Heterogeneity of Occupational Supply Elasticities and Changes in Labour Demand *

Michael J. Böhm† Ben Etheridge‡ Aitor Irastorza-Fadrique§

This version: February, 2024
PLEASE DO NOT CIRCULATE OR CITE.

Abstract

We propose a theoretically founded and interpretable measure of labour supply elasticities across occupations that can be implemented using observed transition rates. We use these to study the heterogeneous impact of demand shifts in Germany during the past decades within an equilibrium model. Employment growth per unit of wage growth is twice as large for occupations ex-ante classified as relatively elastic compared to inelastic occupations. We find that cross-price effects, capturing the role of wage changes in close substitute occupations, are particularly important in explaining employment growth heterogeneity. We validate the estimated elasticities with external correlates, including occupational licensing and task distance. The model explains substantially more of the structural changes during 1985–2010 than an approach with homogeneous labour supplies.

JEL Classification: J21, J24, J31

Keywords: Labour Supply Elasticities, Heterogeneity, Occupational Similarity, Occupational Changes

*We thank Wolfgang Dauth, Emma Duchini, Khalil Esmkhani, Giovanni Gallipoli, Georg Graetz, Florian Lehmer, and Oskar Nordström Skans as well as seminar participants at Uppsala, Darmstadt, Duisburg-Essen, RWI, Mercator School of Management, DFG SPP 1764 Conference in Mannheim, the Tasks VI, the IAB-FDZ User Conferences in Nürnberg, the EALE in Prague, and the FIT Labour Workshop in Helsinki for very 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. 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 are our own.

†TU Dortmund University and IZA. Email: michael.j.boehm@tu-dortmund.de

‡University of Essex, IFS and MiSoC. Email: bsethe@essex.ac.uk

§Institute for Fiscal Studies (IFS). Email: aitor.irastorza-fadrique@ifs.org.uk

1 Introduction

Today it is well-established that sustained shifts in the demand for occupations have led to large changes in employment and wages, with many economic and societal consequences.[1] These shifts in demand have been attributed to a variety of sources, including changes to patterns of trade (see e.g. Autor et al., 2013), to households' consumption of goods (Mazzolari & Ragusa, 2013) and to technology (Dauth et al., 2021; Acemoglu & Restrepo, 2022). In adapting to these shifts, a key concern for policy makers is how responsive is labour supply. While meeting an increased demand for some occupations may be straightforward, supplying workers in other occupations may be much harder. Equally, workers may reallocate from some occupations in decline more easily than from others. Indeed, research has found that the reallocation and supply of labour across sectors is heterogeneous (e.g. Cortes, 2016; Caliendo et al., 2019). The aim of this paper is to characterise the role of this labour supply heterogeneity in explaining the evolution of employment and wages over recent decades.

This paper's main contributions are as follows. (i) We develop a tractable equilibrium model of the labour market that captures heterogeneity in supplies across occupations in a simple yet flexible way. These labour supply functions account for variable substitutabilities across occupations, which induce heterogeneous occupational spillovers to shocks. An important feature of the model is that the labour supply elasticities can be estimated directly from patterns of worker flows. (ii) We document the empirical relevance of the estimated elasticities, displaying their distribution, their relationship to existing metrics used in the literature and, accordingly, their potential for future application. (iii) Using the German labour market as a laboratory, we show that the resulting heterogeneous supplies are quantitatively important for explaining the evolution of occupational employment and wages. In particular, we find a key role for the joint distribution of occupational spillovers with observed demand shifts. (iv) In a related way, we show that neglecting these heterogeneous spillovers leads to a substantial *understatement* of the response of occupational employment to wage changes, as studied in, for example, Mishel et al. (2013) and Hsieh et al. (2019).

As such, we contribute to an extensive literature which has sought to provide a rich characterisation of the distribution of labour market changes, most notably including job

[1] Acemoglu & Autor (2011) provide a comprehensive analysis of the evidence on the labour market over several decades (see also Autor, 2019). For recent analyses of societal consequences beyond the labour market, see, among others, Autor et al. (2020) on political polarization, Adda & Fawaz (2020) on health, and Keller & Utar (2022) on family structure.

'polarization'.[2] Moving forward, our results are important for understanding the likely effect of ongoing changes to demand, such as through advanced automation (Felten et al., 2018; Webb, 2020).

The supply side of our framework is based on a random utility model of workers' occupational preferences related to, among others, Cortes & Gallipoli (2018) and Card et al. (2018). In this model, the choices of occupation can be solved for as standard probability formulae which depend on wages in a destination occupation as well as pairwise occupational switching costs. We show that these probabilities, together with occupational employment shares, are sufficient statistics for the elasticity of each occupation's employment with respect to occupational prices. Given that transition probabilities vary substantially across occupations, the model provides an intuitive reason for why these supply elasticities vary. Generally, the framework is highly tractable and its mechanisms have intuitive economic interpretations, which we explore in detail.

As comes naturally out of the model, we distinguish between 'cross-price' elasticities, which capture the impact on employment of changes in the wage in a *different* occupation, and 'own-price' elasticities, which capture the impact of wage changes in the occupation itself. The model can then be used to theoretically assess the outcome of a set of wage changes across the whole economy. We decompose the predicted employment changes into those coming from own-price and total cross-price *effects*. The own-price effect is determined both by the own-price elasticity and the size of the wage change. Similarly, the total cross-price effect depends on the interaction of cross-price elasticities and outside wage changes: the resulting outcome depends subtly on whether close substitute occupations see wage increases or declines. We find that it is the spillovers resulting from this total cross-price effect which are of particular quantitative importance. The heterogeneous spillovers that arise are typically missing from related analyses, such as those relating to firm-level distortions in, for example, Card et al. (2018), Lamadon et al. (2022) and Berger et al. (2022).

We apply the model using data from Germany, which are uniquely suited for the purpose. In particular, we use the Sample of Integrated Employment Biographies (SIAB), which follows workers over their entire labour market careers for the years 1975–2010 and provides a consistent and fine-grained set of 120 occupations over this period. We estimate the supply elasticities using occupations' employment sizes and workers' transition flows over 1975–1984, in five-year rolling windows. These display substantial heterogeneity: Own-price elasticities vary by a factor of ten between the most elastic (e.g.,

[2] The literature characterizing labour market polarization is extensive. See the discussion towards the end of this Introduction.

Nursery teachers, Other occupations attending on guests) and the most inelastic (Physicians and pharmacists, Bank and building society specialists) occupations. A large part of this is how dispersedly they source employees. Lower employment elasticities are also correlated with occupations' certification requirements and regulations (taken from Vicari, 2014), larger share of university graduates, and higher average age.

Turning to the cross-price elasticities, we find that these are distributed approximately log-normally. As such, although most elasticities are close to zero, there exists a right tail of very similar occupation pairs, often within broader groups (such as Nursery teachers and Social work teachers, or Carpenters and Concrete workers), that lead to high elasticities of employment in one with respect to wage changes of the other. Overall the cross-price elasticities are strongly correlated with task distances between occupations, measured as in Gathmann & Schönberg (2010); Cortes & Gallipoli (2018). However, our empirical measure improves on task distance by having a cardinal quantitative interpretation.

We then use the estimated elasticities to examine wage and employment changes over 1985–2010. To explore mechanisms and aid intuition we first apply the model on the supply side of the labour market only. We begin by documenting a striking relationship between occupations' own-price elasticity and ex-post outcomes: as would be expected, more elastic occupations display significantly higher employment growth per unit of wage growth than do less elastic occupations. We then examine the full effect of labour supply heterogeneities on employment in a simple regression specification. Both own-price and, particularly, cross-price effects are important in explaining the observed patterns of employment changes. Using model R-squared as a simple metric, we find that the explanatory power for employment changes is 33% higher than in standard specifications with homogeneous labour supplies only. By relaxing a key restriction in this regression, we also implement a simple test of the model, which is clearly passed.

Of particular note, these results highlight the importance of allowing for cross-price effects in this type of analysis. Because substitutable occupations experience similar wage shocks, omitting cross-price effects gives a misleading impression of the effect of wages on occupational employment. Specifically, the simple relationships analysed by e.g. Autor et al. (2008); Mishel et al. (2013); Hsieh et al. (2019); Böhm et al. (2024) *understate* the ceteris-paribus response of employment to sectoral wage changes by over 50%. This omitted variable result relates to the similar point made by Borusyak et al. (2022) on the responsiveness to wage shocks of migrants' location.

To analyse labour market equilibrium, we add the demand side to the model using an intentionally-homogeneous CES aggregator. In terms of the theory, we find that the effect of shocks on both sides of the market can be expressed in compact form, featuring a single additional matrix which captures how shocks to *either* demand or supply dissipate across the labour market heterogeneously. Importantly, in terms of empirics, the model also points directly to a strategy to isolate demand shocks with instrumental variables. Following the literature on routine-biased technical change (see e.g., Autor et al., 2003, among many others), we base our instruments on occupations' average task content in the late 1970s and early 1980s. We then interact this we the predicted effects that it has depending on the different supply elasticities. These instruments are strong and, as we show in detail, yield similar conclusions to those from the pure supply-side analysis.

We contextualise these results by using the equilibrium model to decompose the effect of supply heterogeneity from underlying shocks. In terms of the shocks, we show that those on the demand side have been substantially more important than supply shocks over the period we study. This feature largely explains why our analysis of the supply side provides insight on its own. Additionally, we document a noticeable U-shaped pattern for both types of shocks across the distribution of occupations, indicating a clear combined polarizing effect on employment, but with more complex, off-setting effects on wages. In terms of supply heterogeneity, we again show the importance of cross-effects. Overall, and given the ex-post set of demand shocks, we find that supply heterogeneity and supply shocks each explain around half of the total observed variation in occupational employment.

This paper contributes to the analysis of occupational changes. A large body of research has studied whether and what kinds of demand shocks have worked on the occupation and task structures, and what effects this has had on employment (notably job polarization) and wages.[3] We advance this literature by highlighting that a fundamental catalyst of such changes is the flexibility of labour supply to react to them. These results inform a broader debate about how the labour market will generate the jobs of the future. Autor et al. (2023) discuss how institutions and policies may be designed to (re-)train workers in the skills that are needed. Autor et al. (2022) show how new occupations and job types emerge from labour-augmenting and automating innovations. We complement this agenda by studying the ability to shift employment among the existing set of occupations and skills. While our empirical application is oriented around changes to technology, our framework could equally be applied to other demand-side changes.

Our theory extends standard models of sector choice by allowing for variation in the costs to transition between occupation pairs. Cortes & Gallipoli (2018) is a notable precur-

[3]See, e.g., Spitz-Oener (2006); Autor et al. (2008); Acemoglu & Autor (2011); Autor & Dorn (2013); Autor et al. (2013); Goos et al. (2014); Deming (2017); Bárány & Siegel (2018) in addition to the papers cited above.

sor. We show that this leads to heterogeneity of labour supply elasticities being identified directly from job flow data.[4] We then highlight the role of the substitutability between the sectors that workers choose from as driving these elasticities. In this sense, and with the resulting importance of cross-occupation effects, the analysis complements research that focuses on employer size to generate heterogeneity in own-wage labour supply elasticities or market power (Berger et al., 2022; Jarosch et al., 2019).

More generally, we complement a micro-economic research agenda on the labour–supply-side substitutability between occupations. A series of studies have provided advice to job seekers about which alternative but related occupations to their previous employment they should search in (e.g., Belot et al., 2019, 2022; Altmann et al., 2022). Gathmann & Schönberg (2010) analyse the importance of task distance and Eckardt (2023) the specificity of training for the costs of switching occupations. Borusyak et al. (2022) highlight the econometric issues that arise in migration regressions when not taking into account the correlation of shocks to workers' current and substitutable region-industries. We formalise the mechanisms that may underpin such relationships between occupations or sectors, and then study their role for more aggregate wage and employment outcomes.

Following this introduction, Sections 2-6 of this paper are structured as described above. Section 7 then summarises robustness checks, which include modeling workers' transitions between occupations and non-employment states, estimating the model separately by sub-period, and employing an alternative measure for (selection-corrected) occupational wage growth. Extensive appendices contain further theoretical results, proofs, and discussions of the model, together with additional empirical results and details of the data and empirical methodology.

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. In this section, we present a static partial equilibrium model with perfect information, providing a tractable framework for labour mobility decisions under frictions. Labour demand and market equilibrium are modelled in Section 5.

[4]Alternatively, e.g., Bhalotra et al. (2022) exploit the full set of women and men's employment and wages across broad occupation task groups for equilibrium identification.

2.1 Environment

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:

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

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 (we will later simplify our language and refer to this as 'price') 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. Gathmann & Schönberg (2010) and Cortes & Gallipoli (2018) analyse these costs explicitly – we further discuss this in Section 3 – while we let them flexibly affect the labour supply functions that we are after.

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.[5] 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.

[5] 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.

2.2 Occupational Choice and Price Elasticities

By standard arguments (McFadden, 1973), the assumptions on eq. (1) imply that workers' occupational choice probabilities are of 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 will be 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 can omit the index ω from now on.

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 that

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

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 Effect of Individual Price Changes

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.

Equation (4) shows how these price elasticities can be computed using transition probabilities (as we discuss in the next section, the transition probabilities have direct analogues in the data as job flows), baseline employment shares, and an unobserved pecuniary parameter θ. We return to the estimation of θ in Section 4. We 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 of price elasticities, which we also refer to as 'elasticity matrix' or 'matrix D' 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.[6] To gauge the empirical content of Remark 1 further, we derive alternative formulations of these elasticities more explicitly in terms of moments of job flows. This provides further intuition on what determines the elasticities as well as metrics to compare to other related measures used in the literature.

To do this, we define some additional terms. First, and as standard, let $\mathbb{E}_\tau x \equiv \sum \tau_i x_i$ 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 $\mathbb{E}_\tau \tilde{\pi}_{iq} = 1$. Normalising the transition probabilities in this way yields moments that are invariant to occupation size. In this spirit, and in parallel, let $Cov_\tau (x, y) \equiv \sum \tau_i (x_i - \mathbb{E}_\tau x)(y_i - \mathbb{E}_\tau y)$. This leads us to the following result:

Remark 2 (Individual Cross-Price Elasticities) *For all $j \neq k$, the off-diagonal elements of matrix D can be expressed as:*

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

where we examine the negative of d_{jk}, rather than d_{jk} itself, so that we can interpret higher elasticities by larger positive numbers. Appendix A.1 contains the derivation.

Expression (6) above consists of two additive components. First is a substitutability component $\tau_k Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$. It consists of an 'occupational-similarity' term that is symmetric between j and k, is invariant to the fineness of the occupational classification, and

[6] Strictly speaking, these elements should be multiplied by θ as shown in eq. (4).

captures the pure similarity of occupation in-flows: If $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) > 0$, then occupations j and k are 'competing' for workers and the cross-price elasticity (i.e., the responsiveness of employment in occupation j to changes in the price of occupation k) will be higher. This occupational similarity term is then weighted by an 'occupational-importance' term τ_k that depends on the size of the occupation of the price change: Price increases in a smaller competing occupation will have smaller percentage ripple effects than price increases in a larger occupation. Second is an occupation-specific intercept which captures occupation k's contribution to a price index and which, in terms of variability across occupations, turns out to be quantitatively relatively unimportant.

Likewise, we can reformulate the on-diagonal elements of the elasticity matrix D, which capture the own-price elasticities. This leads us to the following result:

Remark 3 (Individual Own-Price Elasticities) *For all $j = k$, the on-diagonal elements of D can be expressed as:*

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

where $Var_\tau(x) \equiv \sum \tau_i \left(x_i - \mathbb{E}_\tau x\right)^2$. Appendix A.1 contains the derivation.

Expression (7) captures the effect of an isolated change in occupation j's own price and has a similar structure to expression (6), though in this case it can be constructively formulated in two ways, each of which provides informative interpretations. The first formulation includes an 'aggregate substitutability' term, that sums substitutability components from all other occupations. This term 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. It is in this sense that the own-price elasticity captures aggregate substitutability with other occupations.

In the second formulation, on the right-hand side of expression (7), the term $\tau_j Var_\tau\left(\tilde{\pi}_{.,j}\right)$ can be interpreted as a 'job-flow dispersion' term, reflecting how dispersed or concentrated are the inflows to occupation j: Sectors hiring from a diversity of sources (in this case, a *small* $Var_\tau\left(\tilde{\pi}_{.,j}\right)$) are more elastic. Following this line of thought, it is useful to consider 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. Finally, in both formulations are 'direct' and price-index effects. As in the discussion following Remark 2, these terms contribute to the *level* of the elasticity, but little to the observed variability.

Remarks 2 and 3 show that we can express the price elasticities in terms of simple moments of the distribution of job flows. Before moving on, it is worth commenting that in eq. (6) we conceptually separate $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ from τ_k in the first summand, while in eq. (7) we interpret $\tau_j Var_\tau\left(\tilde{\pi}_{.,j}\right)$ jointly. We formulate the expressions in this way because it is $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$ and $\tau_j Var_\tau\left(\tilde{\pi}_{.,j}\right)$ (rather than $Var_\tau\left(\tilde{\pi}_{.,j}\right)$) which are invariant to the fineness of the occupational classification. We discuss this point further in Appendix A.1 using both empirical evidence and theoretical justification.

2.2.2 Effect of Multiple Price Changes

We now generalise the formulation given in eq. (4). The response of the vector of employment shares to a change in the vector of prices can be approximated by:

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

with $\Delta \mathbf{e}$ representing 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 approximation 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 application, we rewrite the inner product of elasticity matrix D with the vector of price changes as follows:

$$\begin{aligned} \Delta e_j &\approx \theta \mathbf{d_j} \Delta \mathbf{p} \\ &= \theta \Bigg(\underbrace{d_{jj} \Delta p_j}_{\substack{\text{own-price} \\ \text{effect}}} + \underbrace{\sum_{k \neq j} d_{jk} \Delta p_k}_{\substack{\text{total cross-} \\ \text{-price effect}}} \Bigg), \end{aligned} \tag{9}$$

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-price effect in eq. (9) represents the part of occupations' employment changes that are due to their own price changing. The total cross-price 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 effects on j's employment share. We provide additional formal details in Appendix A.2.

