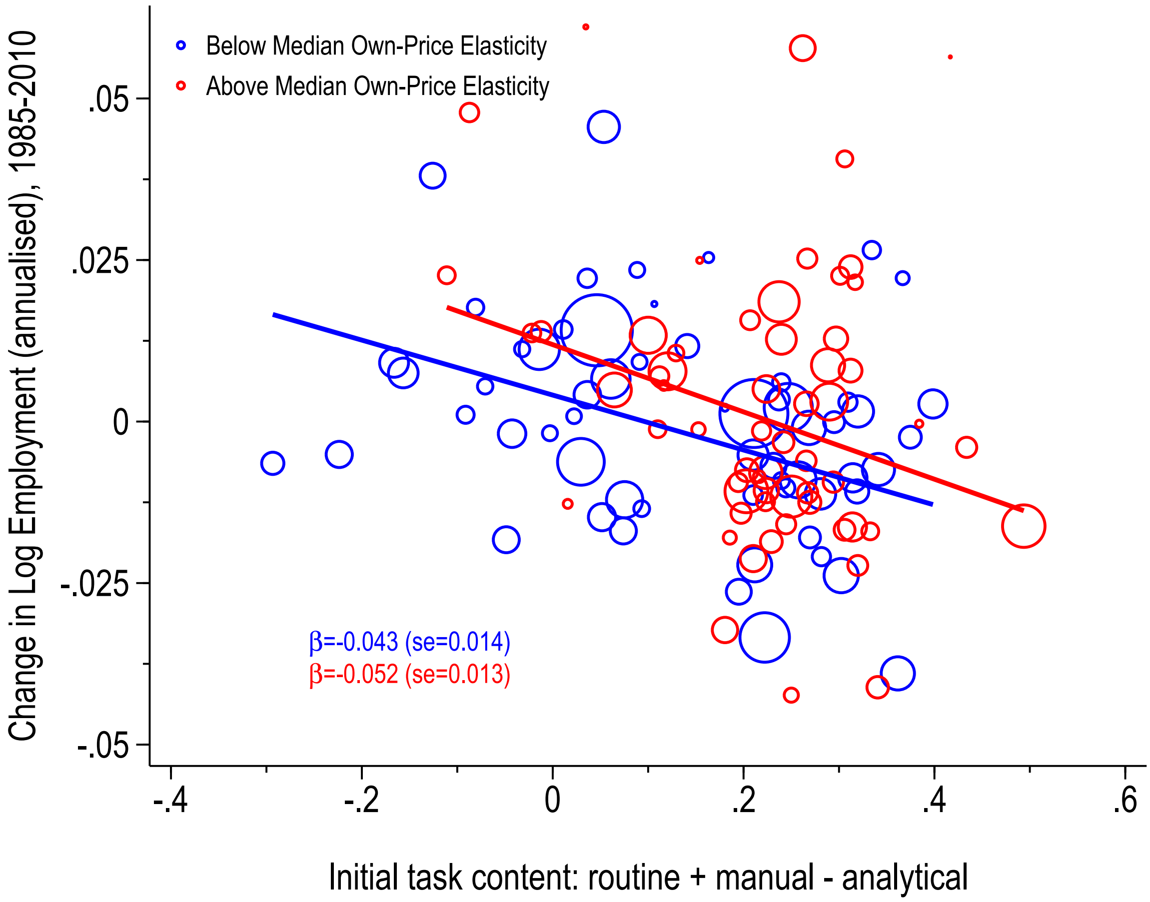

Notes: The figure shows reduced-form regressions of occupations' employment changes on their initial task contents r_j. Colour codes and linear regression lines are split by occupations below (blue, inelastic) and above (red, elastic) the median own-price elasticity (d_{jj}). β and *se* refer to the slope coefficient and standard error, respectively. Marker size indicates the baseline employment (in 1985) in each occupation.

E.3.2 OLS versus IV Estimates

We wish to estimate eq. (15), which is reproduced here for convenience:

$$\Delta \mathbf{e} \approx \theta D \Delta \mathbf{p} + \Delta \mathbf{s}$$

Allowing for a regression constant, we stack parameters into vector $\beta = [\alpha \quad \theta]'$ and regressors into $N \times 2$ matrix $X = [1_N \quad D\Delta\mathbf{p}]$, where 1_N is a vector of ones. The OLS estimate of β is then

$$\hat{\beta}_{OLS} = (X'X)^{-1}X'\Delta\mathbf{e} = \beta + (X'X)^{-1}X'\Delta\mathbf{s}.$$

From (13) and (14), we note that

$$D\Delta\mathbf{p} = DV(\Delta\mathbf{b} - \frac{1}{\sigma}\Delta\mathbf{s})$$

and in the data the relevant covariances and variances are quite similar with $Cov(\Delta b_j, \Delta s_j) = 0.000101$ and $\frac{1}{\sigma}Var(\Delta s_j) = 0.000128$, respectively. It turns out that also the weighting matrix DV does not change this near-equivalence such that $(D\Delta\mathbf{p})'\Delta\mathbf{s} = \Delta\mathbf{b}'V'D'\mathbf{s} - \frac{1}{\sigma}\Delta\mathbf{s}'V'D'\mathbf{s}$ is only slightly negative (close to zero). Since $\Delta\mathbf{s}$ is size-weighted mean zero, also $1'_N\Delta\mathbf{s} \approx 0$ such that

$$(X'X)^{-1}X'\Delta\mathbf{s} \approx [0 \;\; 0]'$$

That is, there happens to be little bias in the OLS estimate.

Therefore, we get from this that

$$\hat{\theta}_{OLS} \approx \theta$$

where, by construction, true θ is identified in (19) from the instrumental variables strategy under the relevant IV assumptions. Put differently, $\hat{\theta}_{OLS} - \theta = 4.15 - 4.78 = -0.63$, which is negative but small relative to the absolute value of θ.

E.3.3 Alternative Estimation from Changing Choice Probabilities

As a second approach, we implement an identification strategy based on changing choice probabilities over time. Building on Cortes & Gallipoli (2018), we rewrite our main eq. (2) to obtain relative switching probabilities as follows:

$$\log\left(\frac{\pi_{ij,t}}{\pi_{ii,t}}\right) = \theta p_{j,t} + a_{ij,t} - (\theta p_{i,t} + a_{ii,t})$$

where $\frac{\pi_{ij,t}}{\pi_{ii,t}}$ are the relative switching probabilities, $\theta p_{j,t}$ is the general pecuniary payoff to occupation j at time t, and $a_{ij,t}$ captures potential pecuniary and non-pecuniary costs of selecting occupation j at time t for individuals initially working in occupation i.

Assume that $a_{ij,t} = a_{ij,t-1} + u_{ij,t}$ evolve randomly over time with $u_{ij,t}$ *i.i.d.* normally distributed, then we can rewrite the previous equation as

$$\triangle\log\left(\frac{\pi_{ij,t}}{\pi_{ii,t}}\right) = \theta\left(\triangle p_{j,t} - \triangle p_{i,t}\right) + \varepsilon_{ij,t} \tag{59}$$

where the regression error $\varepsilon_{ij,t} \equiv u_{ij,t} - u_{ii,t}$ is also *i.i.d.* normally distributed. There are two interpretations of the error term: If this reflects measurement errors in the choice probabilities (e.g. due to sampling), OLS regression is consistent. If the ε_{ij} partly also reflect changes in the underlying relative switching costs, and thereby potential endogeneity, we will need an instrument for the price changes. Subtracting the main diagonal elements (i.e. cases where $i = j$), we end up with $N \times (N-1)$ observations. In our application, this means 120x119=14280. Following our main analysis, the change over time here refers to the 25 years between 1985 and 2010.[60] To calculate π_{ij}, we use 5 yearly flows during a period of 10 years. That is, the initial π_{ij} is based on 5-yearly flows during 1975–1985, whereas the end period π_{ij} is based on 5-yearly flows during 2000–2010.

Columns (1)–(2) of Table E.8 show the results (OLS and IV, respectively) from the estimation of eq. (59), where we replace the zeros with the smallest value observed for π_{ij}.[61] We instrument for relative price changes using initial task content, i.e. $r_{j,t_0} - r_{i,t_0}$ where $r_{j,t_0} = (routine_{j,t_0} + manual_{j,t_0}) - analytical_{j,t_0}$ captures the net routineness of occupation j at the baseline period t_0 (measured in the late 1970s, early 1980s).[62] Alternatively, column (3) shows the IV estimate for a specification where observations with zero occupational flows are dropped from the sample. In all cases, standard errors are clustered at the occupation level (i.e. 120 clusters). The table shows that occupational choice probabilities react to changing prices. Also, the estimates for θ, which represents the wage preference, are overall similar to the ones reported in the main analysis.

Table E.8: Results for Alternative Way of Estimating θ

	Zeros replaced		**No Zeros**
	OLS	IV	IV
Wage preference θ	7.12	4.03	3.11
	(1.84)	(2.54)	(2.23)
First-stage coeff (se)		-0.064 (0.003)	-0.056 (0.003)
Observations	14280	14280	7530

Note: The table presents the results from the estimation of equation (59) over the period 1985-2010. Observations are at the occupation pair level. Standard errors are clustered at the occupation level in all specifications.

[60] The results are similar (statistically more precise) if we consider five 5-year sub-periods instead (e.g. 1985–1990, 1990–1995, and so on).

[61] Here we follow the literature on estimating gravity equations (see, among others, Head & Mayer, 2013; Cortes & Gallipoli, 2018).

[62] Instrumenting relative price changes using the model-based IV (i.e. $V\mathbf{r}$) leads to downward biases in the estimates (the division bias) as π_{ij} appears on both sides of the equation (Borjas, 1980).

F Model-Based Decomposition and Counterfactuals

This section develops the counterfactual elasticity matrices introduced in Section 6.1, relating them to the theory and empirics used in prior literature. We then report additional empirical results on the model solution and counterfactual analyses in Section 6.2.

F.1 Backing Out the Shocks

In Section 6, we use the model solution to construct counterfactuals. Here we show how to obtain the supply and demand shocks for this.

From (53), we immediately see that

$$\Delta \mathbf{s} \approx \Delta \mathbf{e} - \theta D \Delta \mathbf{p} \tag{60}$$

Similarly, from (51)

$$\begin{aligned} \Delta \mathbf{p} &\approx V \Delta \mathbf{b} + S \Delta \mathbf{s} \\ \implies \sigma (I - W) \Delta \mathbf{p} &\approx \sigma (I - W) V \Delta \mathbf{b} + \sigma (I - W) S \Delta \mathbf{s} \end{aligned}$$

Summing with (52) this implies that

$$\begin{aligned} \Delta \mathbf{e} + \sigma (I - W) \Delta \mathbf{p} &\approx (\sigma (I - W) V + \theta D V) \Delta \mathbf{b} \\ &= (\sigma (I - W) + \theta D) V \Delta \mathbf{b} \\ &= \sigma (I - W) \Delta \mathbf{b} \end{aligned}$$

where the last line follows from equation (48). Rearranging gives:

$$(I - W) \Delta \mathbf{b} \approx \frac{1}{\sigma} \Delta \mathbf{e} + (I - W) \Delta \mathbf{p}$$

Given the definition of the $b_j = \ln \frac{\beta_j}{1-\beta_j}$ as logs of relative demands, their (marginal) changes have mean of zero when weighted by employment shares. So we can write

$$\Delta \mathbf{b} \approx \frac{1}{\sigma} \Delta \mathbf{e} + (I - W) \Delta \mathbf{p} \tag{61}$$

without loss of generality. Equations (60)–(61) can be used to construct the shock vectors. Note that in (61) the term $(I - W)$ is retained to de-mean any given vector of price changes. This term is not required in (60) because the D matrix de-means the vector automatically. Additionally, $\Delta \mathbf{e}$ is (weighted) mean zero by construction.

F.2 Counterfactual Elasticities

F.2.1 Heterogeneous Own-Price Elasticities Only

The counterfactual matrix D_{own} considers the case that occupations' aggregate (own-price) elasticities vary but their substitutabilities with other occupations are homogeneous. In particular, we have that $Cov_{\tau}\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) = c \in [-1, 0]$ in eq. (6) and $Var_{\tau}\left(\tilde{\pi}_{.,j}\right) = -\frac{1-\tau_j}{\tau_j} c$ in eq. (7). The main diagonal elements of D_{own} are the actual own-price elasticities, whereas cross-price elasticities reduce to size-weighted fractions of the on-diagonals $\frac{-\tau_k}{1-\tau_j} d_{jj}$.[63]

A specific version of this counterfactual with $c = 0$ can be derived from setups commonly used in the literature on firms, even if their focus is on studying heterogeneity of (own-price) labour supply elasticities facing employers. Consistent with, among many others, Card et al. (2018); Lamadon et al. (2022); Berger et al. (2022), one could take a simpler version of individuals' indirect utility eq. (1) as follows:[64]

$$u_j(\omega) = \theta p_j + a_j + \varepsilon_j(\omega), \tag{62}$$

Note that, in this case, switching costs a_j do not depend on the source employer i.

We derive the versions of Remarks 1–2, which result from eq. (62), by noting that the choice probability $\pi_j = \frac{\exp(\theta p_j + a_j)}{\sum_{k=1}^{N} \exp(\theta p_k + a_k)}$ also no longer depends on sending occupation i. For occupation sizes, we obtain:

$$\begin{aligned} E_j(\mathbf{p}) &= \textstyle\sum_i \tau_i \pi_j &= \pi_j \\ & &= \tau_j \quad \text{if } \mathbf{p} = \mathbf{p}^* \end{aligned}$$

since $\sum_i \tau_i = 1$ in the first line and then $\pi_j = \tau_j$ in baseline stationary equilibrium.

From this, we obtain $\tilde{\pi}_{i,j} = \frac{\pi_j}{\tau_j} = 1$ for all i, j and $Cov_{\tau}\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) = Var_{\tau}\left(\tilde{\pi}_{.,j}\right) = 0$. Without combination-specific access costs, occupations are just all equally substitutable from a labour supply perspective. Remark 2 then lead to

$$d_{jk} = \begin{cases} 1 - \tau_j & \text{if } j = k \\ -\tau_k & \text{otherwise} \end{cases}$$

[63]Forcing fully homogeneous cross-elasticities (i.e. $\frac{-d_{jj}}{N-1}$) yields very similar empirical results to those shown below. In both cases, D_{own} is still a valid elasticity matrix, since $d_{jj} = -\sum_{k \neq j} d_{kj}$.

[64]In Berger et al. (2022) or Lamadon et al. (2022), the substitutability between employers within a market is fixed by what corresponds to our parameter θ. Across predefined markets (region-industries) is an extra substitutability parameter, which leads to a nested CES or logit structure. In contrast, we allow for flexibly heterogeneous occupational substitutabilities as governed by job flows in the data.

which is the result corresponding to Remark 1. The economic model with $a_{ij} = a_j$ thus generates a version of our matrix with homogeneous occupational substitutabilities D_{own}, and with $c = 0$ as mentioned above.

F.2.2 Fully Homogeneous Labour Supplies

The second counterfactual imposes completely homogeneous labour supply elasticities. The main diagonal elements of matrix D_{hom} become $\bar{d}_{diag} = \sum_j \tau_j d_{jj}$ and cross-price elasticities a constant fraction of it $\frac{-\bar{d}_{diag}}{N-1}$.[65] This counterfactual is consistent with specifications in the empirical literature that regress occupations' log employment changes on their log wage changes (e.g. Autor et al., 2008; Dustmann et al., 2009; Cavaglia & Etheridge, 2020; Böhm et al., 2024, or column (1) of Table 2). This is formalised in terms of counterfactuals as follows:

$$\begin{aligned} \Delta e_j &= \theta \sum_{k=1}^{N} d_{jk} \Delta p_k \\ \Rightarrow \Delta e_{j,cf} &= \tilde{\theta} \Delta p_j - \tilde{\theta} \left(\frac{1}{N} \sum_{k=1}^{N} \Delta p_k \right) \end{aligned}$$

where counterfactual employment changes in the second line are obtained by replacing d_{jk} by $\frac{-\bar{d}_{diag}}{N-1}$. The first $\tilde{\theta} \equiv \frac{N}{N-1} \bar{d}_{diag} \theta$ is a single slope parameter on the price change and the second term becomes a regression constant that reflects average wage growth in the economy. In the equilibrium model (15), there is additionally an error term Δs_j, which reflects supply shocks. Alternatively, as in the main text, we can normalise $\Delta \mathbf{p}$ to have a mean of zero without loss of generality, in which case $\Delta e_{j,cf} = \tilde{\theta} \Delta p_j$.

The economic model would generate a specific version of D_{hom} with $\bar{d}_{diag} = \frac{N-1}{N}$ and $\tilde{\theta} = \theta$ if, in addition to substitutabilities, all occupation sizes are also the same. That is, when $\theta p_j + a_j = const.$ in eq. (62).