The thought experiment we imagine is that the economy is hit by a sudden change in the technology of final goods production, which shifts demands for the different occupations' labour inputs either to the right (an increase) or to the left (a decrease).[7] This leads to a new set of equilibrium prices $\{p_j\}$ and quantities $\{e_j\}$ in all occupations. In the absence of supply shocks (eq. (9)), these price changes are sufficient statistics for the implied changes in demand – we consider this case in Section 4. In the presence of supply shocks ($\Delta\mathbf{e} \approx \theta D \Delta\mathbf{p} + \Delta\mathbf{s}$), the equilibrium model presented in Section 5 provides an identification framework which can be implemented with appropriate instruments.

3 Data and the Elasticity Matrix

3.1 Data Sources

Our first objective is to analyse the elasticity components and to estimate the labour supply curves given by eq. (9). To take the model to the data, we use the Sample of Integrated Labour Market Biographies (SIAB, Frodermann et al., 2021), a 2% sample of administrative social security records in Germany since 1975. The SIAB data contains complete employment histories and wage information for more than one million employees. This dataset is representative of all individuals covered by the social security system, roughly 80% of the German workforce. It excludes self-employed, civil servants, and individuals performing military service.

The SIAB data have been used in various prior studies of the labour market and are well-suited for our purpose. First, the panel dimension allows us to measure worker flows over long frequencies. The administrative nature ensures that we observe the exact date of a job change and the wage associated with each job. Second, occupation codes are consistently coded from 1975 to 2010 ($N = 120$ occupations). Since employers are legally required to report the kind of job their employees perform, miscoding of occupations is less likely than in the case of survey-based data collection. Finally, the wage information is highly reliable. The SIAB is based on process data used to calculate retirement pensions and unemployment insurance benefits, so misreporting is subject to severe penalties.

[7] The instrumental variables strategy in Section 5 will exploit occupations' initial task contents at the beginning of our analysis period as proxies for subsequent demand shocks. More generally, forces of occupational demand may include, among others, task-biased technological change and automation (e.g., Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2022), international trade and offshoring (Autor et al., 2013; Goos et al., 2014), transformation of the industry structure (Bárány & Siegel, 2018), changes in consumption patterns (Autor & Dorn, 2013; Mazzolari & Ragusa, 2013), or social skills content (Deming, 2017).

We restrict the main sample of analysis to men aged 25–59 who are working full-time (excluding apprentices and always-foreigners) in West Germany.[8] 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.[9] Following Böhm et al. (2024), we transform the daily spell structure of the SIAB into a yearly panel by using the longest spell in a given year. Our final sample consists of approximately 600,000 unique individuals and 9 million individual $\times$ year observations for the whole period 1975–2010.

Importantly, the SIAB data allows us to compute worker flows (sufficient statistics for the elements of D), changes in occupational employment ($\Delta\mathbf{e}$), as well as changes in occupational prices ($\Delta\mathbf{p}$). For the latter, we follow the literature on this, which emphasises that raw wages need to be selection-corrected (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 changes as the main estimate of changes in occupational prices. We show the robustness of our results using an alternative price estimation procedure following Cortes (2016) that corrects for worker–occupation-spell fixed effects in Section 7. The SIAB data also provides us with other occupational characteristics (e.g., workers' mean age by occupation, the share of workers with university degree by occupation) that we use to 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, and we categorise them into analytical, routine, and manual tasks, assigning values based on response frequency. By averaging responses from pooled QCS data in 1979 and 1985/1986, we compute task intensities among those three categories by occupation, which we also use to construct a measure of task distance between occupations following Gathmann & Schönberg (2010) and Cortes & Gallipoli (2018). We study how these relate to our elasticity measures below. In Section 5, we use task measures to proxy for demand changes across occupations between 1985–2010. Finally, to obtain measures of occupational licencing, we use the indicators for standardised certificates and degree of regulation developed by Vicari (2014).

We report summary statistics for the 120 occupations in Appendix Table B.1. This

[8]Excluding East Germans allows us to define a consistent sample during the whole 1975–2010. We also remove women and individuals who are always foreigners as these groups have experienced some strong and potentially confounding changes during this period (rapidly rising education and employment rates, declining workplace discrimination, changing norms; see e.g. Hsieh et al., 2019; Boelmann et al., 2023).

[9]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). See Appendix B for all the details.

shows that cross-sectional variation of employment is substantial, with occupations at the 10th percentile shrinking by 1.8 log points annually (averaged over the period 1985–2010) while growing by 2.4 log points annually at the 90th percentile (see also Figure 2 below). The annualised wage growth of occupation stayers is positive at 0.59 log points, again with considerable dispersion around this average. Similar variation is found for our alternative measure of occupational prices à la Cortes (2016). Using five-year subperiods, we also show that there is large variation in employment and wage growth over time which is, for example, slower in the economically sluggish early 2000s. More details on the data, variable construction, and descriptive statistics are presented in Appendix B.

3.2 The Elasticity Matrix

As shown in Remark 1, a strength of the elasticity matrix implied by the theory is that it can be computed directly from baseline worker flows. We construct transition rates across all occupation pairs for individuals who are observed 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, and element π_{ij} represents 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. With that, we compute the elasticity matrix D following eq. (5).[10]

Panel A of Table 1 reports occupations at different quantiles of the own-price elasticities d_{jj} (the full list for the 120 occupations is in Appendix Table B.5). These range from 0.07 among physicians and pharmacists to 0.80 among personnel in social, medical, and gastronomy service occupations. Figure 1a correlates the key component of own-price elasticities, 'aggregate substitutability', with education, age, task content, and occupational requirements.[11] The latter, taken from Vicari (2014), measure whether access to exercising a professional activity is linked to a standardised training certificate as well as the extent of legal and administrative regulations in the occupation. Figure 1a shows that

[10]The baseline period (1975–1984) sample consists of 252,309 individuals and 1,794,286 individual $\times$ year observations. Our findings 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). Appendix Table B.3 summarises the transition probability matrix Π and the elasticity matrix D.

[11]As discussed in relation to Remark 3, aggregate substitutability is the key component of own-price elasticities, dominating the price index component. Appendix B.3 reports the variation in the two. As such, the relationship of the own-price elasticity with external characteristics is almost the same as that of its substitutability component (Appendix Figure B.1).

occupations with a higher degree share, more analytical tasks, and higher regulation and certification requirements are less substitutable and by extension less own-price elastic. Together, the table and figure also suggest that occupations which are more sensitive to wages are easier to enter and have less specialised workforces (consistent with Cortes & Gallipoli, 2018), and that occupational licensing can have significant negative effects on worker labour market flows (Kleiner & Xu, 2024).[12]

Panel B of Table 1 shows cross-price elasticities ($-d_{jk}$) between occupation pairs. The highest effects of price changes on employment are naturally among related occupations: of 'home wardens, social work teachers' on 'nursery teachers, child nurses'; of 'non-medical practitioners, masseurs, physiotherapists' on 'medical receptionists'; and of 'office specialists' on 'stenographers, shorthand typists, data typists'. While quantitatively these top pairs are within the range of the own-price elasticities, 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 1b plots the occupational similarity component of the cross-price elasticities against occupational task distance as in Gathmann & Schönberg (2010) and Cortes & Gallipoli (2018). Because it abstracts from the role of occupational importance, 'occupational similarity' is the fitting comparison to task distance as both are symmetric and size-independent. Occupational similarity is also the main driver of variation in cross-price elasticities, which are strongly skewed and approximately log-normally distributed.[13] Figure 1b illustrates the corresponding skewness of occupational similarities and a negative (and significant) relationship with measured task distance. In other words, the figure shows that the higher the distance in task content between two occupations, the lower the cross-price elasticity (i.e., lower 'substitutability' of these occupations). However, we note that task distance is based only on the set of tasks reported in survey responses and it explains at most a subset of occupational similarity, which in contrast contains all information implied by realised worker flows. This last point is underscored by the fact that task distance is essentially an ordinal variable whereas the skewness of occupational similarity and cross-elasticities has a natural quantitative interpretation. Accordingly, Spearman's rank coefficient provides a better fit in Figure 1b than standard linear correlation.

[12] Abraham & Kearney (2020) review the employment and wage effects of occupational licensing. Eckardt (2023) studies the effect of training specificity for the costs of switching occupations in Germany.

[13] A detailed analysis is in Appendix B.3, which depicts the log of cross-price elasticities against the normal distribution and decomposes its variation into occupational similarities versus other factors following Remark 2. We also plot the cross-price elasticities against occupational task distance directly.

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	Other attending on guests
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 $\rightarrow$ Other attending on guests
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: This table shows summary statistics of own- and cross-price elasticities in Panels A and B, respectively.

Figure 1: Elasticity Components: Comparison with External Metrics

(a) Aggregate Substitutability and Occupational Characteristics **(b)** Occupational Similarity and Task Distance

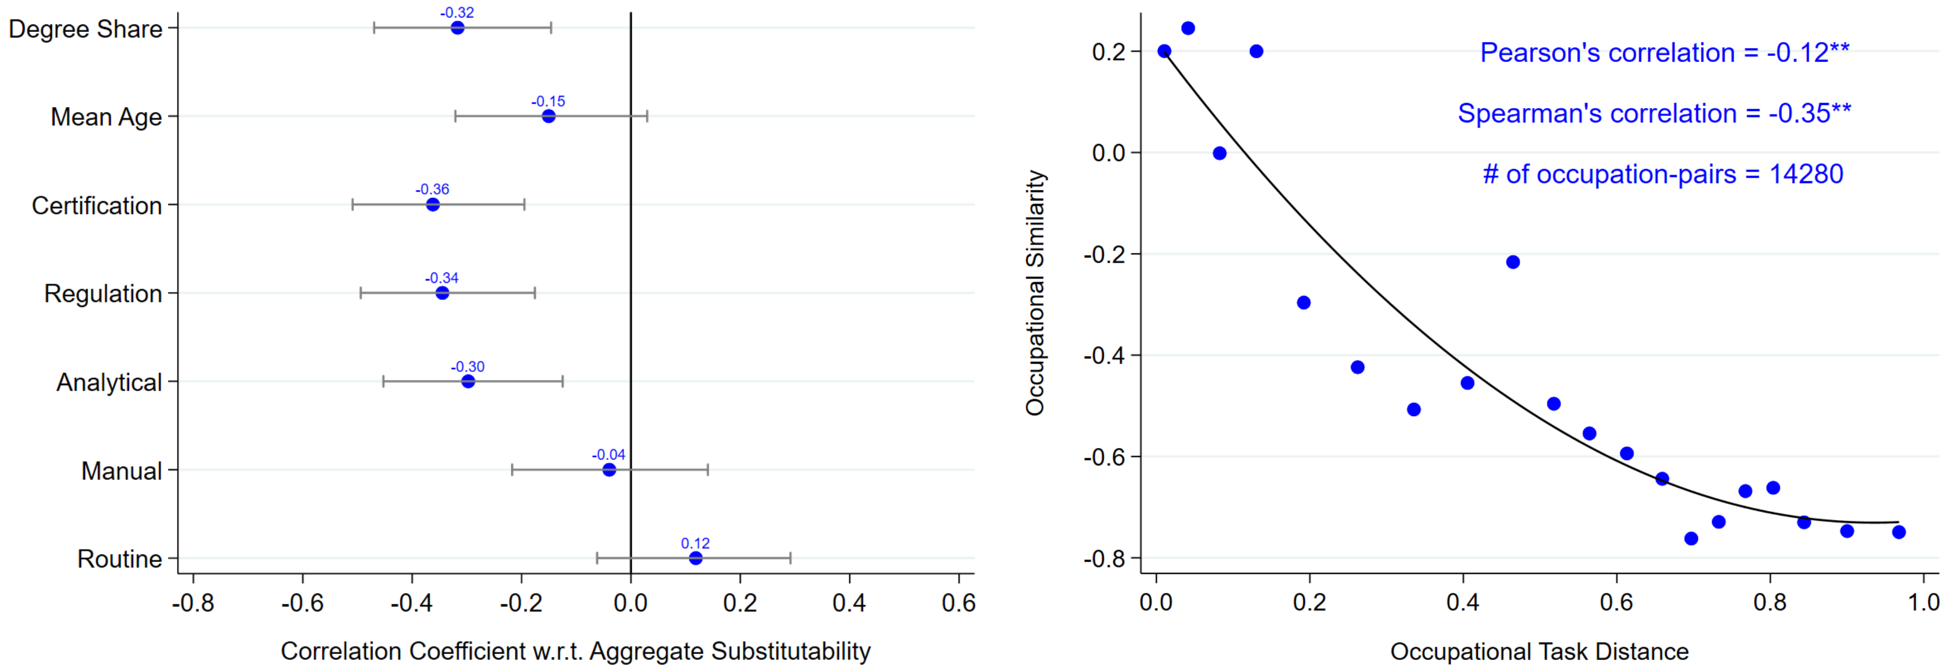

Notes: The figure shows the relationship of the elasticity components with respect to external metrics. Panel (a) reports how the aggregate substitutability component of the own-price elasticity, namely $\sum_{k\neq j} \tau_k Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$, relates to skill requirements captured by the share of university graduates and workers' average age, occupational certification and regulations (taken from Vicari, 2014), and occupational task content (analytical, manual, and routine). These are weighted by initial employment in each occupation. Panel (b) shows the relationship (with a quadratic fit) between the occupational similarity 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 B.1 does the same plots for d_{jj} and $-d_{jk}$ instead.

4 Estimates of Own- and Cross-Price Effects

This section presents our estimates of own- and cross-price effects. First, we study only bivariate relationships of occupations' changing prices with their employment. Then we add the effects of other occupations' price changes taking the respective own- and cross-price elasticities into account.

4.1 Heterogeneity of Own-Price Elasticities

Figure 2a plots occupations' changes in employment – annualised over the period 1985-2010 – against our measure of changes in occupational prices, based on stayers' wage growth. These wage growth rates clearly line up with their employment growth, consistent with earlier work (Cavaglia & Etheridge, 2020; Böhm et al., 2024). However, 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.

This paper's hypothesis is that a significant part of such heterogeneity is due to differences in labour supply curves across occupations. To investigate this empirically, we first consider individual price changes in isolation and reduce equation (9) as follows:[14]

$$\Delta e_j(\mathbf{p}) \quad \approx \quad \underbrace{\theta d_{jj} \Delta p_j}_{\substack{\text{own-price} \\ \text{effect}}} \tag{10}$$

The hypothesis is that the effect of occupations' own price changes on their employment should be governed by the heterogeneity in d_{jj}. In Figure 2b, we split occupations at the median of d_{jj} and draw two separate regression lines. The blue circles, including 'physicians and pharmacists', are the occupations ex-ante predicted to be relatively inelastic in terms of employment response with respect to changes in their own price, while the red circles, including 'assistants', are predicted to be relatively elastic.

Indeed, we find that the relationship between occupational employment and price changes is substantially flatter among the red than among the blue circles. That is, the employment response associated with a given price change is substantially stronger among the above-median d_{jj} (high predicted elasticity) than among the below-median d_{jj} (low elasticity) occupations. The differences on the regression slopes are not only strongly sig-

[14]That is, we focus on the partial derivative of own prices on employment ($\frac{\partial e_j(\mathbf{p})}{\partial p_j} = \theta d_{jj}$) from Remark 1.

nificant (p-value < 0.01), but also economically meaningful. 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 group. Appendix Figure C.1 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 exhibiting the steepest relation of employment vs prices, the highest d_{jj} quartile exhibiting the flattest relationship, and the middle quartiles ranked in between.

4.2 Full Own- and Cross-Effects Implementation. Estimating θ

The analysis above showed that even a simplified version of our model helps explain whether occupational changes are characterised by relatively larger shifts in employment or wages. The full model presented in Section 2 is however equally characterised by price effects that work *across all* occupations.

We take our model to data fully by developing eq. (9) as follows:

$$
\begin{aligned}
\Delta e_j &\approx \theta\left(d_{jj}\Delta p_j + \textstyle\sum_{k\neq j} d_{jk}\Delta p_k \right) \\
&= \underbrace{\theta \bar{d}_{diag}\Delta p_j}_{\text{fixed relationship of price with employment}} + \underbrace{\theta(d_{jj} - \bar{d}_{diag})\Delta p_j}_{\text{heterogeneity of own-price effect}} + \underbrace{\theta \sum_{k\neq j} d_{jk}\Delta p_k}_{\text{total cross-price effect}}
\end{aligned}
\tag{11}
$$

The top line of equation (11) repeats that displayed in eq. (9), and includes own- and cross-price effects. As in the previous section, it is instructive to further split the own-price effect into a fixed relationship that one would obtain when regressing employment onto price changes (Figure 2a) and the additional effect of the pure heterogeneity in elasticities d_{jj} (Figure 2b). This is done in the last line of eq. (11), where $\bar{d}_{diag}$ is the mean of matrix D's main diagonal elements, and the heterogeneity is captured by $d_{jj} - \bar{d}_{diag}$.

Our baseline empirical specification extends eq. (11) in two additional respects:

$$
\Delta e_j = \alpha + \theta_1 \bar{d}_{diag}\Delta p_j + \theta_2(d_{jj} - \bar{d}_{diag})\Delta p_j + \theta_3 \sum_{k\neq j} d_{jk}\Delta p_k + \varepsilon_j \tag{12}
$$

This regression replaces the pecuniary preference parameter common to all effects in eq. (11) by some generic coefficients, which allows us to test the theoretical restriction that $\theta_1 = \theta_2 = \theta_3 = \theta$. Second, 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. The approximation error from eq. (11) is represented by ε_j.