F.3 Further Results on Decomposition and Counterfactuals

This section complements the decomposition and counterfactual analyses in Section 6.

Figure F.1 shows the impact of including labour supply heterogeneity in a counterfactual with no supply shocks ($\Delta \mathbf{s}_{off} = \mathbf{0}$). Figure F.1a, same as Figure 4a, starts by considering the case with only demand shocks in the fully homogeneous model (i.e. D_{hom}). In this case, all occupational changes induced by demand shocks $\Delta \mathbf{b}$ run perfectly along a single supply curve. Figure F.1b then introduces both own- and cross-occupation effects keeping $\Delta \mathbf{s}_{off} = 0$. Relative to F.1a, variation around the regression line increases, such that the R-squared reduces to 69%. The locus of points moves on average *counterclockwise* and the slope of the regression line increases from 0.52 to 0.86. These changes show the

[65] Empirical results below do not change if we size-weight the cross-price elasticities as $\frac{-\tau_k}{(1-\tau_j)} \bar{d}_{diag}$.

importance of allowing for supply heterogeneity (and especially cross-occupation effects, which effectively reduce elasticities) to explain the data.

Figure F.2 plots the distribution of demand and supply shocks by occupation, exhibiting a generally positive correlation between the two (0.22). It shows that, e.g. occupations such as Assistants or Data processors experienced positive demand and supply shocks, while occupations like Bricklayers suffered negative demand and supply shocks. Another interesting example is the occupation Physicians and pharmacists, which experienced a (large) positive demand shock but no supply shock.

Figure F.3 and Figure F.4 display employment and wage changes along the occupational wage distribution (in the initial year 1985), for the full model and the fully homogeneous (counterfactual) model, respectively. We highlight some key points:

First, our period of analysis is characterised by an increase in wage inequality and employment polarisation. This is represented in Figure F.3 by the dashed black line, which reproduces estimates from the raw data. This evidence is consistent with Dustmann et al. (2009), among others. Similar to them, we find that for occupations in the upper half of the wage distribution, employment and wage changes are positively correlated, while they are negatively correlated for occupations in the lower half.

Second, a key strength of our framework is that it allows us to decompose the contribution of demand and supply shocks to the observed wage and employment changes. This decomposition, which follows from equations (13) and (14) in Section 5, reveals the distinct roles played by demand and supply shocks. Demand shocks, depicted in grey, emerge as the primary drivers behind both wage and employment changes. They are, however, more important in explaining wage changes than in explaining employment changes. For the latter, as we extensively discuss in Section 6.2 and Table 3, supply shocks and supply heterogeneity also play a role.

Finally, and related to the last point, switching off supply heterogeneity and considering counterfactual outcomes from the fully homogeneous model (i.e. comparing Figure F.4 to F.3) result in smaller wage changes and larger employment changes across occupations. The intuition for this is, as we discuss in the main text, that heterogeneous cross effects make occupations less price elastic. As such, realised labour supply elasticities captured by the full model are lower than those captured in the homogeneous model.

Figure F.1: Counterfactual Changes of Prices and Employment (II)

(a) Demand shocks only; Fully-homog. supply

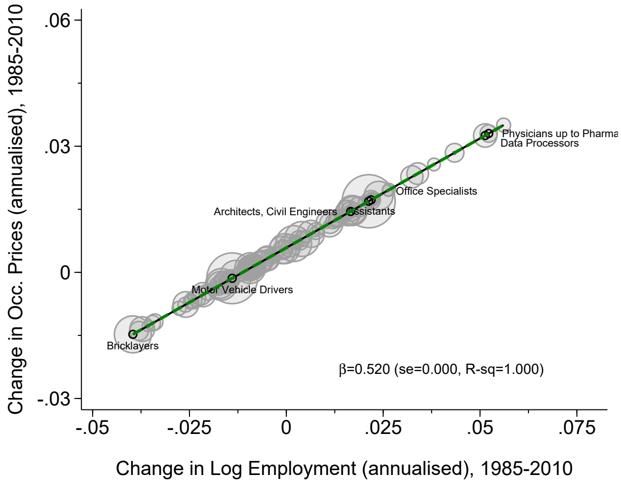

(b) ... Fully heterogeneous supply

Notes: The figure shows occupational price and employment changes for different manipulations of $\Delta\mathbf{s}$ and the elasticity matrix D. In F.1a, both supply shocks and heterogeneity in D are switched off (i.e. $\Delta s = \Delta s_{off} = 0$ and D_{hom}), leaving only demand shocks. F.1b introduces heterogeneous own- and cross-price elasticities (i.e. full matrix D is used). For the exact description of the counterfactuals, see Section 6. The OLS with slope coefficients, standard errors, and R-squared is shown for each panel. The regression line in F.1a is repeated as green-dashed in both panels. Marker size indicates the baseline employment (in 1985) in each occupation.

Figure F.2: Distribution of Demand and Supply Shocks by Occupation

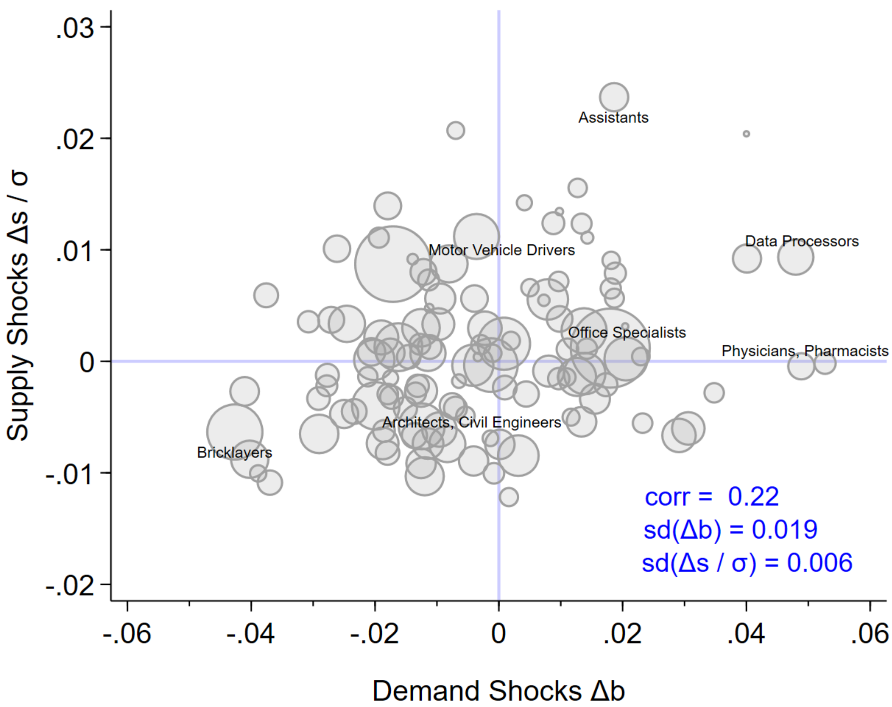

Notes: The figure shows the distribution of demand and supply shocks by occupation. Marker size indicates the baseline employment (in 1985) in each occupation. The standard deviations and correlation of demand and supply shocks are reported in the figure.

Figure F.3: Contribution of Demand and Supply Shocks (Full Model)

(a) ... to Wage Changes **(b)** ... to Employment Changes

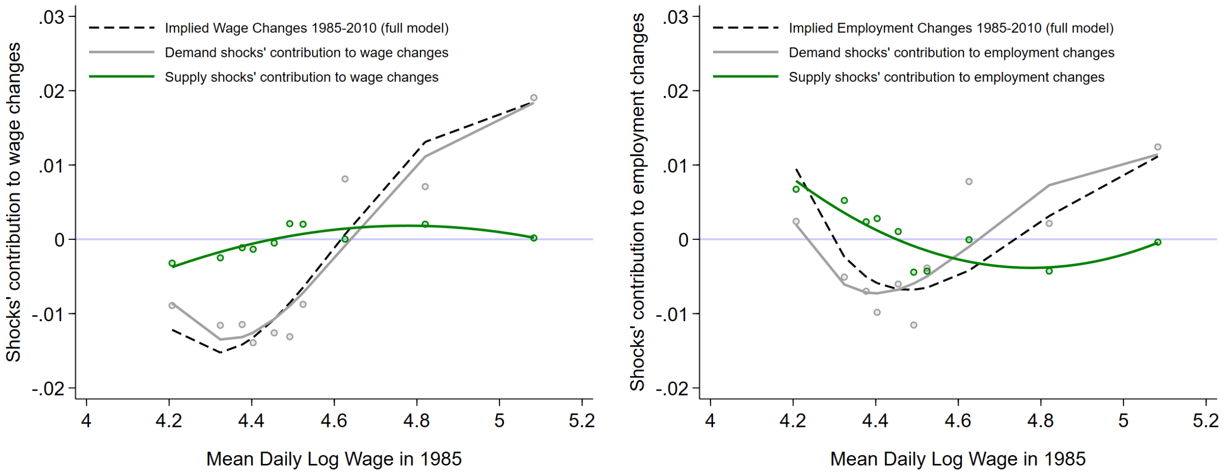

Notes: The left panel F.3a shows the contributions to price changes of demand and supply shocks across the occupational wage distribution for the full model. Each point represents the relevant decile of occupations ranked by mean wages in 1985. Changes are given by $V\Delta\mathbf{b}$ and $-\frac{1}{\sigma}V\Delta\mathbf{s}$, as in eq. (13). The right panel F.3b shows the contributions to employment changes of demand and supply shocks across the wage distribution for the full model. These are given by $\theta DV\Delta\mathbf{b}$ and $V\Delta\mathbf{s}$, as in eq. (14). For supply, a quadratic is used for the smoothed fit. For demand, a fractional cubic is used.

Figure F.4: Contribution of Demand and Supply Shocks (Fully Homogeneous Model)

(a) ... to Wage Changes **(b)** ... to Employment Changes

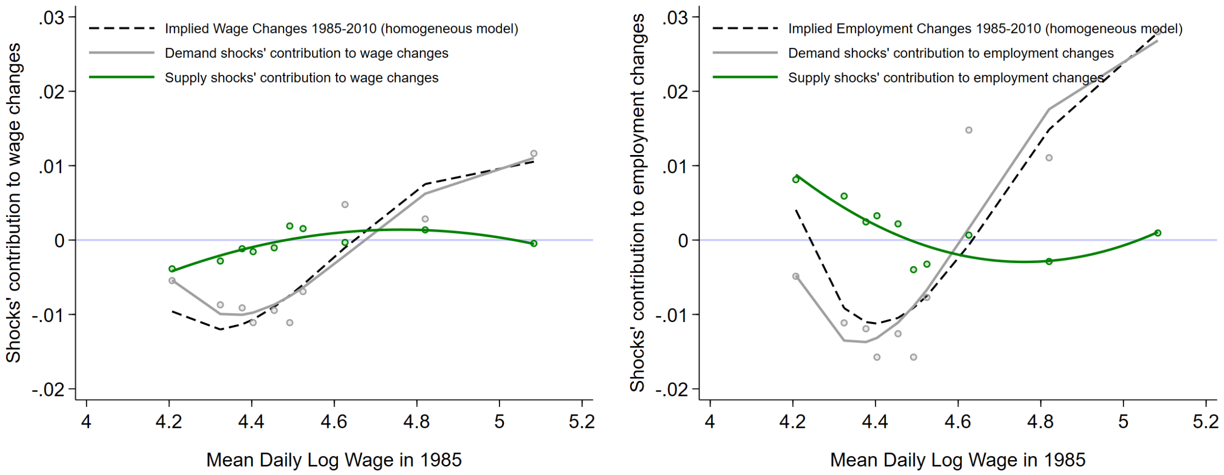

Notes: The left panel F.3a shows the contributions to price changes of demand and supply shocks across the wage distribution for the fully homogeneous model. Each point represents the relevant decile of occupations ranked by mean wages in 1985. Changes are given by $V_{hom}\Delta\mathbf{b}$ and $-\frac{1}{\sigma}V_{hom}\Delta\mathbf{s}$, parallel to eq. (13) and where $V_{hom} = \left(\frac{\theta}{\sigma}D_{hom} + I\right)^{-1}(I - W)$. The right panel F.3b shows the contributions to employment changes of demand and supply shocks. These are given by $\theta D_{hom}V_{hom}\Delta\mathbf{b}$ and $V_{hom}\Delta\mathbf{s}$, parallel to eq. (14). For supply, a quadratic is used for the smoothed fit. For demand, a fractional cubic is used.

G Robustness: Supplementary Material

In this section, we present details on the robustness checks of the main findings. We first study changes in occupational prices and employment by five-year sub-period. Second, we introduce an alternative method for estimating changes in occupational prices.

G.1 Analysis in Five-Year Sub-Periods

In the main analysis, we study changes in occupational prices and employment over the period 1985–2010. In this section, we split this longer interval into five-year sub-periods (1985–1990, 1990–1995, 1995–2000, 2000–2005, and 2005–2010), to explore robustness and potential temporal heterogeneity.

The pooled panel sample containing 600 observations (120 occupations × 5 sub-periods) is used to estimate an extended version of eq. (16):

$$\Delta e_{jt} = \alpha + \theta d_{jj} \Delta p_{jt} + \theta \sum_{k \neq j} d_{jk} \Delta p_{kt} + \delta_t \ (+\gamma_j) + \varepsilon_{jt} \tag{63}$$

where t refers to a five-year period, and the matrix of elasticities D can be obtained using the baseline period 1975–1984 as previously or using the lagged matrix from the preceding five-year period (e.g. for the period 1995–2000, the matrix of elasticity is computed using employment transitions over the period 1990–1995).[66] The period fixed effects (δ_t) capture unobserved time-specific shocks or trends that affect all occupations uniformly within each sub-period. A more demanding specification additionally includes occupation fixed effects (γ_j), removing average occupational growth over 1985–2010 and identifying only from accelerations/decelerations in the respective sub-period.

Figure G.1 plots prices against employment growth for the pooled sample of 600 occupation–sub-periods (G.1a) as well as separately for each sub-period (G.1b), analogous to the main text Figure 2b. The previous finding is strengthened in the sense that each regression slope for above-median own-price elastic occupations (in blue) is flatter than any slope for below-median own-price inelastic occupations (in red). Similarly, Table G.1 shows that linear OLS and IV estimation on the pooled data essentially reproduce the results obtained in the main text. Even in estimations with occupation fixed effects (γ_j), which only use deviations of price changes from their 1985–2010 averages interacted with the price elasticities, results are broadly similar to before.[67] In sum, estimation in a series of shorter intervals shows that the role of occupational price elasticities persists,

[66] Consistent with the high autocorrelation of matrix D over time discussed in Table E.3, results are similar whether we use the baseline or the lagged matrix.

[67] Note that we can only do the OLS for this as our instrument does not vary by period.

with some evidence that even acceleration/deceleration of price growth in different subperiods is translated into employment growth according to these elasticities.

G.2 Alternative Occupational Price Estimation

The main results in Section 5 use wage changes of occupation stayers' (i.e. workers who do not switch occupations from one year to the next) as the main estimate of changes in occupational prices. This accounts flexibly for the selection into occupations based on observable and unobservable individual characteristics. In this section, we use an alternative price estimation that also controls for the occupation-specific effect of time-varying observable characteristics on wages.

In this approach, originally proposed by Cortes (2016), observed log wages for individual ω in period t are modeled by

$$\ln w_t(\omega) = \sum_j Z_{jt}(\omega)\varphi_{jt} + \sum_j Z_{jt}(\omega)X_t(\omega)\zeta_j + \sum_j Z_{jt}(\omega)\kappa_j(\omega) + \mu_t(\omega) \tag{64}$$

where $Z_{jt}(\omega)$ is an occupation selection indicator that equals one if individual ω chooses occupation j at time t, φ_{jt} are occupation-time fixed effects, and $\kappa_j(\omega)$ are occupation-spell fixed effects for each individual. The model allows for time-varying observable skills (e.g. due to general human capital evolving over the life cycle) by including in the control variables X_t a set of dummies for five-year age bins interacted with occupation dummies.[68] Finally, $\mu_t(\omega)$ reflects classical measurement error, which is orthogonal to $Z_{jt}(\omega)$. It may be interpreted as a temporary idiosyncratic shock that affects the wages of individual ω in period t regardless of their occupational choice. The estimated occupation-year fixed effects (φ_{jt}) are the parameters of interest, which allow studying changes over time in occupation's log prices ($\Delta p_j = \varphi_{j,2010} - \varphi_{j,1985}$).

The results using occupational prices à la Cortes (2016) turn out similar to our main results. The main figures of the paper using this alternative measure for changes in occupational prices are replicated in Figure G.2. The main regression results, shown in Table G.2, including those when accounting for non-employment transitions, turn out very similar. Our findings remain consistent and robust to this alternative price estimation.

[68] The bins are for ages 25–29, 30–34, 35–39, 40–44, 45–49, 50–54, and 55–59.