Figure 2: Stayers' Wage Growth and Employment Changes (1985–2010)

(a) All 120 Occupations

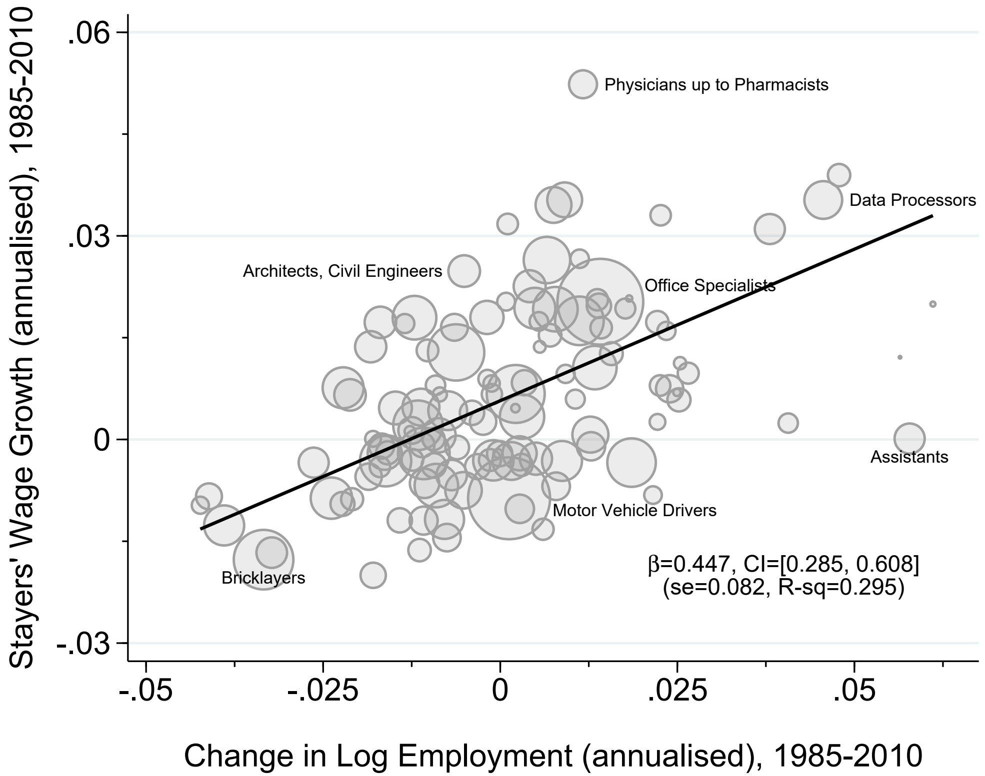

(b) By Own-Price Elasticity (Median Split Illustration)

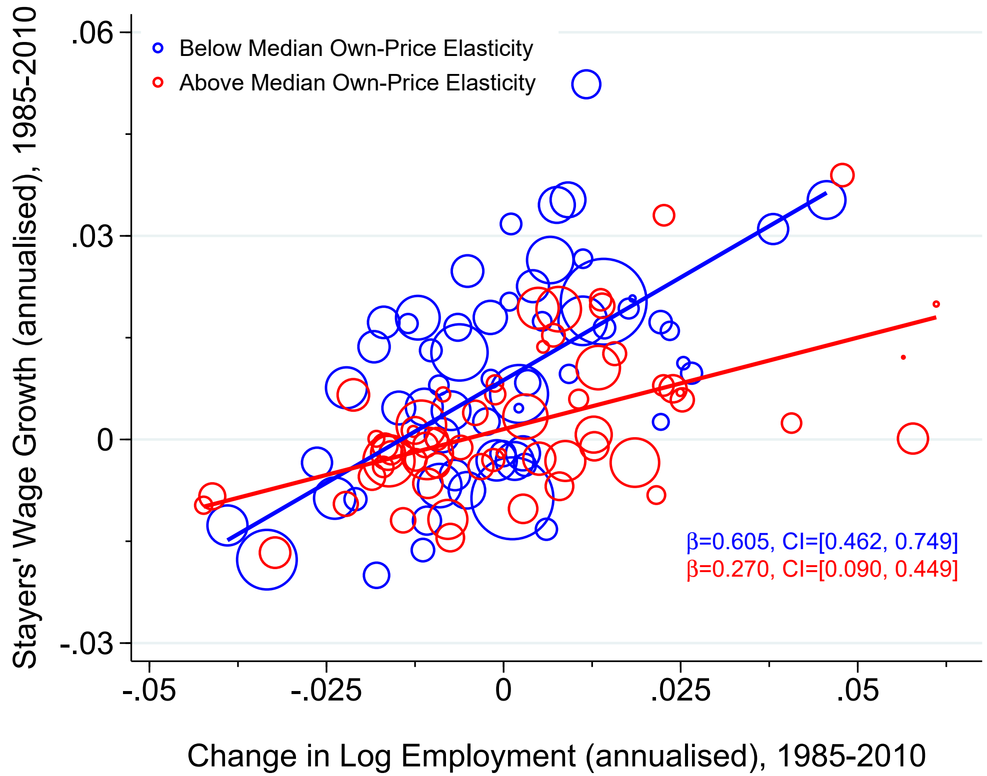

Notes: Scatter of occupations' change in log of total employment (x-axis) and stayers' wage growth (y-axis) between 1985–2010. Panel (a) shows the overall regression line for the 120 occupations. Panel (b) shows colour codes and linear regression lines by occupations below (blue, inelastic) and above (red, elastic) the median predicted elasticity of labour supply with respect to own price (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.

Table 2 reports the estimates from different versions of regression (12). Observations are weighted by occupations' initial employment so that coefficients represent effects as faced by the typical worker.[15] Column (1) shows the regression of Δe_j onto $\bar{d}_{diag}\Delta p_j$ only. As seen in Figure 2a and in prior work, this fixed relationship of employment with price changes results in a positive and significant slope parameter 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 the strong implications of Figure 2b.

Table 2: Full Model with 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	1.79	4.09		
		(0.30)	(0.31)	(0.89)	1.81	
heter. own effect:	$(d_{jj} - \bar{d}_{diag})\Delta p_j$		1.25	4.07	(0.32)	4.15
			(0.36)	(1.00)		(0.70)
total cross effect:	$\sum_{j \neq k} d_{jk}\Delta p_k$			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: Regressor in column (4) is $d_{jj}\Delta \ln p_j$. In column (5), the regressor is $\sum_k d_{jk}\Delta \ln p_k$, i.e., corresponding to the full model. 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.

Column (3) of Table 2 then adds the cross effects of price changes in other occupations that may be more or less substitutable. Consistent with theory, the coefficient on this is also positive and significant. This should be so, since $d_{jk} < 0$ 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 discussed above, a stronger implication of the theory is that coefficients θ_1–θ_3 should all capture the same pecuniary preference parameter. Although econometrically they are allowed to differ, estimated coefficients turn out almost identical across regressors. We examine the equality of coefficients more formally in columns (4) and (5). Consistent with $\theta_1 = \theta_2 = \theta_3$ being fulfilled, the estimates do not change much when we run the restricted models (10) and (11).[16]

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 highly cross-

[15] Unweighted regressions and additional specifications are reported in Table C.1 and C.2.

[16] To be precise, column (4) estimates $\Delta e_j = \alpha + \theta d_{jj}\Delta p_j + \varepsilon_j$ and (5) estimates $\Delta e_j = \alpha + \theta \sum_k d_{jk}\Delta p_k + \varepsilon_j$.

elastic occupations tended to experience similar price changes. That is, for $-d_{jk}$ large, Δp_j and Δp_k tended to move together such that $Cov(\Delta p_j, d_{jk}\Delta p_k) < 0$. Adding this up for all $k \neq j$, the own-price and total cross-price effects are negatively correlated, and including the latter raises the coefficient on the former in our estimation. Borusyak et al. (2022) find a related result in migration regressions across Brazilian region-industries. They highlight the omitted variables bias that results when not taking into account that shocks will often be correlated between workers' current and potentially substitutable employment options. In theirs as well as our case, estimated pecuniary parameters are indeed substantially larger in the full model which accounts for the fact that actual wage opportunities from moving across substitutable options are not that large.[17]

Simulations in Borusyak et al. (2022) show that migration responses can be underestimated by over half when not taking correlated shocks into account. We find that our pecuniary preference parameter more than doubles, to 4.15, once we include total cross-price effects. This number is broadly comparable to Cortes & Gallipoli (2018), who estimate θ using US wage data and obtain estimates in the range of 2 to 8.87.[18] 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 $\theta d_{jj} = 1.8$ as $\overline{d}_{diag} = 0.43$). The novelty of our approach lies in the heterogeneity around the average for own-price (from $0.07 \cdot 4.15 = 0.3$ to $0.80 \cdot 4.15 = 3.3$) as well as cross-price elasticities (from essentially 0 to 1.9). This stems from the worker flows and substitutabilities between occupations that we model explicitly.[19]

A final noteworthy feature of Table 2 is that the R-squared rises by 2 percentage points when including heterogeneous own effects and by another 8 percentage points when adding the total cross-price effects. This latter substantial increase, together with the change in the coefficients, indicates that cross-price effects are the critical components of the effective labour supply elasticities prevailing in the economy. We shall see this in further detail below. In particular, Section 6 will solve for the full economic model, including all shocks to demand and supply, to quantify the overall contribution of labour supply heterogeneity to the occupational changes.

[17] The variation in the right-hand side of the full model ($\sum_{k=1}^{N} d_{jk}\Delta p_k$) is half the variation in the right-hand side of the model with own-effects only ($d_{jj}\Delta p_j$).

[18] Cortes & Gallipoli (2018) set $\theta = 1$ in what corresponds to eq. (1) but estimate it via the dispersion of $\varepsilon_j(\omega)$, which is equivalent (see also footnote 5).

[19] Berger et al. (2022) and Jarosch et al. (2019) model heterogeneity based on employer size differences (granularity), a feature contained in but not dominating our elasticities.

5 Labour Demand, Equilibrium, and IV Estimation

To study labour market equilibrium, we add occupations' labour demand to the model. We characterise the resulting system of equations for prices and quantities and study its reaction to shocks. We then implement an instrumental variables estimation strategy that exploits relative demand shocks across occupations interacted with their predicted differential effects according to the heterogeneity in labour supplies.

5.1 Labour Demand and Equilibrium

Having primarily addressed the supply side thus far, we proceed to close the model by specifying an explicit theory of occupational labour demand. We provide a discussion of the main issues here leaving details to Appendix D.

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

$$Y = A\left(\sum_j \beta_j E_j^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \text{ s.t. } \sum \beta_j = 1 \tag{13}$$

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 1, which focused on supply, competitive behaviour results in labour demand elasticities of the form:

$$\frac{\partial e_j^d}{\partial p_k} = \sigma \begin{cases} -(1-\tau_j) & \text{if } j = k \\ \tau_k & \text{otherwise} \end{cases} \tag{14}$$

In equation (14), own-elasticities of labour demand are negative but attenuated by an occupation's size. The latter is equivalent to the price index terms in Remarks 2 and 3. The cross-price elasticities are positive and, after occupation size adjustment, constant.

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

$$e_j(\mathbf{b},\mathbf{s}) = e_j^s(\langle p(\mathbf{b},\mathbf{s})\rangle, \mathbf{s}) = e_j^d(\langle p(\mathbf{b},\mathbf{s})\rangle, \mathbf{b}) \tag{15}$$

where $\mathbf{b}$ is the vector of relative productivities (i.e., demand shifters $\left(\ln \frac{\beta_j}{1-\beta_j}\right)$), $\mathbf{s}$ is the vector of supply shifters, j indexes the occupation as before, and both supply (s) and demand (d) curves depend on the full set of prices.

Our focus is on the response of this system to shocks to the structural parameters, given by changes to $\left(\ln \frac{\beta_j}{1-\beta_j}\right)$ and $\mathbf{s}$. Appendix D shows that a linear approximation to the changes in prices and employment can be expressed as:

$$\Delta \mathbf{p} \approx V\Delta \mathbf{b} - \frac{1}{\sigma} V\Delta \mathbf{s} \tag{16}$$

and

$$\Delta \mathbf{e} \approx \theta D V \Delta \mathbf{b} + V\Delta \mathbf{s} \tag{17}$$

where $V = \left(\frac{\theta}{\sigma}D + I\right)^{-1}(I - W)$ and W is the matrix of stacked occupation sizes with j, kth element τ_k. Equations (16) and (17) 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 positive supply shocks increase employment but reduce prices.

Given that the matrix V plays a 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 V plays a parallel role here to that which matrix D plays in our analysis of the supply side of the market: it governs the dissipation of shocks across the economy. As in Section 3, we can summarise its effect most simply by examining its diagonal elements. Appendix Table D.1, which displays summary statistics for V, shows that the correlation of its diagonal with that of D is -0.96. Accordingly, for example, the diagonal elements of V tend to be *lower* for more elastic occupations. As implied by eq. (16), for these occupations, ceteris-paribus demand shocks induce relatively muted changes to prices.[20]

The right-hand side of Table D.1 summarises relevant features of the matrix product DV, which similarly has rank $N - 1$ with all non-negative eigenvalues. Most 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

[20] Appendix Table D.1 also provides summary statistics of the elements of V off the diagonal. 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 occupations. As indicated by eq. (16), this scarcity can then lead to an increase in prices in *both* occupations.

implied by eq. (17).[21] The parallel effects to those just discussed can be traced through shocks to supply.

Finally, combining eq. (16) and eq. (17), we obtain our basic regression equation

$$\Delta \mathbf{e} \approx \theta D \Delta \mathbf{p} + \Delta \mathbf{s} \tag{18}$$

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, $\mathbf{d_j} \Delta \mathbf{p}$.

5.2 Instrumental Variables Estimation

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. Equation (16) implies proportionality of the form

$$D\Delta \mathbf{p} \sim DV\mathbf{r} = D \left(\frac{\theta}{\sigma} D + I \right)^{-1} \check{\mathbf{r}} \tag{19}$$

where vector $\check{\mathbf{r}} \equiv (I - W)\,\mathbf{r}$ is the weighted-demeaned version of $\mathbf{r}$. Equation (19) represents an IV first-stage relationship for the relevant regressor, the product of elasticities with price changes. 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 our initial estimates of θ from Table 2, we calibrate $\frac{\theta}{\sigma} = 2.3$ as a benchmark.[22] As we shall see below, the resulting estimate of θ is consistent with this choice. Moreover, the robustness of our results to different values of $\frac{\theta}{\sigma}$ is shown in Appendix Table D.2. More immediately, we turn to the vector $\mathbf{r}$.

Our instrument for relative productivity shocks is based on initial task content. As discussed in Section 3, 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), several important papers have 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

[21] In terms of off-diagonal elements, the Table D.1 shows that for matrix DV these are all negative. Following through the example just given in footnote 20, 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. (17), this cross-effect always reduces employment.

[22] Relating relative wages to relative occupational inputs in US-CPS data, Burstein et al. (2019) estimate the elasticity of substitution between occupational inputs to be within the range $\sigma \in [1.81, 2.10]$.

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.[23] We thus approximate occupation j's (negative) demand shocks during 1985–2010 as

$$r_j = (routine_j + manual_j) - analytical_j$$

The idea is that occupations initially scoring high on routine and manual relative to analytical tasks will decline during the sample period, in terms of wages and employment, compared to occupations that score low on our measure r_j.

5.2.1 Estimation Without Cross-Price Effects

Following the exposition in Section 4, we illustrate the IV estimation by first implementing the model removing cross-price effects. As in Figure 2b we capture the heterogeneity of own-price effects by splitting estimation by the median value of d_{jj}. In this case, the instrument $DV\mathbf{r}$ within each sub-sample of 60 occupations reduces to a scalar multiple of the pure proxy vector $\mathbf{r}$.

The relationship between r_j and Δp_j is displayed in Figure 3a. 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 D.1 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 Figure 2b, 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.