Table G.1: Full Model Pooled Sub-Periods (OLS–IV)

		Dependent Variable: Δe_j			
		Unrestricted model	Restricted model		
		(1)	(2)	(3)	(4)
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	3.90			
		(0.67)			
heterogeneous own effect:	$(d_{jj} - \bar{d}_{diag})\Delta p_j$	4.04	4.01	3.18	4.17
		(0.83)	(0.56)	(0.51)	(1.32)
total cross effect:	$\sum_{k\neq j} d_{jk}\Delta p_k$	3.69			
		(1.09)			
R-squared		0.492	0.491	0.791	-
Number of occupations		600	600	600	600
Estimation method		OLS	OLS	FE	IV
F-statistic 1st Stage		-	-	-	13

Notes: The table presents the estimates from different versions of eq. (63). Pooled panel sample containing 600 observations (120 occupations × 5 sub-periods). Sub-periods are: 1985–1990, 1990–1995, 1995–2000, 2000–2005, and 2005–2010. All regressions include dummies for the respective five-year estimation period. The regressor in columns (2)–(4) is the full $\sum_k d_{jk}\Delta p_k = \mathbf{d}_j\Delta\mathbf{p}$ and in column (4) this is instrumented by $\mathbf{d}_j V\mathbf{r}$ (see eq. (19)). Column (3) uses occupation fixed effects. Observations weighted by occupation j's initial employment size (e.g. for the period 1985-1990, this is 1985; for the 2000-2005 period, this is 2000, and so on). Standard errors clustered at the occupation level in parentheses; all coefficients shown are significant at the 1% level.

Figure G.1: Occupational Price and Employment Changes (by Own-Price Elasticity Median Split)

(a) Pooled Sub-Periods. 600 Occupations × Sub-Periods

(b) By Sub-Period

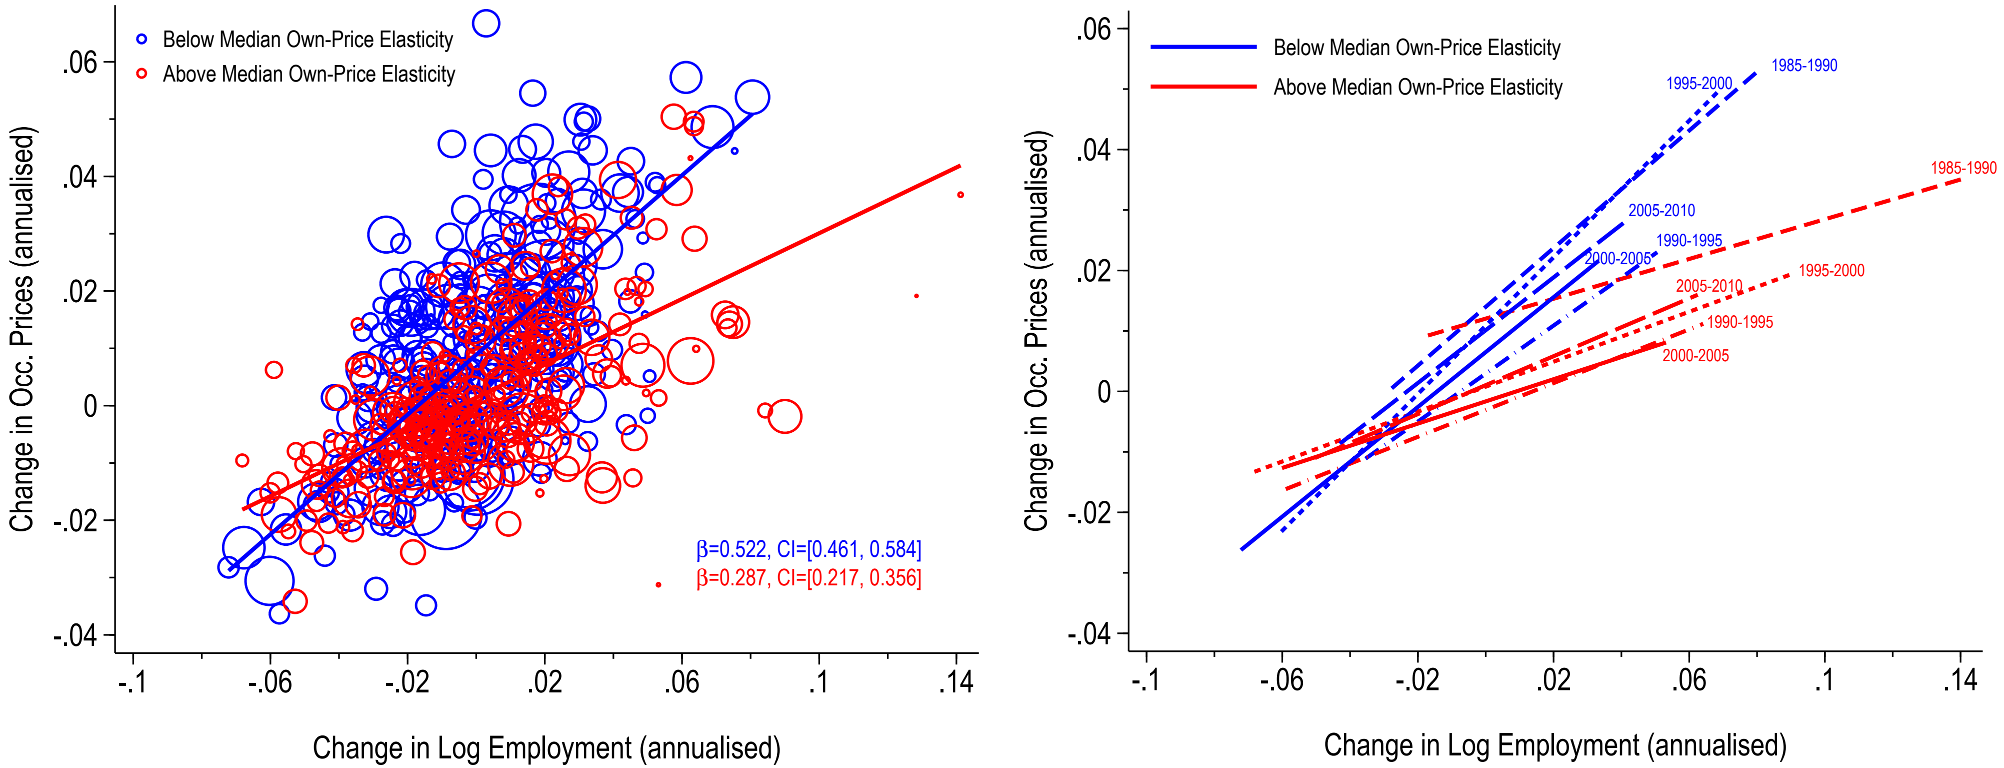

Notes: The figure shows the lines from an occupation-size weighted regression of price change on employment change, split by occupations below (blue, inelastic) and above (red, elastic) the median own-price elasticity (d_{jj}). Figure G.1a shows this for the pooled sample of 600 occupation–sub-periods. Figure G.1b shows this separately for each sub-period. Sub-periods are: 1985–1990, 1990–1995, 1995–2000, 2000–2005, and 2005–2010. β refers to the slope coefficient and CI stands for the 95% confidence interval. Marker size indicates the baseline employment in each occupation.