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

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

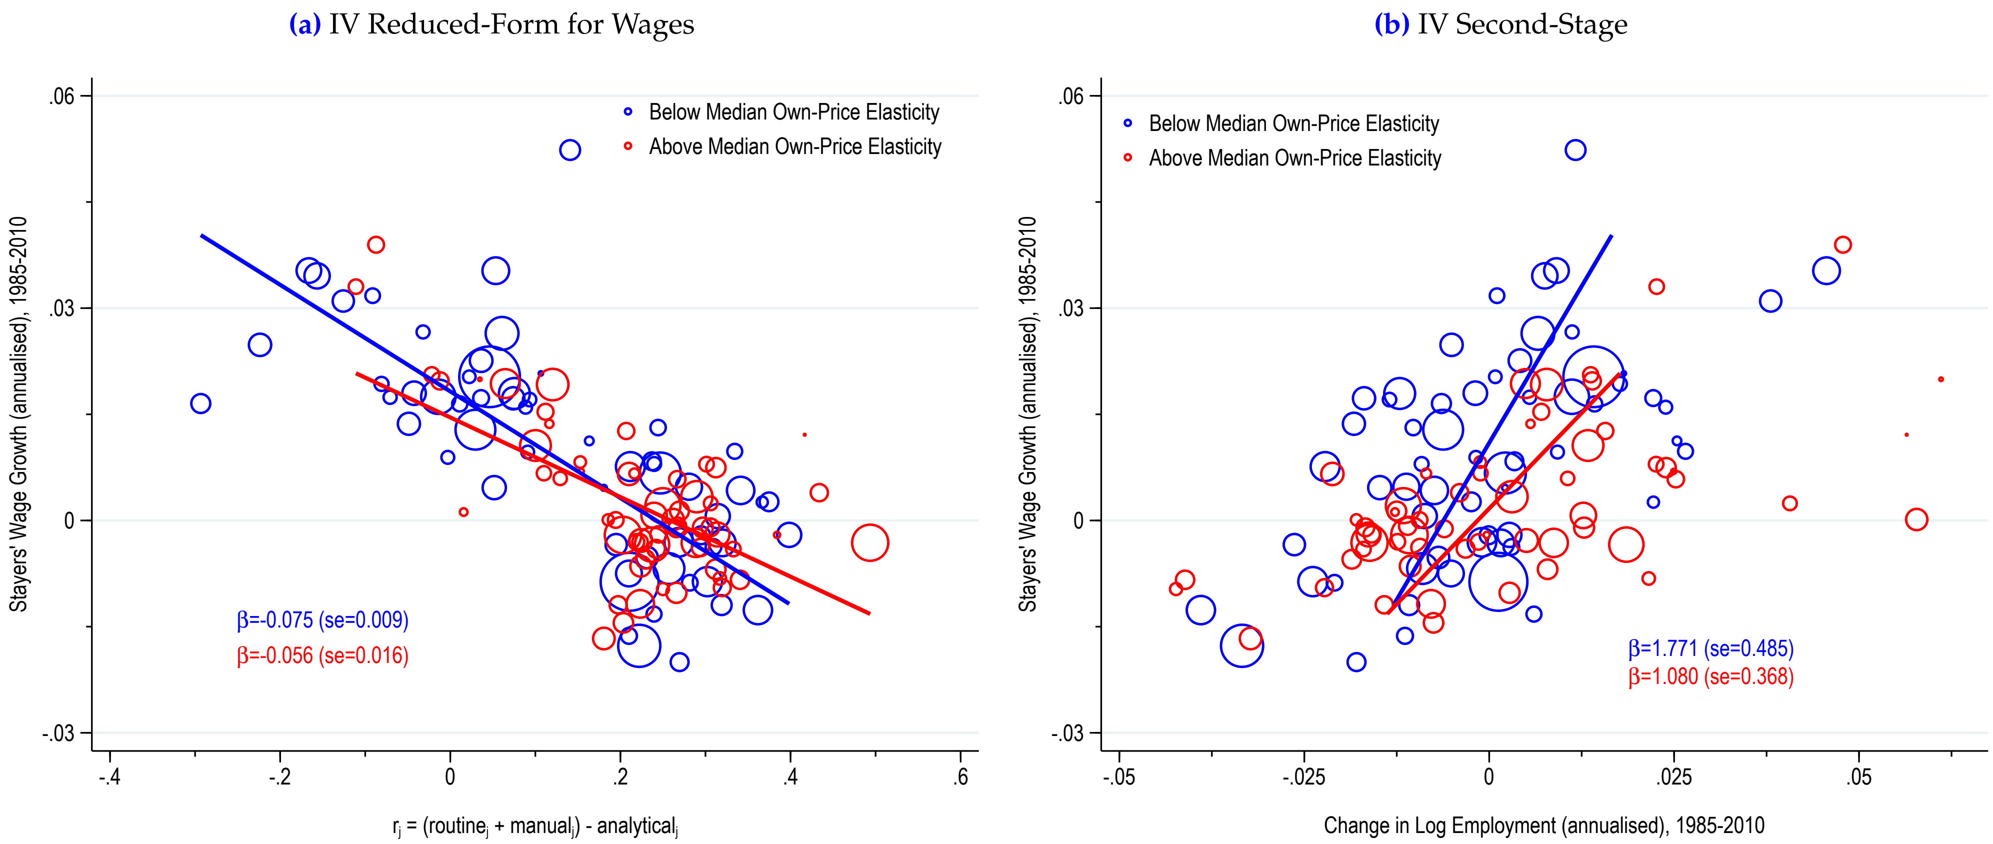

Notes: The left panel shows reduced-form regressions of occupations' price changes on their initial task contents r_j. The right panel shows second-stage IV-2SLS regressions of occupations' price changes on predicted employment changes using initial task contents as the instrument. Colour codes and linear regression lines are separated by occupations below (blue, inelastic) and above (red, elastic) the median predicted elasticity of labour supply with respect to own price (d_{jj}). β and *se* refer to the slope coefficient and standard error, respectively.

5.2.2 Full Model Estimation

Finally, expression (20) reports the results from the two-stage least squares regression, using (19) as first and (18) as second stage:

$$\begin{aligned} \Delta e_j &= \underset{(1.30)}{4.78}\ \mathbf{d}_j \Delta \mathbf{p} + constant + error_j \\ \mathbf{d}_j \Delta \mathbf{p} &= \underset{(0.0125)}{-0.046}\ \mathbf{d}_j V \mathbf{r} + constant + error_j \end{aligned} \tag{20}$$

We focus directly on the theoretical model with a single value of θ. 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 $\left(\frac{-0.046}{0.0125}\right)^2 = 13.5$. In the second stage, the estimated θ parameter is 4.78, or about 15% higher than the OLS estimate of 4.15 from Table 2, and again statistically significant.

Standard intuition implies that, if price changes are correlated with supply shocks, then OLS should be attenuated and biased downwards. We see this here, but the IV estimate in (20) is still relatively similar to that from the OLS. As we discuss through a formal analysis in Appendix D.4, in this case two relevant and opposing forces are at play: i) demand shocks were positively correlated with supply shocks, ii) the variance of supply shocks was relatively small. The second factor would, on its own, lead to only a small attenuation of OLS estimates. And, in fact, this attenuation is partly offset by the first factor.[24] We now turn to providing further insights into the full solution of the underlying supply and demand model.

6 Model-Based Decomposition and Counterfactuals

The previous section showed how to solve for the equilibrium of the full supply and demand model. We now use this to decompose the changes in employment and wages into contributions of different factors: shocks to occupational demand and supply as well as the heterogeneities in labour supply elasticities that we emphasise.

[24] It is important to note that in the reverse regression of prices on employment, this broad alignment of IV and OLS estimates need (and does) not hold. In fact, we find much steeper regression lines in Figure 3b than Figure 2b. More generally, a positive correlation of shocks does not preclude the IV estimation of the supply and demand model. What matters is that the subset of demand shocks used for identification (via the instrument) are uncorrelated with supply shocks.

6.1 Construction of Counterfactuals

We use equations (16)–(17) to express the changes of prices and employment in terms of parameters and exogenous shocks as follows:[25]

$$\Delta\mathbf{p} = \left(\frac{\theta}{\sigma}D + I\right)^{-1}\Delta\mathbf{b} - \frac{1}{\sigma}\left(\frac{\theta}{\sigma}D + I\right)^{-1}\Delta\mathbf{s} \tag{21}$$

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

The equilibrium solution treats equations (21) and (22) 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 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 similarities with other occupations are homogeneous. For example, employment in service industries may be responsive to price, but 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 firms literature (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.[26] We term this the model with 'heterogeneous own-price effects'.

Another counterfactual imposes completely homogenous labour supply elasticities. The main diagonal elements of matrix D_{hom} become an average $\bar{d}$ 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. (18), 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.[27] We term

[25] Equations (21) and (22) are obtained by inserting the solution for V into equations (16)–(17) and then using the fact that demand and supply shocks are weighted mean zero by construction (see Appendix D.2).

[26] We use size-weighted $d_{jk} = \frac{-\tau_k}{1-\tau_j} d_{jj}\ \forall k$, which is more theory-consistent, but $d_{jk} = \frac{-1}{N-1} d_{jj}$ fully homogeneous yields the same results as those shown below. See Appendix E.1 for details.

[27] Again, see Appendix E.1 for details.

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, setting $\Delta\mathbf{s} = \mathbf{0}$ in equations (21)–(22), which allows us to assess the variation in wages and employment that these shocks account for.

6.2 Results

Throughout this section we use our baseline parameter estimates: $\theta = 4.8$, $\sigma = 2.10$, and $\frac{\theta}{\sigma} = 2.3$.[28] We begin with a decomposition to uncover the drivers of overall employment changes. The first row in Table 3 shows that demand shocks in the fully homogeneous model (i.e., $\Delta\mathbf{s} = 0$ and D replaced by D_{hom}) explain 64% of the variance of employment changes.[29] This is consistent with the literature on job polarization (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.

The second row of Table 3 adds supply shocks, still under D_{hom}, to create a counterfactual price change according to eq. (21) in the homogeneous model. This explains 86% of employment changes in an R-squared sense, or roughly half of the remaining variance in $\Delta\mathbf{e}$. Similarly, adding heterogeneity of supply, and using full matrix D with the demand shocks in eq. (22), accounts for 85% of employment changes and again roughly half of the remaining variance.[30] Together, supply shocks and heterogeneity, by construction, explain the full variation in actual employment changes (last row of Table 3).[31] 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.

Figure 4 displays proper counterfactuals where, in keeping with the figures throughout the paper, we relate implied $\Delta\mathbf{p}$ from eq. (21) to implied $\Delta\mathbf{e}$ from eq. (22) in different scenarios. We start again with demand shocks in the fully homogeneous model. In this case, all occupational changes emanating from pure $\Delta\mathbf{b}$ run perfectly along a single supply curve (panel a). We can see from this plot that the explicitly labelled 'physicians and

[28] As discussed earlier, Appendix Table D.2 shows that θ estimates are largely insensitive to the exact $\frac{\theta}{\sigma}$. The results below are also very similar for the range of θ and σ values in Table D.2.

[29] Regarding the table labels, regressing $\Delta\mathbf{e}$ on raw $\Delta\mathbf{b}$ gives the same fit as regressing it on $\theta D_{hom} V_{hom} \Delta\mathbf{b} = constant + slope \cdot \Delta\mathbf{b}$, which is demand shocks' implied employment impact from eq. (22).

[30] A formal comparison to the generally lower levels of R-squared in the OLS regressions on price changes and heterogeneity (Table 2) is rather tedious. Intuitively, error terms related to supply shocks in eq. (22) are substantially less dispersed than in OLS estimation of eq. (18). In the latter, they are also negatively correlated with the regressor due to simultaneity, lowering the estimated contribution of prices.

[31] The two standalone contributions sum to $59.9\% + 57.9\% > 100\%$, which implies a -18% interaction effect. This is because eq. (22) is not purely additive.

Table 3: Decomposition of Overall Employment Changes

	(1) R-sq. between model & data	(2) Remainder explained
Base $\Delta\mathbf{e}$ with $\Delta\mathbf{b}$	0.641	
Adding supply shocks	0.856	59.9%
Adding supply heterogeneity	0.849	57.9%
Full model	1.000	

Notes: The table shows the results of the decomposition of overall employment changes. The first row starts by considering demand shocks in the fully homogeneous model. The second row adds supply shocks. The third row adds supply heterogeneity. The final row considers the full model.

pharmacists' as well as 'data processors' are among the occupations with the largest relative demand increases over time. 'Bricklayers' are among the occupations with the largest negative demand shocks.[32]

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 65%. The regression line moves *clockwise* and its slope reduces 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 effects, but retains homogeneity in cross-occupation elasticities (i.e., uses matrix D_{own} 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 regression-based analyses of Sections 4–5. A strong ex-

32. 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.

Figure 4: Counterfactual Changes of Prices and Employment

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

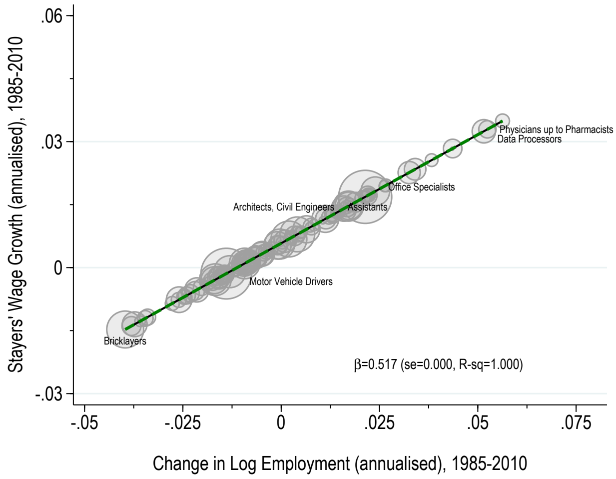

(b) ... with supply shocks

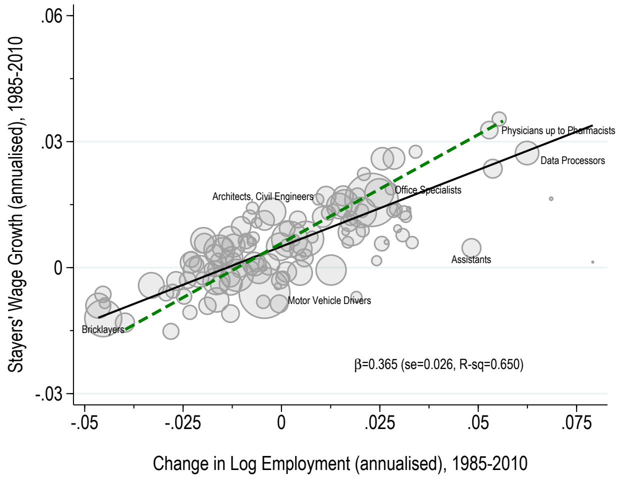

(c) ... with supply shocks; Het. own-price supply

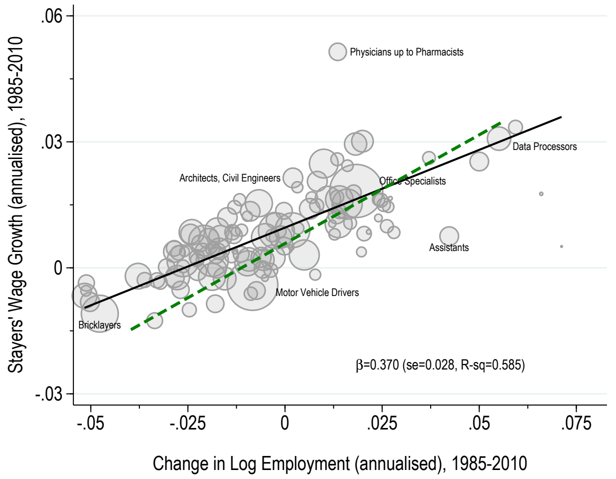

(d) Both shocks; Fully heterogeneous supply

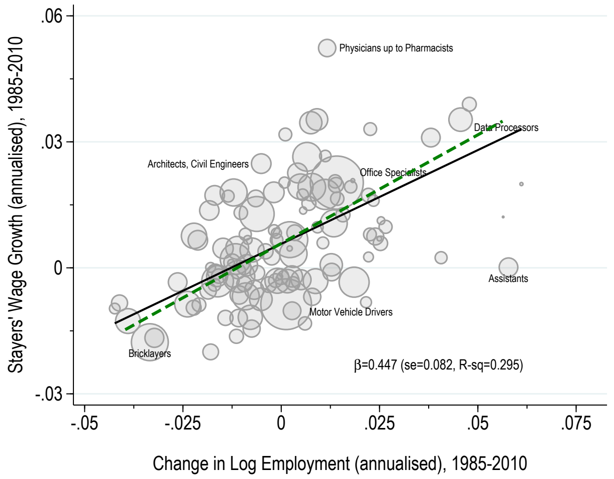

Notes: The figure shows price and employment changes for different manipulations of the elasticity matrix D and $\Delta\mathbf{s}$, as described in Section 6.1. In Panel 4a, both supply shocks and heterogeneity in D are switched off (i.e. $\Delta s = 0$ and D_{hom}), leaving only demand shocks. Panel 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 E.1. 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.

ception 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). In short, 'physicians and pharmacists' is the most notable occupation with an extreme own-elasticity (high or low) that also had an extreme demand shock.

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 'architects with civil engineers' and 'motor vehicle drivers' move away from one another. In addition, 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 cross-occupation elasticities to explain the data. As discussed previously, cross-occupation effects (also called 'total cross-price effects' in Section 4) make occupations less price elastic. In effect, realised cross-price elasticities captured by the full matrix D are lower than those captured by matrix D_{own} or the fully homogeneous model, since 'clusters' of occupations, which within them are relatively elastic, are equally shocked. Meanwhile, substitutabilities between the (differently shocked) clusters are relatively low.[33]

The impact of including the total cross-price effect is a key difference of the exposition in Figure 4 compared to earlier Table 3. It is seen even more starkly in Appendix Figure E.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 E.3–E.4. Among other things, these show that the lower effective labour supply elasticities have led to even larger between-occupation inequality than in a model without cross-price effects. For Figure 4, as was the case in the table, it is worth noting that, 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.

We finish this discussion by starting again at panel a and, from there, assessing the relative importance to dispersion in price and employment changes of supply shocks versus supply heterogeneity. Using changes in R-squared as a metric, we see that this is roughly equal. Panels a and b show that supply shocks cause a decline in the R-squared of 35 percentage points, while b and d show that supply heterogeneity accounts also for a decline of 35 ppt. 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 are important for explaining idiosyncratic outcomes of particular occupations.

[33] An interesting exception to this is 'assistants', which moves further Southeast in panel d compared to panel c. For this occupation, close substitute occupations saw strong relative demand declines, and its positive shock was therefore accentuated, making working as an assistant even more attractive. Our model thus provides a novel explanation for the expansion of this occupation over this period: the large increase in the number of assistants was not only due to an increase in the number of individuals suited for this type of work (positive supply shock) but also due to a strong increase in demand, not in absolute terms but *relative* to occupations requiring similar skills.

7 Extensions and Robustness

This section summarises findings from important extensions and robustness checks of the main results in the paper. The model is extended to non-employment transitions in Section 7.1. We then discuss estimates in five-year sub-periods and finally, in Section 7.3, with an alternative method of estimating changes in occupational prices.

7.1 Accounting for Non-Employment Transitions

A driver of heterogeneity in occupational growth that we have omitted so far is the extensive margin of employment. This may be particularly important if young workers' entry and old workers' exit from the labour market affects specific occupations' growth (in the case of US routine occupations this was prominently shown by Autor & Dorn, 2009). The secular decline of German unemployment from the mid-2000s may also be relevant in this respect.