Figure G.2: Occupational Prices à la Cortes (2016) and Employment

(a) Overall Relationship

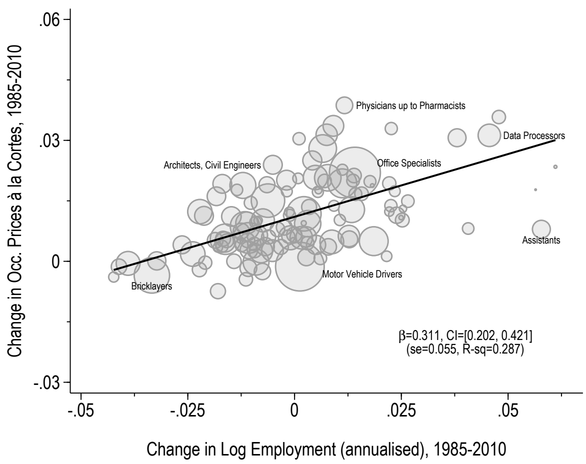

(b) By d_{jj} Median Split

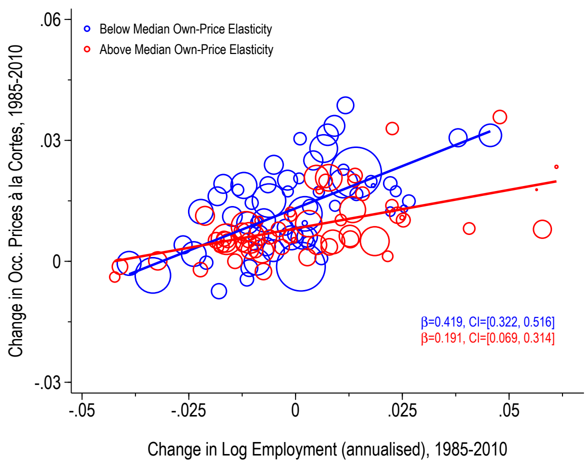

(c) By d_{jj} Quartiles

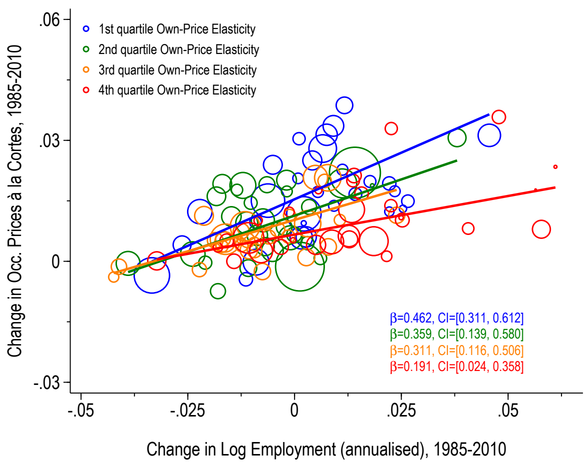

(d) IV Second-Stage: Inverse Supply Curve

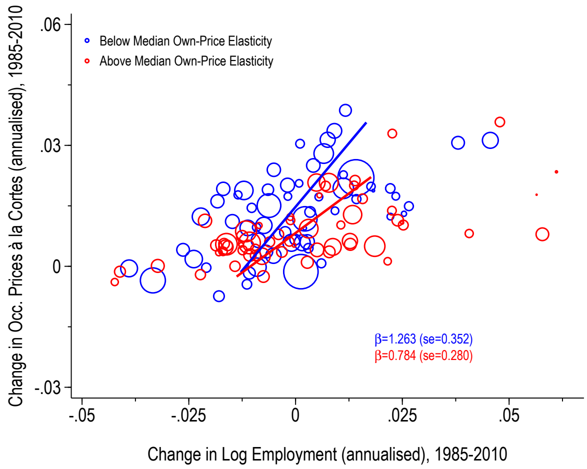

(e) Pooled Sub-Periods. By d_{jj} median split

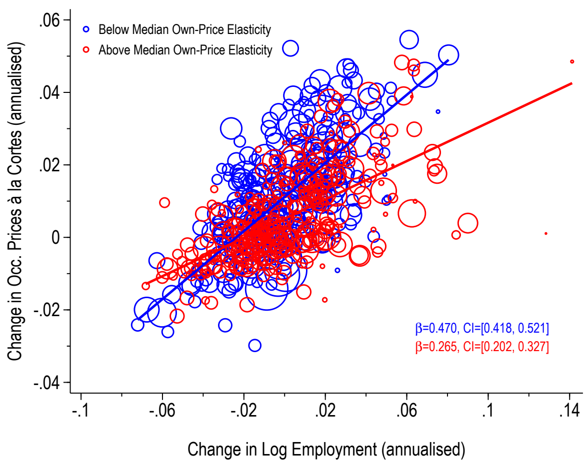

(f) Sub-Periods. By d_{jj} median split

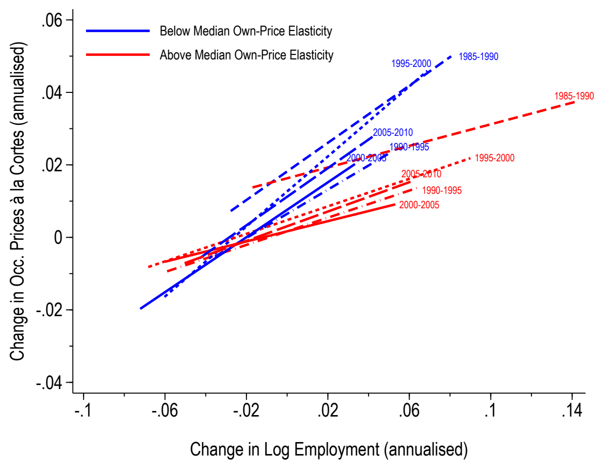

Notes: Panel (a) shows the line from an occupation-size weighted regression of price change on employment change. Panel (b) shows a split by occupations below (blue, inelastic) and above (red, elastic) the median own-price elasticity d_{jj}. Panel (c) shows a split by occupations in the lowest (blue), second (green), third (orange), and highest (red) quartile of d_{jj}. Panel (d) shows, by d_{jj} median split, the IV-2SLS second-stage of occupations' price on employment changes using initial task contents as the instrument. Panel (e) shows the overall regression line for the pooled 600 occupations × sub-periods case. Finally, panel (f) splits by d_{jj} median the pooled occupations × sub-periods sample. β refers to the slope coefficient, *CI* to the 95% confidence interval, *se* refers to standard error, and *R-sq* stands for R-squared of the regression. Marker size indicates the baseline employment in each occupation.

Table G.2: Occupational Prices à la Cortes (2016) and Changes in Employment. Main Results.

		Dependent Variable: Δe_j						
		Unrestricted Model	Restricted Full Model					
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
fixed relationship:	$\overline{d}_{diag}\Delta p_j$	4.46						
		(1.30)						
heterogeneous own effect:	$(d_{jj} - \overline{d}_{diag})\Delta p_j$	4.65	5.18	6.48	4.76	5.45	4.43	4.92
		(1.73)	(1.15)	(2.12)	(1.10)	(1.78)	(0.70)	(1.56)
total cross effect:	$\sum_{k \neq j} d_{jk}\Delta p_k$	3.23						
		(1.81)						
R-squared		0.371	0.350	-	0.402	-	0.486	-
Number of occupations		120	120	120	120	120	600	600
Estimation method		OLS	OLS	IV	OLS	IV	OLS	IV
F-stat 1st Stage		-	-	10	-	23	-	11
Accounting for non-employment transitions		no	no	no	yes	yes	no	no
Analysis pooling five-year sub-periods		no	no	no	no	no	yes	yes

Notes: Regressor in columns (2)–(7) is $\sum_k d_{jk}\Delta p_k$, i.e. corresponding to the full model. In columns (3), (5), and (7), regressor $\sum_k d_{jk}\Delta p_k = \mathbf{d}_j\Delta\mathbf{p}$ is instrumented by $\mathbf{d}_j V\mathbf{r}$ (see eq. (19)). In columns (4)–(5), we consider $M = 3$ different non-employment sectors: unemployment, out of the labour force (during the career and including part-time as well as employment with benefit receipt), and entry or exit due to newly joining the labour force at age 25–32 or retiring at age 52–59. In columns (6)–(7), we use the pooled panel sample containing 600 observations (120 occupations $\times$ 5 sub-periods). Sub-periods are: 1985–1990, 1990–1995, 1995–2000, 2000–2005, and 2005–2010. These regressions include dummies for the respective five-year estimation period and cluster standard errors at the occupation level. Observations weighted by occupation j's initial employment size.

H Further Analysis on the Future Projections

This section presents further analysis to complement the one in Section 7.

H.1 Descriptive Statistics in the Prediction Sample

As discussed in the main text, we use our prediction sample during 2012–2021 as well as the estimated aggregate parameters ($\theta = 4.8$ and $\sigma = 2.10$) to construct new matrices D (i.e. the elasticity matrix) and V (i.e. the spillover matrix). Parallel to the historical period, we select workers aged 25–59 but now include females, Eastern Germans, and part-time workers in the sample.[69] The latter are weighted to one-half of a full-time worker in the construction of all statistics. The prediction sample also features a new classification system (KldB 2010) of 126 occupations with more recent job titles and finer detail in rising occupational fields such as information technology, education, and health.

We supplement these data with estimates about what share of tasks in each occupation could in principle be replaced with current frontier technologies. The underlying information is from a triennial survey of experts run by the IAB and we use the most recent available wave in 2022. Based on it, the IAB also publishes 'Job Futuromat', an online tool that advises graduates and job seekers about which occupations will face rising (declining) demand in the future. We scale the share of tasks that may be replaced by such enhanced automation to have mean zero and the same standard deviation as the demand shocks backed-out over the period 1985–2010. That is, we analyse the impact of a specific set of predicted, scaled, and cardinally interpreted relative automation shocks. Although we thereby leave out other potential drivers of occupational labour demand or supply, our model could be used to analyse such shocks, too, as long as data are available.