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 in order to model-consistently control for switches to and from different non-employment states.[34]

We start by computing a new elasticity matrix that includes all transitions to and from non-employment states. Then consider equation (11) with $N+M$ occupations, M of which refer to non-employment sectors:

$$\Delta e_j \approx \theta \sum_{k=1}^{N+M} d_{jk} \Delta p_k = \theta \sum_{k=1}^{N} d_{jk} \Delta p_k + \sum_{m=N+1}^{N+M} (\theta \Delta p_m) d_{jm} \tag{23}$$

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 d_{jm} are control variables for the occupation j's predicted 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 of non-employment 'prices'. This product cannot be disentangled, as Δp_m is unobserved, but other than that the model is again identified.

Appendix F.1 shows the results from these estimations with $M = 3$ different non-

[34] As noted above, regressions so far included a constant that captures employment growth from sources other than direct occupational transitions (e.g., due to general growth of the working-age population). Now we allow for such contributions to vary by occupation.

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.[35] The R-squared is consistently higher in all specifications as more of the heterogeneity in employment growth can be explained when allowing for occupations' different elasticities with respect to non-employment states. However, the estimated role of own- and cross-price effects turn out similar to before (see Table F.1 and Table F.2). Results also do not substantively change when further separating part-time work and work with benefit receipt from out of labour force, or when merging the three states into one single non-employment sector (results not reported but available upon request).

7.2 Analysis in Five-Year Sub-Periods

In the main analysis, we have studied changes of occupational prices and employment over the period 1985–2010. We now 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 x 5 sub-periods) is used to estimate an extended version of (11):

$$\Delta e_{jt} = \alpha + \theta d_{jj}\Delta p_{jt} + \theta \sum_{k\neq j} d_{jk}\Delta p_{kt} + \delta_t \ (+\gamma_j) + \varepsilon_{jt} \tag{24}$$

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).[36] 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 or decelerations in the respective sub-period.

The results are shown in Appendix F.2. Graphically, Figure F.1 plots prices against employment growth for the pooled sample of 600 occupation-sub-periods (panel a) as well

[35] A limitation of the records from unemployment insurance is that we cannot observe the exact reasons for individuals entering or leaving the dataset (e.g. health shock, discouraged worker, emigration, self-employment, military service or becoming a civil servant). Outside the age range for labour market entry or retirement, these are all treated as out of the labour force for our purposes.

[36] Consistent with the autocorrelation of matrix D over time (see Appendix Table B.3), results are similar whether we use the baseline or the lagged matrix.

as separately for each sub-period (panel b), 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 is flatter than any slope for below-median own-price inelastic occupations. Linear OLS (Table F.3) and IV estimation (Table F.4) on the pooled data essentially reproduce the results obtained in Section 4. 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.[37] Overall, estimation in a series of shorter intervals shows that the role of occupational price elasticities persists, with some evidence that even acceleration or deceleration of price growth in different sub-periods is translated into employment growth according to these elasticities. These results are also robust to using alternative estimates of occupational prices, which we now turn to discuss.

7.3 Alternative Occupational Price Estimation

The results so far use the annual wage growth of workers who do not switch occupations (i.e., occupation stayers) as the main estimate of an occupation's changing log price or wage rate per efficiency unit of skill. This accounts flexibly for the selection into occupations based on observable and unobservable individual characteristics. Now 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 estimated by

$$\ln w_t(\omega) = \sum_j Z_{jt}(\omega)\varphi_{jt} + \sum_j Z_{jt}(\omega)X_t(\omega)\zeta_j + \sum_j Z_{jt}\kappa_j(\omega) + \mu_t(\omega) \tag{25}$$

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.[38] 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 our case, 1985-2010) in occupation's log prices (Δp_j).

[37] We can only do the OLS for this as the instrument does not vary by period.
[38] The bins are for ages 25–29, 30–34, 35–39, 40–44, 45–49, 50–54, and 55–59.

The results using prices à la Cortes (2016) are presented in Appendix F.3. The main figures of the paper are replicated using these alternative prices in Figure F.2. The main regression results (Table F.5–F.8), including those when accounting for non-employment transitions, turn out very similar. Our findings hence remain consistent and robust to this alternative price estimation.

8 Conclusion

Shifts in the demand for occupations have led to large changes in employment and wages. (Goos & Manning, 2007; Cavaglia & Etheridge, 2020). One important aspect that remains unexplored is the responsiveness of labour supply (i.e., the ability of the workforce to react) to the changing demand for jobs. In this paper, we study the role of the heterogeneity of occupational labour supply in explaining the variation of employment and wage growth between 1985 and 2010.

We propose a measure of occupation-specific labour supply elasticities, capturing the impact on employment of changes in the wage structure across occupations. These include wage changes in the occupation itself (own-price elasticities) and wage changes in other occupations (cross-price elasticities). We show how these price elasticities can be interpreted in terms of moments of the job flows and study how they relate to several occupational characteristics such as occupational certification or task content. We implement our framework in administrative panel data from Germany with long-running occupation information. Findings show that the heterogeneity in labour supply elasticities matters, and that accounting for wage changes in other occupations is important for explaining the evolution of the employment structure of the economy.

We close by discussing several avenues that this research opens up for further study. First, our framework allows us to conduct "what-if" scenarios to better understand the potential impact of specific changes. For example, one counterfactual might consider a scenario where occupations with lower employment elasticities were brought up to the median (own-price) elasticity. This could provide insights into the potential effects of policy interventions to enhance labour market flexibility and responsiveness. Second, while this paper studies the period 1985–2010, there exists an opportunity to extend the analysis to more recent years, specifically from 2012 onward, with a new classification of occupations.[39] Finally, exploring the ripple effects of labour supply elasticities on individuals' careers presents another future avenue. This entails understanding how shifts in demand and supply dynamics influence occupational mobility, career trajectories, and

[39] The introduction of the new occupation code KldB2010 in 2011 led to several serious problems (e.g. an increase in missing values for variables relevant in our framework) for the year 2011.

wage growth prospects for different segments of the workforce.

References

Abraham, K. G. & M. S. Kearney (2020). Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence. *Journal of Economic Literature 58*(3), 585–643. doi: 10.1257/jel.20191480.

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.

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

Adda, J. & Y. Fawaz (2020). The health toll of import competition. *The Economic Journal 130*(630), 1501–1540.

Altmann, S., A. Glenny, R. Mahlstedt, & A. Sebald (2022). The Direct and Indirect Effects of Online Job Search Advice.

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.

Autor, D., C. Chin, A. M. Salomons, & B. Seegmiller (2022). New Frontiers: The Origins and Content of New Work, 1940–2018. Technical report, National Bureau of Economic Research.

Autor, D. & D. Dorn (2009). This Job is 'Getting Old': Measuring Changes in Job Opportunities using Occupational Age Structure. *American Economic Review 99*(2), 45–51. doi: 10.1257/aer.99.2.45.

Autor, D., D. Dorn, G. Hanson, & K. Majlesi (2020). Importing political polarization? the electoral consequences of rising trade exposure. *American Economic Review 110*(10), 3139–3183.

Autor, D., L. Katz, & M. Kearney (2008). Trends in US Wage Inequality: Revising the Revisionists. *The Review of economics and statistics*. Available from: `http://www.mitpressjournals.org/doi/abs/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.

Autor, D. H. (2019). Work of the Past, Work of the Future. In *AEA Papers and Proceedings*, Volume 109, pp. 1–32. JSTOR. doi: 10.1257/pandp.20191110.

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.

Autor, D. H., D. Dorn, & G. H. Hanson (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review 103*(6), 2121–2168. doi: 10.1257/aer.103.6.2121.

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

Bárány, Z. L. & C. Siegel (2018). Job Polarization and Structural Change. *American Economic Journal: Macroeconomics 10*(1), 57–89.

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.

Belot, M., P. Kircher, & P. Muller (2022). Do the Long-term Unemployed Benefit from Automated Occupational Advice during Online Job Search?

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.

Bhalotra, S., M. Fernández Sierra, & F. Wang (2022). The Distribution of the Gender Wage Gap: An Equilibrium Model.

Boelmann, B., A. Raute, & U. Schonberg (2023). Wind of Change? Cultural Determinants of Maternal Labor Supply.

Böhm, M. J., H.-M. von Gaudecker, & F. Schran (2024). Occupation growth, skill prices, and wage inequality. *Journal of Labor Economics 42*(1), XX. doi: 10.1086/722084.

Borusyak, K., R. Dix-Carneiro, & B. Kovak (2022). Understanding migration responses to local shocks. *Available at SSRN 4086847*.

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.

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.

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.

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.

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 (2021). Technological Change, Firm Heterogeneity and Wage Inequality. Technical report, Working Paper.

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.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics 132*(4), 1593–1640.

Dustmann, C., J. Ludsteck, & U. Schönberg (2009). Revisiting the German Wage Structure. *Quarterly Journal of Economics 124*(2), 843–881. Available from: https://www.jstor.org/stable/40506245.

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

Eckardt, D. (2023). Training Specificity and Occupational Mobility: Evidence from German Apprenticeships. *mimeo*. Available from: https://www.dropbox.com/s/y07s57vfzba9067/Eckardt_paper.pdf?dl=0.

Felten, E. W., M. Raj, & R. Seamans (2018). A method to link advances in artificial intelligence to occupational abilities. In *AEA Papers and Proceedings*, Volume 108, pp. 54–57. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

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.

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.

Goos, M. & A. Manning (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *Review of Economics and Statistics 89(1)*, 118–133. Available from: https://www.jstor.org/stable/40043079.

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.

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.

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

Jarosch, G., J. S. Nimczik, & I. Sorkin (2019). Granular search, market structure, and wages. Technical report, National Bureau of Economic Research.

Keller, W. & H. Utar (2022). Globalization, gender, and the family. *The Review of Economic Studies 89*(6), 3381–3409.

Kleiner, M. M. & M. Xu (2024). Occupational licensing and labor market fluidity. *Journal of Labor Economics XX*(XX), XX. Available from: https://www.mingecon.com/uploads/8/6/8/6/86866790/kleinerxu_2022_sep.pdf.

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.

Mazzolari, F. & G. Ragusa (2013). Spillovers from high-skill consumption to low-skill labor markets. *Review of Economics and Statistics 95*(1), 74–86.

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.

Mishel, L., J. Schmitt, & H. Shierholz (2013). Assessing the job polarization explanation of growing wage inequality. *Economic Policy Institute, Working Paper 11*, 2013.

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.

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.

Vicari, B. (2014). Degree of Standardised Certification of Occupations. Technical report, FDZ Methodenreport.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Technical report.

Appendix for:

Heterogeneity of Occupational Labour Supply Elasticities

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

Draft version: February 2024

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, and finally develop the model decomposition for counterfactual analysis in detail.

A.1 Derivation of Formal Results: Remarks 1–3

We start by formally deriving Remarks 1–3.

A.1.1 Remark 1 (Elasticities and Job Flows)

We define the inverse 'choice index' as $\lambda_i(\mathbf{p}) = \frac{1}{\sum_{k=1}^{N} \exp(\theta p_k + a_{ik})}$, 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 follows:

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

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:

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

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

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

By combining these results, we derive the following expression:

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

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

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

This is equation (4) in Section 2. It shows 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.1.2 Remark 2 (Individual Cross-Price Elasticities)

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:

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

$$= -\frac{1}{\tau_k} d_{jk} - 1$$

Rearranging gives equation (6).

A.1.3 Remark 3 (Individual Own-Price Elasticities)

Turning to the on-diagonal elements of the elasticity matrix D. These are:

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

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

$$\begin{aligned} d_{jj} &= \frac{1}{\tau_j} \sum_i \tau_i \pi_{ij} - \frac{1}{\tau_j} \sum_i \tau_i \pi_{ij}^2 \\ &= 1 - \frac{1}{\tau_j} \left(Var_\tau \left(\pi_{.,j}\right) + \left(\mathbb{E}_\tau \pi_{.j}\right)^2 \right) \\ &= 1 - \frac{1}{\tau_j} \left(Var_\tau \left(\pi_{.,j}\right) + \tau_j^2 \right) \\ &= 1 - \tau_j \left(1 + \frac{1}{\tau_j^2} Var_\tau \left(\pi_{.,j}\right) \right) \\ &= 1 - \tau_j \left(1 + Var_\tau \left(\tilde{\pi}_{.,j}\right)\right) \end{aligned} \tag{26}$$

Rearranging gives expression (7).

A.1.4 Further Discussion on Remarks 1–3

We turn now to justifying our choices of normalisations. We first consider $Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) \equiv \sum_i \tau_i \left(\tilde{\pi}_{ij} - \mathbb{E}_\tau \tilde{\pi}_{.,j}\right) \left(\tilde{\pi}_{ik} - \mathbb{E}_\tau \tilde{\pi}_{.,k}\right)$. 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 compute this term for occupational classifications at various levels of coarseness (i.e., 4 main occupation groups, 10 occupation groups, 120 occupations), see Table A.1 below.

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

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

$$= \sum_{k \neq j} \tau_k + \sum_{k \neq j} \tau_k Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$$
$$= 1 - \tau_j + \sum_{k \neq j} \tau_k Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) \tag{27}$$

Equating equations (26) and (27) we see that

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

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.[40] 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, we compute this term for occupational classifications at various levels of coarseness. In particular, Table A.1 reports the median for the covariance, and two measures of the variance. It shows this for three levels of aggregation: 4 main occupation groups as described below in the data appendix section, 10 occupation groups corresponding to one-digit occupation categories, and the 120 occupations considered in the analysis (see Table B.5 for the full list).

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: 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 detailed occupations and their aggregations see Section B.1.

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

A.2 Remark 4: Vector of Price Changes

This section develops a further result on the aggregation of Remarks 2–3 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 in terms of distributions of worker flows.

Remark 4 (Vector of Price Changes) *Matrix D can be expressed as follows*

$$D = I - W - W \otimes C \tag{28}$$

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

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

$$= \theta \left(\underbrace{\Delta p_j - \Delta \mathbb{E}_\tau p}_{\substack{\text{real price} \\ \text{change}}} + \underbrace{Cov_\tau\left(c_{.,j}, \Delta p_j - \Delta p_.\right)}_{\substack{\text{occupational} \\ \text{substitutability}}} \right) \tag{29}$$

$$= \theta \left(\underbrace{\left(1 - \tau_j - \tau_j c_{jj}\right) \Delta p_j}_{\substack{\text{own-price} \\ \text{effect}}} + \underbrace{\sum_{k \neq j} \left(-\tau_k - \tau_k c_{jk}\right) \Delta p_{k.}}_{\substack{\text{total cross-} \\ \text{-price effect}}} \right) \tag{30}$$

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

Remark 4 complements the interpretations contained in Remarks 2–3. In the formulation in equation (29), the effect of a vector of price changes on a given sector 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 (30) then builds on this formulation by relating it back to equation (9), which forms the basis of our empirical application. Equation (30) 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

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

follows directly from Remarks 2–3. 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 Data Appendix

This section provides a more detailed presentation of the main data and supplementary descriptive statistics. We first discuss the SIAB data and outline the procedures for sample selection and wage imputation. We then present additional descriptive statistics on the main variables.

B.1 The SIAB Data

We use the Sample of Integrated Labour Market Biographies (*Stichprobe der Integrierten Arbeitsmarktbiographien* – SIAB) for our analyses.[41] 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 (up to 2010), the corresponding nominal daily wage, and socio-demographic variables such as age, gender, or level of education. Data is 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 (this follows from Böhm et al., 2024).

B.1.1 Sample Selection and Variable Description

To work with a homogeneous sample throughout, the main sample is restricted to West German full-time male workers aged 25–59. Since the level and structure of wages differ substantially between East and West Germany, we drop from our sample all workers who were ever employed in East Germany. Our focus on full-time jobs is driven by the absence of data on hours worked. Excluding younger workers, we ensure the vast majority of our

[41] Access to the data is subject to signing a contract with the Research Data Center of the German Federal Employment Agency. See Frodermann et al. (2021) for an up-to-date documentation of the data.

sample will have concluded their formal education by the time they enter the sample. Besides, we stop relatively early (at 59) because early retirement programs were common in Germany, particularly in the late 1970s and the 1980s.

We further 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 prior to 1999. After the first reform in 1999, the tax-free wage threshold was fixed during the period 1999 to 2003 at 325 euro per month. In 2003, the range of exempted earnings was expanded up to 400 euro, which was effective until 2012. The minimum threshold for mini-jobbers increased in 2013 from 400 to 450 euro 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).[42] Finally, workers without information on their occupation or wages are dropped from the analysis.

Occupations. 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 (from the detailed KldB 1988 classification system), available during the long time period of 1975–2010, and listed in Table B.5. After 2010, SIAB uses a new classification system, which results in a relatively sharp break of the occupation codes. 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 workers, operators, and craftsmen (Prod-Oper-Crafts), and (4) workers in services and care occupations (Serv-Care).

Wages. The available wage variable is the employee's gross daily nominal wage in euro. 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 East and West Germany, where the decisive factor is the location of the establishment. Second, the income components being subject to social security tax were extended

[42] Workers who are German at some point but foreign at another are not dropped from the sample.

in 1984. Prior to that, one-time payments such as bonuses were not included in the daily wage benefit measure. We further 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.

B.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 main sample (25–59 years old, 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).[43] 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.[44] 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:

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

[43]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).

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

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 further details.[45] 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.

B.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 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. Instead of doing that, 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.

[45] 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. (2021) and Böhm et al. (2024).

B.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.

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:[46]

$$\text{AngSep}_{jk} = \frac{\sum_{a=1}^{A} (x_{aj} \cdot x_{ak})}{\left[\sum_{a=1}^{A} (x_{aj})^2 \cdot \sum_{a=1}^{A} (x_{ak})^2 \right]^{\frac{1}{2}}} \tag{31}$$

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. 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.

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

To obtain measures of occupation's certification requirements and degree of regulation, we use the indicators for standardised certificates and 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. The degree of standardised certification of an occupation is 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. 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. Occupations that are not initially categorised as standardised, if they are regulated, they are also considered to have standardised certification. These 0-1 values are finally aggregated at the 3-digit occupational classification (i.e. the 120 occupations used in our analysis), weighted by the number of individuals employed in each occupation. Intuitively, occupational certification indicates whether access to exercising a professional activity is linked to a standardised training certificate. 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. These indicators are constructed as a metric value between 0 and 1, with the indicator increasing in the degree of certification and regulation.

B.3 Descriptive Statistics

This section presents descriptive statistics to complement the analysis in Section 3.

Table B.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.[47] 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).

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

The middle panel of Table B.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 B.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
Stayers' Wage Growth	0.586	0.516	1.354	−0.959	0.408	2.168	120
Prices á 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 (%)	0.135	0.117	0.232	0.006	0.018	0.463	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 $\mathbf{r}$	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 $\mathbf{r}$ used in Section 5. The weighted mean is weighted by each occupation's employment share in 1985.

Table B.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 B.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 (á 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 (á 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 (á 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 (á 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 (á 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 and occupational price changes for different 5-year periods. The last column refers to the autocorrelation between that period and the earlier 5-year period, e.g. the autocorrelation of employment changes between 1990-1995 relative to employment changes in 1985-1990.

Table B.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 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 B.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 B.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) are 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 in Section 7 when estimating our model pooled in these five-year sub-periods.

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

	Elasticity Matrix D		**Transition Probability Matrix Π**	
	Own-Price Elasticity (d_{jj})	Cross-Price Elasticity ($-d_{jk} \times 100$)	Diagonal Elements (Π_{jj})	Off-Diagonal Elements ($\Pi_{jk} \times 100$)
Mean	0.434	0.364	0.746	0.214
Std. Dev.	0.128	0.939	0.090	0.660
Variance	0.016	0.882	0.008	0.436
Skewness	0.177	14.672	−0.722	17.449
Kurtosis	3.634	493.494	4.393	585.670
p10	0.294	0.007	0.627	0.000
p50	0.430	0.111	0.754	0.046
p90	0.604	0.867	0.839	0.516
p99	0.796	4.021	0.931	2.585
Average autocorr. 5-year	0.876	0.881	0.868	0.806
Autocorrelation 25-year	0.761	0.768	0.761	0.660
Number of Observations	120	14,280	120	14,280

Notes: The table presents summary statistics for the elasticity matrix D (Remark 1) and the transition probability matrix. The average (5-year period) autocorrelation is computed by averaging autocorrelations of reported variables between 1985–1990 and 1980–1985, 1990–1995 and 1985–1990, 1995–2000 and 1990–1995, and so on. The 25-year period autocorrelation refers to the autocorrelation between the later period 2005–2010 and the earlier period 1980–1985.

We show how own-price elasticities d_{jj} relate to several occupational characteristics in Figure B.1a. 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 B.1 plots cross-price elasticities against occupational task distance. Table B.5 below offers the full list of the 120 occupations ranked by their respective own-price elasticities, together with their employment sizes in 1985 and 2010.

Figure B.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

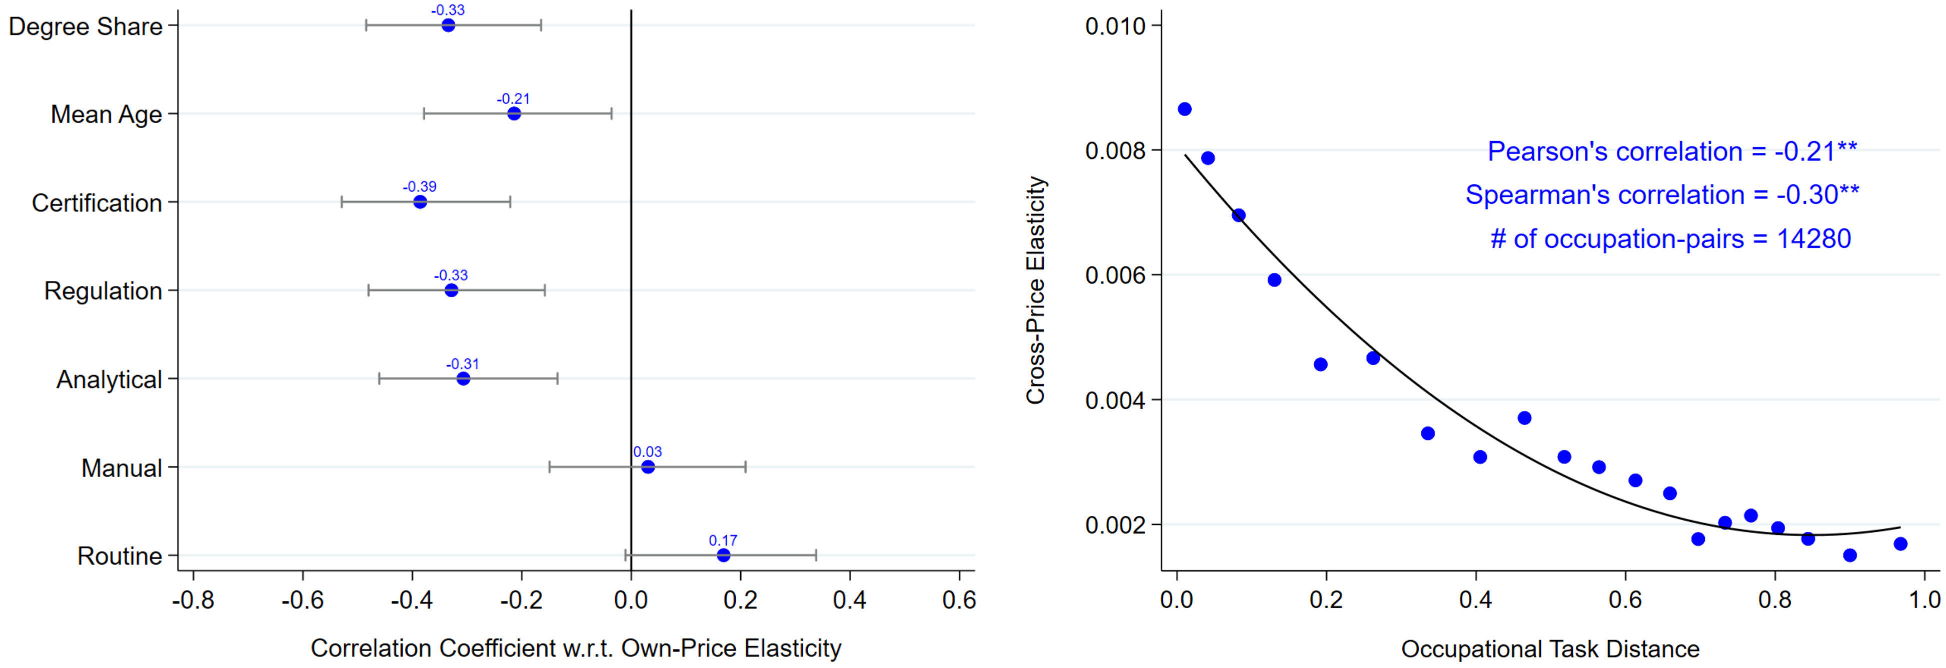

Notes: This table shows the relationship of the own- and cross-price elasticities with respect to external metrics. Panel (a) reports how the own-price elasticity, namely d_{jj}, relate to skill requirements captured by the share of university graduates and workers' average age, occupational certification and regulations (taken from Vicari (2014), and occupational task content (analytical, interactive, manual, and routine). These are weighted by initial employment in each occupation. Panel (b) shows the relationship (with a quadratic fit) between the cross-price elasticity, namely $-d_{jk}$, and occupational task distance measured as in Cortes & Gallipoli (2018).

Finally, we decompose the variation of labour supply elasticities. Column two of Table B.3 showed that occupational cross-price elasticities are strongly skewed and with high kurtosis. As Figure B.2 shows, they are distributed approximately *log*-normally. Accordingly, we decompose the log of the cross-price elasticities using the expression in Remark 2 as follows:

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

Here, the variance of log differences in cross-price elasticities can be decomposed into variances of log differences in sector sizes and occupational similarities (plus one, to make them all positive). Table B.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 $\ln\left(Cov_\tau(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}) + 1\right)$, while the dispersion in log occupational sizes contributes less than 30%. Although not shown in Table B.4, but can be easily inferred, the covariance of log occupational size with the similarity 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 similarities.

Figure B.2: Kernel Density of Log Cross-Price Elasticity

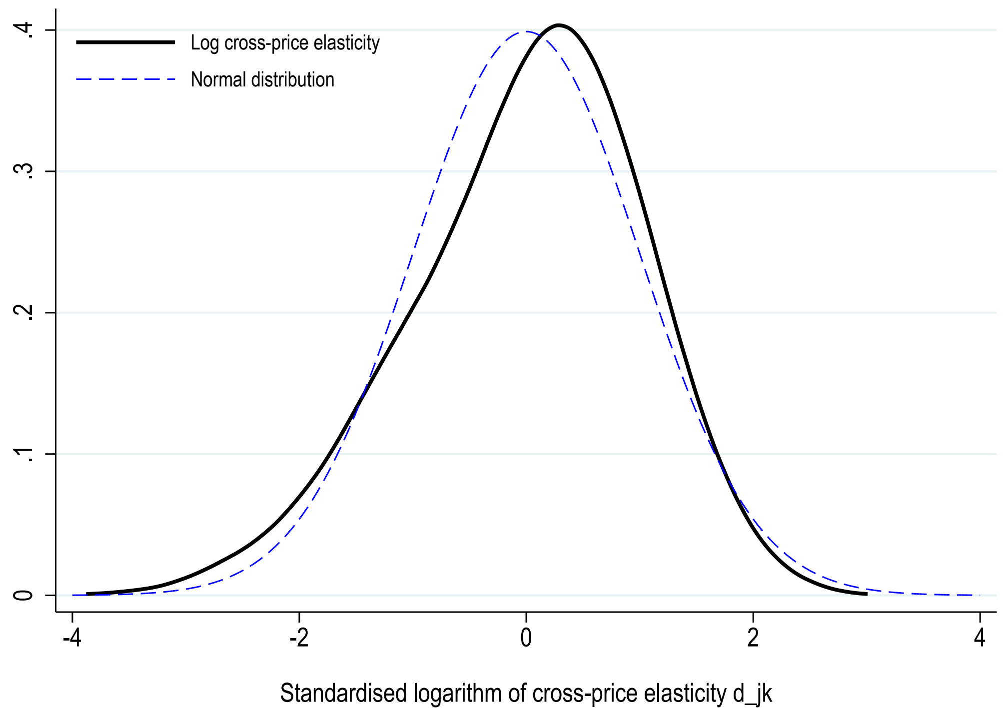

Notes: The figure shows the kernel density of the (standardised to have mean zero and standard deviation one) log of cross-price elasticity, superimposing a normal distribution. Bandwidth is 0.3.

Own-elasticities are distributed approximately normally in *levels* (see reported skewness and kurtosis in Table B.3 above). Although we do not explore the reason for this feature rigorously here, it seems to be because own-price elasticities comprise the sum of many apparently independently-distributed terms, as Remark 3 indicates. In line with this feature, and with the first expression of Remark 3, 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 Table B.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.

Table B.4: Summary Statistics. Elasticity Components

	Mean	Std. Dev.	Variance	Min	p10	p50	p90	Max	Skewness
Own-Price Elasticity (d_{jj})									
Aggregate Substitutability	−0.558	0.126	0.0161	−0.918	−0.692	−0.565	−0.389	−0.202	0.131
Stationary Employment Size	0.008	0.012	0.0001	0.001	0.002	0.004	0.017	0.090	4.360
Cross-Price Elasticity ($-d_{jk}$)									
Occupational Similarity	−0.429	2.418	5.848	−0.999	−0.979	−0.779	0.203	129.849	36.912
Log of components:									
Cross-Price Elasticity		1.917	3.676						−0.479
Stationary Employment Size		1.072	1.150						−0.488
Occupational Similarity + 1		1.582	2.504						−0.259

Notes: The table presents summary statistics for the elasticity components of the own-price elasticity (as discussed in Remark 3) and cross-price elasticity (as discussed in relation to Remark 2). The number of observations is 120 for own-price elasticity and its components, while 14280 for cross-price elasticity and its components.

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

Occupations (based on German KlDB 1988 Classification)	Own-Price Elasticity		% Share of Employment	
	d_{jj}	d_{jj}^{NE}	1985	2010
Physicians up to Pharmacists	0.07	0.27	0.65	0.81
Bank specialists up to building society specialists	0.13	0.18	1.79	1.98
Nurses, midwives	0.16	0.24	0.37	0.67
Dental technicians up to doll makers, model makers, taxidermists	0.18	0.31	0.32	0.24
Non-medical practitioners up to masseurs, physiotherapists and related occupations	0.19	0.27	0.13	0.22
Journalists up to librarians, archivists, museum specialists	0.20	0.28	0.28	0.35
Hairdressers up to other body care occupations	0.23	0.37	0.06	0.06
Architects, civil engineers	0.23	0.30	0.83	0.69
Soldiers, border guards, police officers up to judicial enforcers	0.27	0.36	0.38	0.51
Musicians up to scenery/sign painters	0.28	0.39	0.29	0.31
Foremen, master mechanics	0.29	0.37	1.39	0.75
Health insurance specialists (not social security) up to life, property insurance specialists	0.29	0.37	0.85	0.89
Chemical laboratory assistants up to photo laboratory assistants	0.30	0.33	0.26	0.25
Doormen, caretakers up to domestic and non-domestic servants	0.30	0.40	0.97	0.97
Type setters, compositors up to printers (flat, gravure)	0.31	0.35	0.75	0.36
Gardeners, garden workers up to forest workers, forest cultivators	0.31	0.40	1.18	1.15
Social workers, care workers up to religious care helpers	0.31	0.39	0.42	0.68
Carpenters	0.32	0.39	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	0.38	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	0.39	0.35	0.34
Bricklayers up to concrete workers	0.34	0.43	2.95	1.20
Home wardens, social work teachers	0.34	0.41	0.28	0.46
Music teachers, n.e.c up to other teachers	0.34	0.41	0.27	0.32
Electrical engineers	0.34	0.37	1.00	1.18
Entrepreneurs, managing directors, divisional managers	0.34	0.43	2.63	2.11
Data processing specialists	0.35	0.38	1.18	3.46
Members of Parliament, Ministers, elected officials up to association leaders, officials	0.36	0.46	0.33	0.48
Measurement technicians up to remnining manufacturing technicians	0.36	0.41	0.81	0.48
Painters, lacquerers (construction)	0.36	0.43	1.11	0.91
Office specialists	0.36	0.43	6.10	8.15
Dietary assistants, pharmaceutical assistants up to medical laboratory assistants	0.36	0.38	0.03	0.05
Chemical plant operatives	0.36	0.43	1.25	0.97
Navigating ships officers up to air transport occupations	0.37	0.45	0.39	0.28
Paper, cellulose makers up to other paper products makers	0.37	0.44	0.53	0.50
Artistic and audio, video occupations up to performers, professional sportsmen, auxiliary artistic occupations	0.37	0.44	0.27	0.25
Motor vehicle drivers	0.38	0.44	5.57	5.39
Toolmakers up to precious metal smiths	0.38	0.43	1.13	0.80
Cost accountants, valuers up to accountants	0.38	0.45	0.82	0.51
Railway engine drivers up to street attendants	0.39	0.47	0.77	0.61
Bakery goods makers up to confectioners (pastry)	0.39	0.46	0.41	0.41
Other technicians	0.39	0.45	1.96	2.43
Commercial agents, travellers up to mobile traders	0.39	0.45	1.58	1.10
Miners up to shaped brick/concrete block makers	0.40	0.47	1.33	0.47
Roofers	0.40	0.49	0.37	0.40
Survey engineers up to other engineers	0.40	0.46	0.75	1.82
Plumbers	0.40	0.46	1.35	1.23
Technical draughtspersons	0.40	0.45	0.60	0.48

Table B.5—continued

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

Table B.5—continued

Occupations (based on German KlDB 1988 Classification)	Own-Price Elasticity d_{jj}	d_{jj}^{NE}	% Share of Employment 1985	2010
Stowers, furniture packers up to stores/transport workers	0.56	0.64	1.95	2.91
Stenographers, shorthand-typists, typists up to data typists	0.56	0.61	0.11	0.12
Other mechanics up to watch-, clockmakers	0.56	0.59	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	0.64	0.86	0.92
Locksmiths, not specified up to sheet metal, plastics fitters	0.59	0.63	1.32	1.54
Salespersons	0.60	0.65	1.57	2.06
Laundry workers, pressers up to textile cleaners, dyers, and dry cleaners	0.60	0.66	0.06	0.06
Building labourer, general up to other building labourers, building assistants	0.61	0.70	1.26	0.97
Electrical appliance, electrical parts assemblers	0.62	0.66	0.22	0.20
Other assemblers	0.63	0.68	0.31	0.81
Household cleaners up to glass, building cleaners	0.63	0.73	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	0.71	0.35	0.58
Metal workers (no further specification)	0.67	0.71	1.07	1.38
Assistants (no further specification)	0.71	0.75	0.75	3.00
Other attending on guests	0.74	0.80	0.21	0.12
Medical receptionists	0.80	0.83	0.01	0.02
Nursery teachers, child nurses	0.80	0.79	0.02	0.09

Notes: The table provides diagonal elements of the elasticity matrix D, not accounting (column (1)) and accounting for non-employment (columns (2)). Columns 3-4 report the occupation's percentage share of employment in 1985 and 2010, respectively.

C Empirical Results on the Labour Supply Model

This section provides tables and figures to complement the main estimation results.