Table H.1 shows summary statistics on the 126 occupations. The top panel presents the distribution of occupational employment sizes, mean age and fractions of females, Eastern Germans, part-timers, and workers with a university degree. The middle panel shows the distribution of the elements of the elasticity and spillover matrices D and V. These do not look too different from the historical period in Table E.3. The bottom panel summarises the demand shocks used that stem from predicted enhanced automatibility. We also report the shocks scaled in terms of effects on employment and wages under a homogenous labour supply elasticity (D_{hom}). Figure 5a in the main text displays these changes of equilibrium outcomes in occupations resulting from the automation demand shocks, which are computed using D_{hom} and parallel to those shown in Section 6.

[69]This sample consists of approximately 770,000 unique individuals and 5.5 million individual $\times$ year observations for the period 2012–2021.

Table H.1: Summary Statistics for the 126 Occupations.

	Mean	Std.Dev.	Variance	p10	p50	p90	Observ.
Occupational Characteristics							
Stationary employment size	0.794	1.274	1.625	0.033	0.304	2.383	126
Employment size in 2021	0.794	1.133	1.285	0.042	0.393	2.096	126
Share of Females (%)	39.54	28.23	796.88	3.82	36.11	78.11	126
Share of Eastern Germans (%)	20.57	6.79	46.13	13.15	19.69	26.99	126
Share of Part-time workers (%)	18.66	15.25	232.53	3.26	14.80	40.16	126
Share of University Degree (%)	27.26	26.39	696.23	2.64	18.88	68.04	126
Mean Workers' Age	42.42	2.25	5.06	39.54	42.74	45.14	126
Elasticity and Spillover Matrices							
Diagonal elements of matrix D (d_{jj})	0.448	0.138	0.019	0.252	0.460	0.612	126
Off-diagonal elements of matrix D ($-d_{jk} \times 100$)	0.358	0.938	0.882	0.003	0.066	0.866	15,750
Diagonal elements of matrix V (v_{jj})	0.502	0.087	0.007	0.415	0.484	0.632	126
Off-diagonal elements of matrix V ($v_{jk} \times 100$)	−0.401	0.866	0.750	−1.159	−0.136	0.007	15,750
Demand Shocks and Implied Annualised Employment and Price Changes under D_{hom}							
Demand shocks	0	0.019	0.004	−0.025	−0.0004	0.028	126
Log Employment	0	1.955	3.821	−2.480	−0.042	2.867	126
Log Prices	0.019	1.006	1.012	−1.257	−0.002	1.495	126

Notes: The top panel of the table presents summary statistics for occupational characteristics, such as the stationary employment size or share of females. Recall that, when calculating the shares, part-timers are weighted to one-half of a full-time worker (this also reduces the share of females here). The middle panel shows the distribution of the elements of the elasticity and spillover matrices D and V. The first line of the final panel summarises the predicted enhanced automation demand shocks, scaled to mean zero and standard deviation of past shocks as described in the text. The next two lines display the impact of the shocks on employment and wages when occupational labour supplies are homogenous (i.e., under D_{hom} and de facto an alternative scaling of the shocks).

H.2 Employment Flows between Occupations

A new analysis in the predictions is the projected employment flows between occupations. First, we start with a formal representation of how the equilibrium price changes in occupation j come about. From eq. (13), we have at the level of the occupation j that:

$$\begin{aligned}\Delta p_j &\approx \sum_k v_{jk}\Delta b_k \\ &= v_{jj}\Delta b_j + \sum_{k\neq j} v_{jk}\Delta b_k \\ &= \underbrace{v_{jj}\Delta b_j}_{\text{own occupation effect}} + \underbrace{\sum_{k\in V_j^+} v_{jk}\Delta b_k}_{\text{near occupations}} + \underbrace{\sum_{k\in V_j^-} v_{jk}\Delta b_k}_{\text{distant occupations}} \end{aligned} \tag{65}$$

where v_{jk} are the individual elements of V which translate demand shocks in occupation k into contributions to wage changes in j. The products $v_{jk}\Delta b_k$ can thus interpreted as 'price pressures' of k on j. The second line of eq. (65) uses the fact that V is of rank $N-1$ and each row sums to zero across columns. The third line splits the price pressures into those coming from near ($v_{jk} > 0$ for V_j^+) as opposed to distant ($v_{jk} < 0$ for V_j^-) occupations. In the data, near occupations are often also with high substitutability and cross-elasticity of occupation j with respect to occupation k.

Table H.2 shows the price pressures coming from the respective five closest occupations k (i.e. most positive v_{jk}) for a selection of impacted occupations j. Positive price pressures on IT consulting and sales arise from positive demand shocks on nearby Software developers and IT network engineers, who compete with it for workers, whereas Electrical engineers and Technical developers experience a somewhat negative demand shock. In general, for most occupations, price pressures are quite spread over a range of impacting occupations. One exception is Laboratory occupations in medicine, which is close to Human medicine and dentistry (high v_{jk}) and where the latter experiences a very positive demand shock (as seen e.g. in Figure 5). As a result, the price pressure of Human medicine on Laboratory medicine is large, compensating about two thirds of the automation demand shock's negative own effect on wages in Laboratory medicine.

Table H.2: Predicted Wage Change and Price Pressures from Five Nearest Occupations

	Wage change			Five most positive v_{jk} occupations	Price pressure
	Total	Own effect	Cross effect		
IT analysis, IT consulting, IT sales	0.54	0.66	-0.11	Software development and programming	0.019
				IT-network engineering, -coordination, -admin, -organisation	0.023
				Computer science	0.000
				Electrical engineering	-0.008
				Technical research and development	-0.004
Machine building & operating	-1.25	-0.87	-0.38	Metalworking	-0.029
				Metal constructing and welding	-0.013
				Tech occup in energy technologies	-0.005
				Tech occup in automotive, aeronautic, aerospace, shipbuilding	-0.004
				Plumbing, sanitation, heating, ventilating, air conditioning	-0.003
Laboratory occupations in medicine	0.03	-0.35	0.38	Human medicine and dentistry	0.244
				Doctors' receptionists and assistants	0.001
				Biology	0.008
				Other occupations	0.003
				Teachers and researchers at universities and colleges	0.002
Building construction	0.84	1.01	-0.17	Civil engineering	0.041
				Construction scheduling and supervision, and architecture	-0.012
				Interior construction, dry walling, insulation, carpentry, glazing	0.016
				Building services engineering	-0.015
				Painters and varnishers, plasterers, and related	0.013

Notes: The table displays the predicted equilibrium wage changes, as well as the constituting own and cross effect, due to enhanced automation shocks of four selected occupations j. Each panel then shows, as contributors to the cross effects, the price pressures emanating from the five nearest occupations k (i.e. those with the highest v_{jk}). For details on how the price pressures are constructed, as well as definitions of own and cross effects, see the text. All the numbers are shown in per cent changes per year.

We use the projected price changes to compute the equilibrium employment flows between occupations. From eq. (15) (even more directly, eq. (24)), we obtain a first-order approximation for the change in occupation j's share of total employment as: [70]

$$\begin{aligned} \Delta E_j &\approx \theta E_j^{base} \textstyle\sum_k d_{jk} \left(\Delta p_k - \Delta p_j\right) \\ &= \theta E_j^{base} \textstyle\sum_{k \neq j} d_{jk} \left(\Delta p_k - \Delta p_j\right) \\ &= \underbrace{\theta E_j^{base} \sum_{\Delta p_k \geq \Delta p_j} d_{jk} \left(\Delta p_k - \Delta p_j\right)}_{\text{net flows out of occupation } j} + \underbrace{\theta E_j^{base} \sum_{\Delta p_k < \Delta p_j} d_{jk} \left(\Delta p_k - \Delta p_j\right)}_{\text{net flows into occupation } j} \end{aligned} \tag{66}$$

where $\theta E_j^{base} d_{jk} \left(\Delta p_k - \Delta p_j\right)$ captures the net flows from occupation k to j. The second line then uses the fact that also in D each row sums to zero across columns, and the third line splits the summation into net outflows versus inflows to occupation j. Finally, we scale these annualised changes to obtain 10-year flows.[71]

Figure 6 in the main text displayed the main bilateral flows resulting from eq. (66) for IT consulting and sales. Panel (a) of Figure H.1 displays the corresponding chart for Machine building and operating, which turns out substantially less elastic. This is because highly substitutable occupations with it, like Metal workers, Draftspersons, or Welders experience at least as negative demand shocks. As a result, wages of Machine operators will drop substantially and they are projected to move to somewhat more distant occupations like Business organisation, Driving, and Construction.