Figure 2b in the main text splits occupation at the median of d_{jj} and draw two separate regression lines. Figure C.1 below alternatively splits occupation 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 C.1: Stayers' Wage Growth and Employment by d_{jj} Quartiles

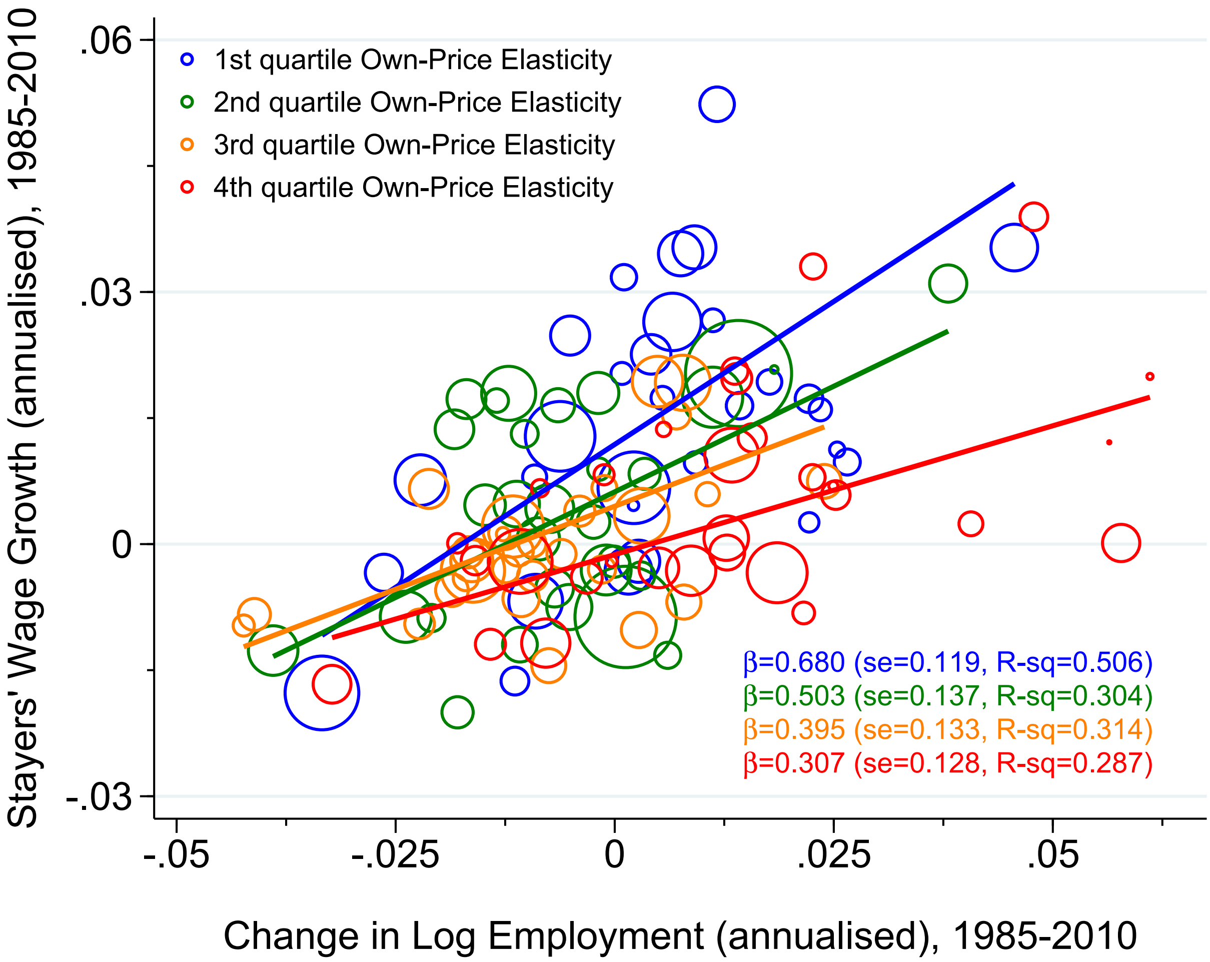

Notes: Scatter of occupations' change in log of total employment (x-axis) and task prices (y-axis, measured by average occupation-stayer wage growth) during 1985–2010. Panel (a) shows the overall regression line. Panel (b) shows colour codes by occupations in the lowest (blue), second (green), third (orange), and highest (red) quartile of the predicted elasticity of labour supply with respect to own price d_{jj}.

Table C.1 considers the case in which own-price effects are not split into a fixed relationship and the additional effect of the heterogeneity in elasticities d_{jj}. That is, it directly implements and 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): for highly substitutable occupations we have $Cov(d_{jj}, d_{jk}) \ll 0$. At the same time, prices for similar occupations have moved in the same direction ($Cov(p_j, p_k) > 0$). In the short regression only on cross-effects $\Delta e_j = \theta_2 \sum_{k \neq j} d_{jk} \Delta p_k + \varepsilon_j$, this leads to an OVB for θ_2 of $\frac{Cov(d_{jj} pj, \sum_{k \neq j} d_{jk} \Delta p_k)}{Var(\sum_{k \neq j} d_{jk} \Delta p_k)} \ll 0$.

Table C.1: Full Model: Two-Type Decomposition (OLS)

Two-Type Decomposition		Dependent Variable: $\Delta \mathbf{e}$ (1)	(2)	(3)	(4)
own effect:	$d_{jj} \Delta p_j$	1.81^{***}		4.10^{***}	
		(0.32)		(0.88)	4.15^{***}
total cross effect:	$\sum_{j \neq k} d_{jk} \Delta p_k$		-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: Regressor in column (4) is $\sum_j d_{jk} \Delta \ln p_k$, i.e., including own effect. Standard errors in parentheses; all coefficients shown are significant at the 1% level. Observations weighed by j's initial employment size. Period 1985–2010.

Table C.2: Full Model: Three-Type Decomposition (OLS). Unweighted.

Three-Type Decomposition		Dependent Variable: Δe_j — Unrestricted Model (1)	(2)	(3)	Restricted Model (4)	(5)
fixed relationship:	$\bar{d}_{diag} \Delta p_j$	1.70^{***}	1.93^{***}	3.82^{***}		
		(0.27)	(0.26)	(0.69)	1.93^{***}	
heter. own effect:	$(d_{jj} - \bar{d}_{diag}) \Delta p_j$		1.94^{***}	4.27^{***}	(0.27)	4.18^{***}
			(0.63)	(1.01)		(0.54)
total cross effect:	$\sum_{j \neq k} d_{jk} \Delta p_k$			3.32^{***}		
				(1.11)		
R-squared		0.264	0.318	0.377	0.318	0.366
Number of occupations		119	119	119	119	119

Notes: Regressor in column (4) is $d_{jj} \Delta \ln p_j$. In column (5), the regressor is $\sum_j d_{jk} \Delta \ln p_k$, i.e. corresponding to the full model. Standard errors in parentheses; all coefficients shown are significant at the 1% level. The number of occupations is 119 because the very tiny occupation 'Medical Receptionists' has been dropped. Period 1985–2010.

D Labour Demand, Equilibrium, and Estimation Strategy

This section extends the model by incorporating occupational labour demand. In what follows, we present the main features of the demand and supply sides, characterise equilibrium, and discuss its practical implementation.

D.1 Labour Demand and Equilibrium

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

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

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,

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

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

$$\tilde{E}_i^d\left(\tilde{p}_i(\mathbf{b},\mathbf{s}),\tilde{\beta}_i\right) = \ln\left(\tilde{\beta}_i\frac{1}{\tilde{\mathfrak{p}}_i}\right)^{\sigma} \tag{32}$$

where $\mathbf{b}$ is the $(N-1)$ vector of relative productivities, $\mathbf{s}$ 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 will be interested in perturbations around the steady state, so in keeping with the rest of the paper, we will 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 sectors 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 $j \neq i$:

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

A similar logic implies that $\frac{\partial e_i^d}{\partial \ln \frac{\beta_j}{1-\beta_j}}$ follows a similar structure.

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.2. The matrix of derivatives with respect to demand shifters is given by $\sigma \left(I - W \right)$, equally of rank

[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. (32) 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.

$N-1$.

D.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 $I-W$.

D.1.2 Equilibrium Characterisation

Similarly to before, we can write

$$e_i\left(\mathbf{b},\mathbf{s}\right) = e_i^s\left(\left\langle p\left(\mathbf{b},\mathbf{s}\right)\right\rangle, \mathbf{s}\right) = e_i^d\left(\left\langle p\left(\mathbf{b},\mathbf{s}\right)\right\rangle, \mathbf{b}\right) \tag{33}$$

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

In what follows, for ease of exposition, we 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. (33) with respect to $\ln\frac{\beta_j}{1-\beta_j}$ we obtain:

$$\sum_k \frac{\partial e_i^s}{\partial p_k}\frac{\partial p_k}{\partial\left(\ln\frac{\beta_j}{1-\beta_j}\right)} = \sum_k \frac{\partial e_i^d}{\partial p_k}\frac{\partial p_k}{\partial\left(\ln\frac{\beta_j}{1-\beta_j}\right)} + \frac{\partial e_i^d}{\partial \ln\frac{\beta_j}{1-\beta_j}} \tag{34}$$

Expressing this in matrix notation gives

$$\theta D V = \sigma\left(W-I\right)V + \sigma\left(I-W\right)$$

$$\implies (\theta D + \sigma (I - W)) V = \sigma (I - W) \tag{35}$$

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 $(\theta D + \sigma (I - W))$ has rank $N - 1$. However, we can also notice that $(I - W)$ is the de-meaning operator, such that for vector x, then $(I - W) x = x - 1_N \sum_i \tau_i x_i$, where 1_N is a column vector of ones. Therefore, we can solve eq. (35) 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

$$V = \left(\frac{\theta}{\sigma} D + I \right)^{-1} (I - W) \tag{36}$$

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:

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

Similarly to above, we can solve for S using a normalization 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:

$$\begin{aligned} S &= - (\theta D + \sigma I)^{-1} (I - W) \\ &= - \frac{1}{\sigma} V \end{aligned} \tag{37}$$

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

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

and then differentiating the identity $e_i\left(\mathbf{b},\mathbf{s}\right) = e_i^s\left(\langle p\left(\mathbf{b},\mathbf{s}\right)\rangle, \mathbf{b}\right)$ 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

$$\Xi = \theta D V$$

D.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

$$\Delta\mathbf{p} \approx V\Delta\mathbf{b} + S\Delta\mathbf{s} \tag{38}$$

and

$$\begin{aligned}\Delta\mathbf{e} &\approx \Xi\Delta\mathbf{b} + \Gamma\Delta\mathbf{s}\\ &= \theta D V\Delta\mathbf{b} - \sigma S\Delta\mathbf{s}\end{aligned} \tag{39}$$

These expressions describe changes to labour shares and prices in terms of demand and supply shocks, price elasticities, and model parameters θ and σ.

D.2 Estimation and Extraction of Shocks

D.2.1 Estimation Strategy

Expressions (38) and (39) also inform the regression framework. From eq. (38), 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. (39) yields:

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

$$= \theta D\Delta\mathbf{p} - (\theta D + \sigma)\, S\Delta\mathbf{s}$$
$$= \theta D\Delta\mathbf{p} - (-(I - W))\,\Delta\mathbf{s}$$
$$= \theta D\Delta\mathbf{p} + \Delta\mathbf{s} \tag{40}$$

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

Equation (40) is our basic regression equation, extending eq. (9) 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 $\Delta b_j \equiv \ln \frac{\beta_j}{1-\beta_j}$ but not with 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. (38):

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

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, $\bar{\epsilon} \equiv V\bar{\eta}$.

In terms of regressing Δe_j on the vector of price changes, this implies that, if $G = D\left(\frac{\theta}{\sigma}D + I\right)^{-1}$ and $\mathbf{g}_j$ is the ith row of this matrix, then an appropriate instrument for $\mathbf{d}_j\Delta\mathbf{p}$ is $\mathbf{g}_j\tilde{\mathbf{r}}$, or equivalently $\mathbf{g}_j\,(I - W)\,\mathbf{r}$.

For cases in which we ignore the cross-price effects and focus only on the own-price effects (i.e. assuming the off-diagonal elements of the elasticity matrix equal zero), we have that the vector $\tilde{\mathbf{r}}$ will be pre-multiplied by $G_{diag} = D_{diag}\left(\frac{\theta}{\sigma}D + I\right)^{-1}$.

We assume $\frac{\theta}{\sigma} = 2.3$ throughout the paper. In Table D.2, we show the robustness of our results to different values of $\frac{\theta}{\sigma}$.

D.2.2 Backing Out the Shocks

From (40) we immediately see that

$$\Delta\mathbf{s} \approx \Delta\mathbf{e} - \theta D \Delta\mathbf{p} \tag{41}$$

Similarly, from (38)

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

Summing with (39) this implies that

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

which follows from equation (35). 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

$$\Delta\mathbf{b} \approx \frac{1}{\sigma}\Delta\mathbf{e} + (I-W)\Delta\mathbf{p} \tag{42}$$

without loss of generality. Equations (41) and (42) can be used to construct the shock vectors.

D.3 Empirical Results on the Equilibrium Model

In this section we present empirical results from the equilibrium estimation.

Table D.1 presents summary statistics on the matrices V and DV that govern the dissipation of shocks to wages and employment in equations (16) and (17).[49] Parallel to the

[49] We use θ and σ from the equilibrium solution to the model. Table D.2 shows that different calibrations of σ hardly change estimated θ. Qualitative conclusions from Table D.1 also do not depend on the particular value of $\frac{\theta}{\sigma}$ that is used.

elasticity matrix in Table B.3, diagonal elements of V and DV (i.e., own-effects of shocks) are on average substantially larger than off-diagonal elements (cross-effects from shocks in other occupations) while, relative to the mean, standard deviations in the off-diagonal elements are higher. Off-diagonals in DV inherit some of the high skewness of D, whereas off-diagonals in V are not particularly skewed compared to the on-diagonals.

Table D.1: Summary Statistics. Matrices V and DV

	Matrix V		**Matrix DV**	
	Diagonal Elements	Off-Diagonal Elements ($\times 100$)	Diagonal Elements	Off-Diagonal Elements ($\times 100$)
Mean	0.508	-0.427	0.210	-0.177
Std. Dev.	0.079	0.916	0.035	0.319
Variance	0.006	0.839	0.001	0.102
Skewness	1.203	-1.482	-1.090	-7.490
Kurtosis	5.881	54.605	5.565	113.507
p10	0.418	-1.096	0.174	-0.397
p50	0.501	-0.226	0.215	-0.082
p90	0.587	0.016	0.249	-0.017
p99	0.746	1.109	0.279	-0.001
Correlation with D	-0.959	-0.246	0.968	0.935
Correlation with V			-0.989	0.016
Number of Observations	120	14,280	120	14,280

Notes: The table presents summary statistics for the matrices $V = \left(\frac{\theta}{\sigma}D + I\right)^{-1}(I - W)$ and DV, where we use the equilibrium solution for $\frac{\theta}{\sigma}$. See also the discussion in the text.

Interestingly, in contrast to matrix D and as discussed in the main text, off-diagonal elements in V can have opposite signs. This reflects that demand shocks in a given sector may have positive effects on prices in close substitute occupations while they have negative effects in more distant occupations. Table D.1 also reports that, overall, matrix elements from D are negatively correlated with those in V but positively with those in DV. This indicates that, ceteris paribus, larger own-price elasticities are associated with lower price changes and higher employment changes in response to demand shocks. Larger cross-price elasticities (more negative d_{jk}) tend to lead to more positive price responses and more negative employment responses from a demand shock in the respective other occupation.