Panel (b) of Figure H.1 shows Laboratory occupations in medicine which, as we saw before, endures a substantially negative demand shock. While many health-related occupations are relatively inelastic (see also Figure 5), this occupation differs because of workers' ability to move to Human medicine. Accordingly, wages of Laboratory occupations in medicine remain broadly constant and all the other flows, even to quite substitutable occupations such as Nurses or Doctors' receptionists and assistants, are small.

[70] Alternatively, in matrix notation, the net flows from all k to occupation j can be expressed as

$$\theta \left(\mathbf{E}^{base} * \mathbf{1}_N^T\right) \circ D \circ \left(\mathbf{1}_N * \Delta\mathbf{p}^T - \Delta\mathbf{p} * \mathbf{1}_N^T\right)$$

where $\circ$ is the element by element Hadamard product, $\mathbf{E}^{base}$ a column vector of employment (shares), and $\mathbf{x}^T$ the transpose of vector $\mathbf{x}$. In more compact notation this can be written as

$$\mathcal{N} \equiv W \circ \theta D \circ P$$

where $\mathcal{N}$ is the matrix of net flows, $W \equiv \mathbf{E}^{base} * \mathbf{1}_N^T$ is the stacked matrix of shares, that we have used before, and $P \equiv \mathbf{1}_N * \Delta\mathbf{p}^T - \Delta\mathbf{p} * \mathbf{1}_N^T$ is the skew-symmetric matrix of relative price changes. Then the matrix of net flows, $\mathcal{N}$, is also skew-symmetric. (The net flow from k to j is the negative of the net flow from j to k.) One can look at just the positive net flows by examining only the positive entries of $\mathcal{N}$.

[71] To be precise, eq. (66) constructs a counterfactual: how employment shares will shift between occupations when prices change compared to if prices stayed the same as in the steady state.

Figure H.1: Projected Net In and Out Flows by Top Senders and Receivers (%)

(a) Machine Building and Operating

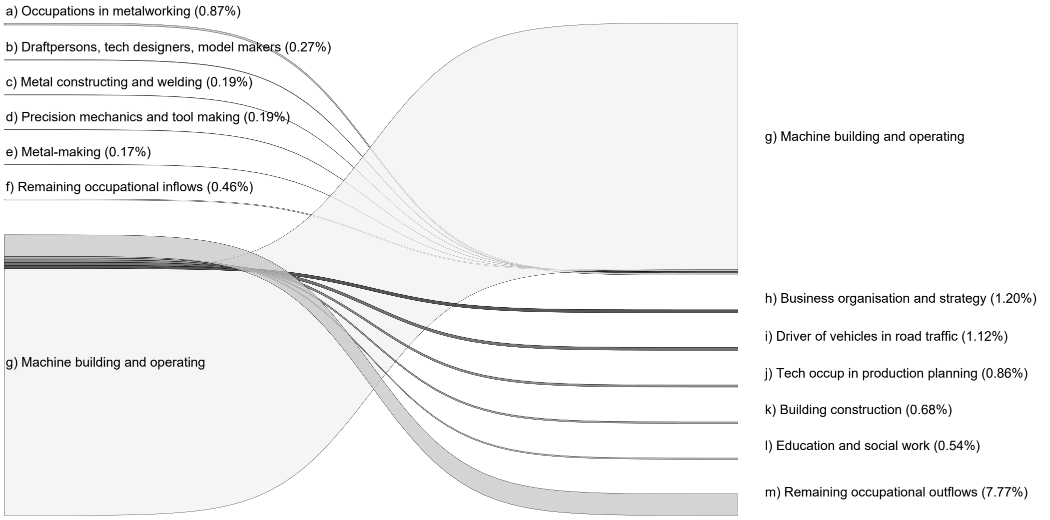

(b) Laboratory Occupations in Medicine

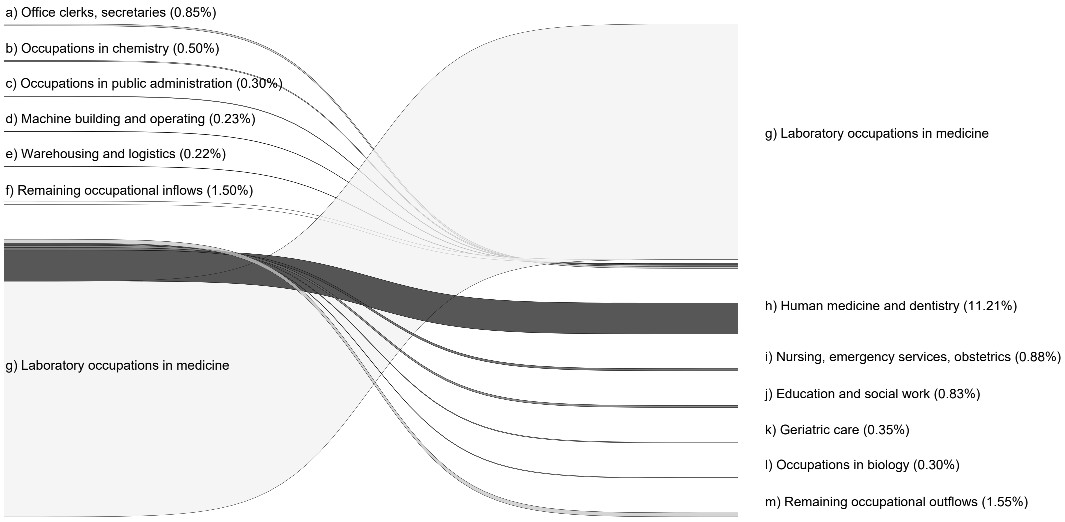

(c) Building Construction

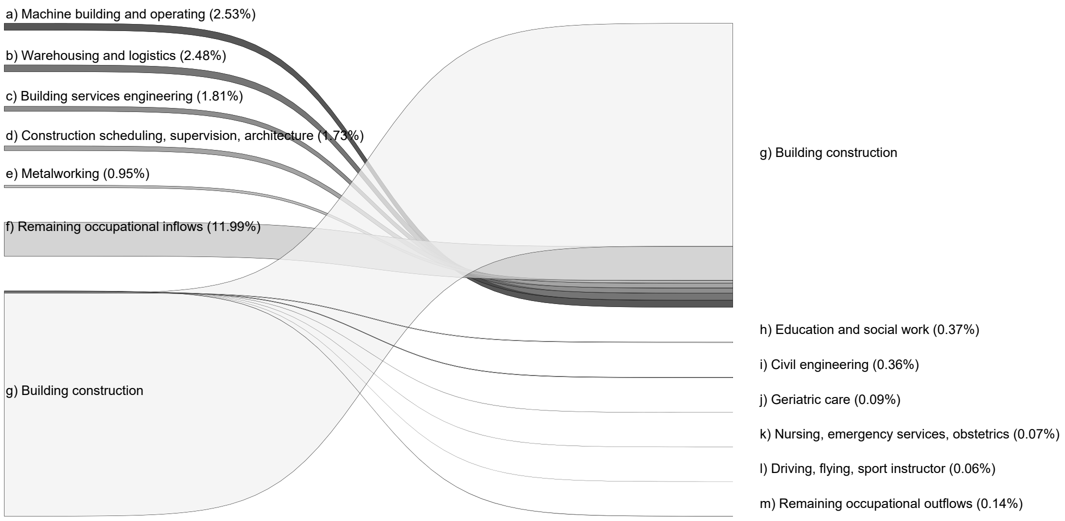

Notes: The figure depicts the projected flows over 10 years of the respective occupation with its top net sending and receiving occupations. The numbers shown are percentages of final (for inflows) and initial (for outflows) employment. Panel (a) shows Machine building and operating, which makes up 3.70% of total employment initially and 3.32% after ten years, i.e. is projected to shrink by around 10% as a share of its initial employment. Panel (b) shows Laboratory occupations in medicine, which makes up 0.39% of total employment initially and 0.35% after ten years, i.e. is projected to shrink by around 10%. Finally, Panel (c) shows Building construction, which makes up 0.79% of total employment initially and 0.99% after ten years, i.e. is projected to grow by over 25%.

Finally, Panel (c) of Figure H.1 shows the main bilateral flows of the Building construction occupation. We see that this is able to draw in workers from close occupations in construction as well as a range of other declining occupations including Metalworking, Machine operating, or Warehousing and logistics. Accordingly, Building construction turns out effectively quite elastic (see also Figure 5) and is projected to grow by more than a quarter over 10 years.