Figure D.1: IV Reduced-Form for Employment

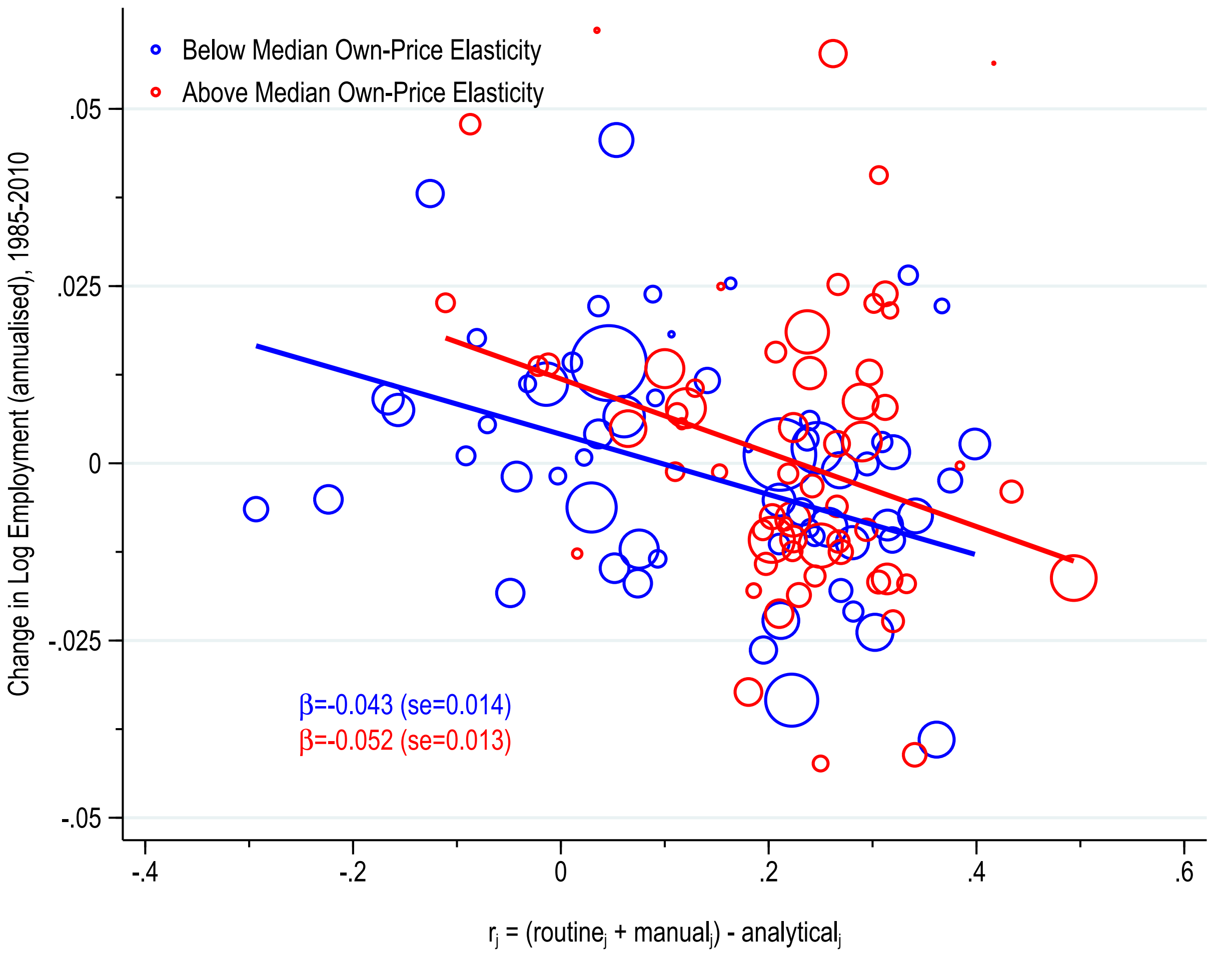

Notes: The figure shows reduced-form regressions of occupations' employment changes on their initial task contents r_j. Colour codes and linear regression lines are separated by occupation below (blue, inelastic) and above (red, elastic) the median predicted elasticity of labour supply with respect to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

Table D.2: Structural IV: Different Values of $\frac{\theta}{\sigma}$

	$\frac{\theta}{\sigma} = 1$	$\frac{\theta}{\sigma} = 1.5$	$\frac{\theta}{\sigma} = 2$	$\frac{\theta}{\sigma} = 2.3$	$\frac{\theta}{\sigma} = 2.5$	$\frac{\theta}{\sigma} = 3$	$\frac{\theta}{\sigma} = 4$
IV estimate for θ	4.95	4.87	4.81	4.78	4.76	4.72	4.66
Implied σ	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 is the benchmark used in the paper.

D.4 OLS versus IV Estimates

We wish to estimate (40), which is reproduced here for convenience:

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

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 (37) and (38) we note that

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

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 (20) 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 Model-Based Decomposition

This section develops in detail the counterfactual elasticity matrices introduced in Section 6.1, specifying how they relate 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. Finally, for the interested reader, we discuss an alternative way to construct counterfactuals that decomposes own-price and cross-price elasticities into a matrix product where each factor can be separately manipulated.

E.1 Counterfactual Elasticities

E.1.1 Heterogeneous Own-Price Elasticities Only

Counterfactual matrix D_{own} considers the case that occupations' aggregate (own-price) elasticities vary but their similarities with other occupations are homogeneous. In particular, we have $Cov_\tau\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right) = c \in [-1, 0]$ in eq. (6) and $Var_\tau\left(\tilde{\pi}_{.,j}\right) = -\frac{1-\tau_j}{\tau_j}c$ in eq. (7). Main diagonal elements of D_{own} continue to be the actual own-price elasticities, whereas cross-price elasticities reduce to size-weighted fractions of the on-diagonals $\frac{-\tau_k}{1-\tau_j}d_{jj}$.[50]

A specific version of this counterfactual can be derived from setups commonly used in the literature on firms, even if they study 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), take a simpler version of individuals' indirect utility eq. (1):[51]

$$u_j(\omega) = \theta p_j + a_j + \varepsilon_j(\omega), \tag{43}$$

Here switching costs a_j do not depend on the source employer i.

We derive the versions of Remarks 1–3, which result from eq. (43), 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 sector 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.

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

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

From this, we immediately 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, sectors are just all equally substitutable from a labour supply perspective. Remarks 2 and 3 then lead to

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

which is the result corresponding to Remark 1. The economic model with $a_{ij} = a_j$ thus generates a version of our matrix with homogenous occupational similarities D_{own}.

E.1.2 Fully Homogenous Labour Supplies

The second counterfactual imposes completely homogenous labour supply elasticities. The main diagonal elements of matrix D_{hom} become $\bar{d} = \sum_j \tau_j d_{jj}$ and cross-price elasticities a constant fraction of it $\frac{-\bar{d}}{N-1}$.[52] 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):

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

where $\tilde{\theta} \equiv \frac{N}{N-1} \bar{d} \theta$ is a single slope parameter on the price change and the second a regression constant that reflects overall average wage growth in the economy. In the equilibrium model (18) 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 $\Delta e_j = \tilde{\theta} \Delta p_j$.

The economic model would generate a specific version of D_{hom} with $\bar{d} = \frac{N-1}{N}$ and $\tilde{\theta} = \theta$ if, in addition to similarities, all occupation sizes are also the same. That is, when $\theta p_j + a_j = const.$ in eq. (43).

E.2 Results

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

Figure E.1 shows the impact of including labour supply heterogeneity in a counterfactual with no supply shocks ($\Delta \mathbf{s} = 0$). Figure E.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,

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

all occupational changes induced by demand shocks $\Delta\mathbf{b}$ run perfectly along a single supply curve. Figure E.1b then introduces both own- and cross-price effects keeping $\Delta\mathbf{s} = 0$. Relative to E.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 0.85. These changes show the importance of allowing for supply heterogeneity (and especially cross-occupation effects, which effectively reduce elasticities) to explain the data.

Figure E.2 plots the distribution of demand and supply shocks by occupation, showing a general positive correlation between the two (0.23). 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, pharmacists', which experienced a (large) positive demand shock but no supply shock.

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

First, our period of analysis is characterised by an increase in wage inequality and employment polarisation. These are represented in the figures by the dashed black lines and the 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 (16) and (17) 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 the full homogeneous model (i.e. comparing Figure E.4 to E.3), results in smaller wage changes and larger employment changes across occupations. The intuition for this is that, as we discuss in the main text, 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 E.1: Counterfactual Changes of Prices and Employment (II)

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

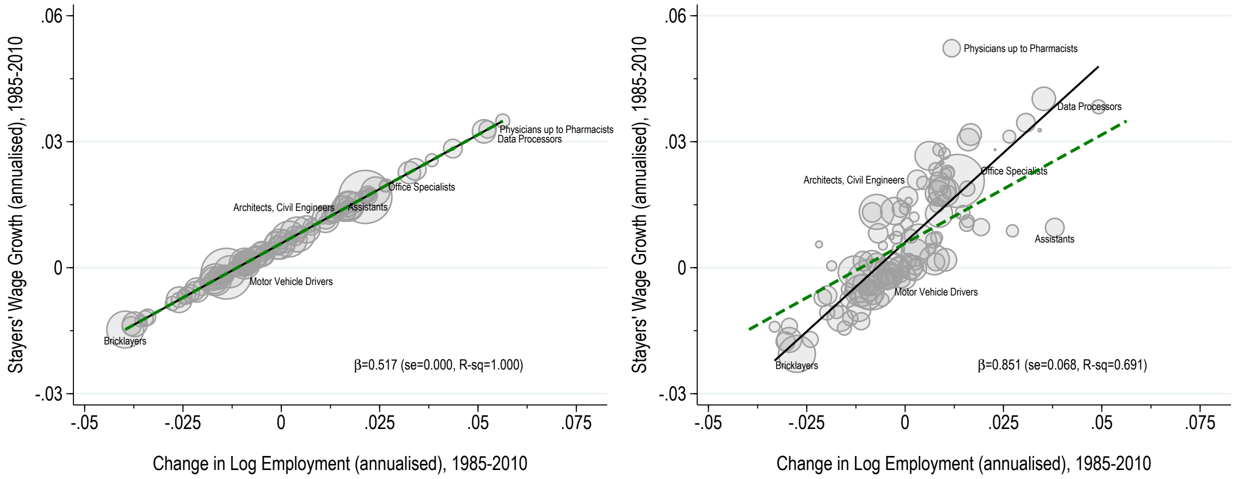

Notes: The figure shows price and employment changes for different manipulations of $\Delta\mathbf{s}$ and the elasticity matrix D. In Panel E.1a both supply shocks and heterogeneity in D are switched off, leaving only demand shocks. Panel E.1b introduces heterogeneous own- and cross-price elasticities (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 Panel E.1a is repeated as green-dashed in both panels.

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

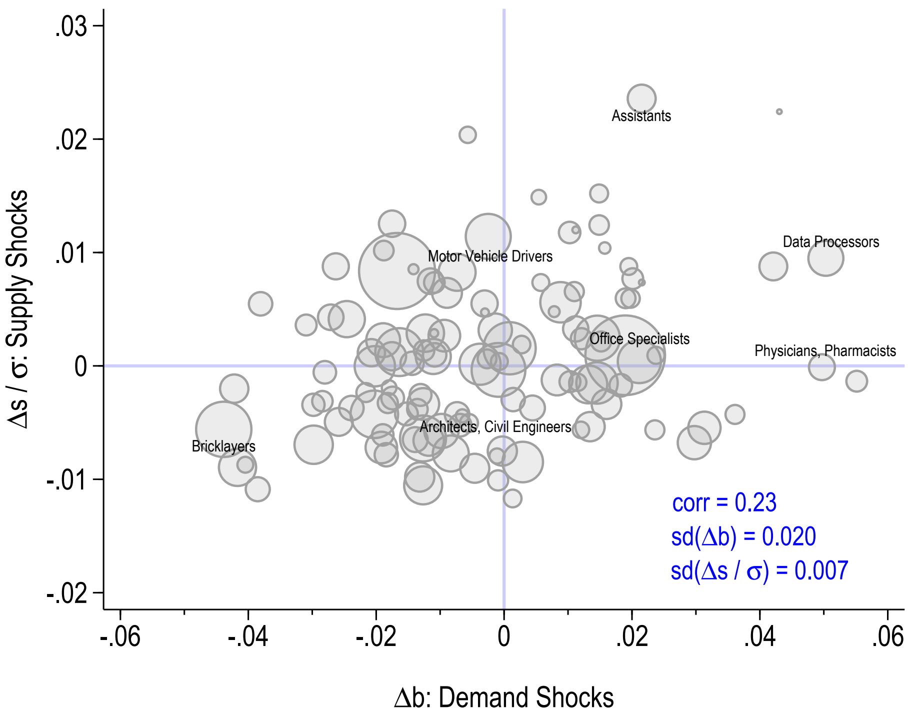

Notes: The figure shows the distribution of demand and supply shocks by occupation. Marker size indicates the baseline employment (in 1985) in each occupation. The standard deviations and correlation of demand and supply shocks are reported in the figure.

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

(a) ... to Wage Changes

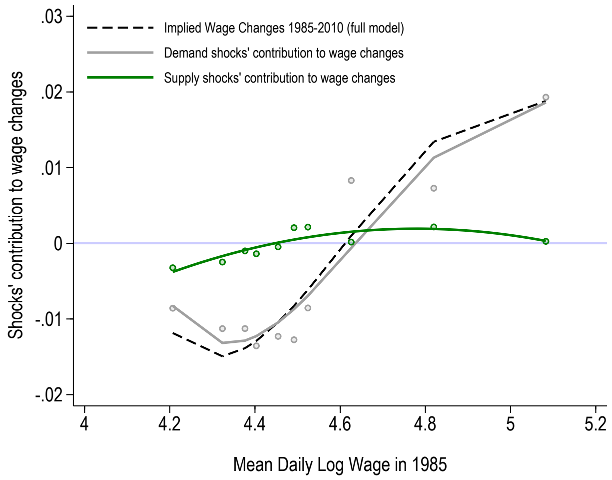

(b) ... to Employment Changes

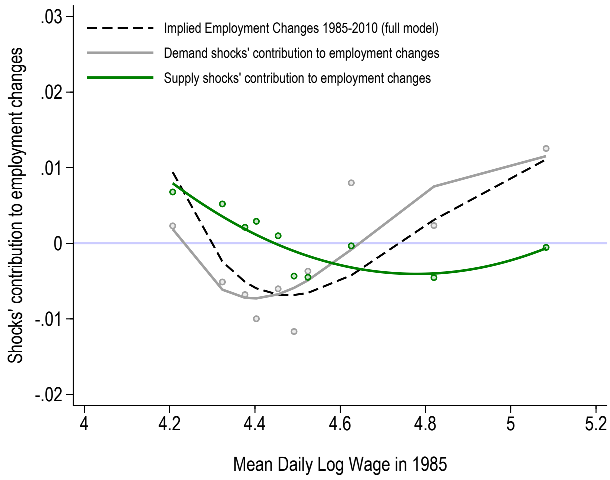

Notes: Figure E.3a shows the contributions to price changes of demand and supply shocks across the wage distribution for the full model. These are given by, $V\Delta\mathbf{b}$ and $-\frac{1}{\sigma}V\Delta\mathbf{s}$, as in eq. (16). Figure E.3b shows the contributions to employment changes of demand and supply shocks. These are given by $\theta DV\Delta\mathbf{b}$ and $V\Delta\mathbf{s}$, as in eq. (17). For supply, a quadratic is used for the smoothed fit. For demand, a fractional cubic is used.

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

(a) ... to Wage Changes

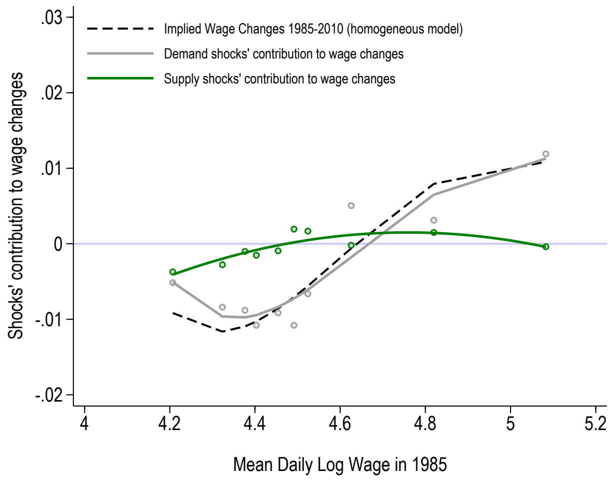

(b) ... to Employment Changes

Notes: Figure E.3a shows the contributions to price changes of demand and supply shocks across the wage distribution for the fully homogeneous model. These are given by, $V_{hom}\Delta\mathbf{b}$ and $-\frac{1}{\sigma}V_{hom}\Delta\mathbf{s}$, parallel to eq. (16). Figure E.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. (17). For supply, a quadratic is used for the smoothed fit. For demand, a fractional cubic is used.

E.3 Construction of Alternative Counterfactuals ** KEEP THIS? **

For the interested reader, this section suggests an alternative way of constructing counterfactual elasticity matrices. An attractive feature of these is that they are multiplicatively decomposed into separate matrices reflecting own- and cross-elasticity components. Empirically they lead to results very similar to those currently shown in the paper.

In our regression analysis in Section 4 we provided an informal additive decomposition of the matrix of elasticities. We can refine this by providing a full multiplicative decomposition, which is more useful in constructing counterfactuals. Specifically, Section E.3.1 below shows that the elasticity matrix can be decomposed as

$$D = \tilde{D}_{diag}\tilde{W} \tag{44}$$

where $\tilde{D}_{diag}$ is a diagonal matrix with jth diagonal element $\tilde{d}_j \equiv \frac{d_{jj}}{(1-\tau_j)}$, an occupation-size adjusted own-price elasticity, and $\tilde{W}$ the associated weighting matrix. Intuitively $\tilde{D}_{diag}$ can be interpreted as isolating heterogeneity in own-price elasticities, and $\tilde{W}$ can be interpreted as capturing the pure heterogeneity in cross-price effects between occupations j and k. Section E.3.1 further shows that we can express $\tilde{W}$ in terms of the occupational similarity measures, $Cov_{\tau_i}\left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$.

A more detailed way of constructing the first counterfactual is as follows. We consider the case that occupations' aggregate (own-price) elasticities vary but their similarities with other occupations are homogeneous. For example, employment in service industries may be responsive to price, but suppose that flows of workers into services come equally from any other occupation according to its size. Accordingly we construct a counterfactual matrix $\tilde{W}_{alt}$ derived from $\tilde{W}$ but imposing that the occupational similarity measures are constant for $j \neq k$. In this case, the weighting matrix reduces to $\tilde{W}_{alt} = (I - W)$, where, as previously, W is the matrix of stacked shares with j, kth element τ_k.

Building on this we can express the resulting counterfactual vector of employment changes as follows:

$$\begin{aligned} \Delta\mathbf{e}_{own} &= \theta D_{own}\Delta\mathbf{p} \\ &= \theta\tilde{D}_{diag}\left(I - W\right)\Delta\mathbf{p} \\ &= \theta\tilde{D}_{diag}\Delta\mathbf{p} \end{aligned} \tag{45}$$

where D_{own} is the alternative matrix of elasticities, generated as described. In the last line of eq. (45), this simplifies further, since $(I - W)$ is a de-meaning operator and because we can consider a vector of relative price changes with a (weighted) mean of zero

$(I - W) \Delta \mathbf{p} = \Delta \mathbf{p}$ without loss of generality. Equation (45) therefore shows that we can express this model using a diagonal matrix only. We term this the model with 'heterogeneous own-price effects' or 'homogeneous cross-price elasticities'.

Subsequently we use the fully homogenous case where all labour supply elasticities are identical across occupations. Accordingly we construct $\tilde{D}_{hom}$ by replacing diagonal elements in $\tilde{D}_{diag}$ with $\frac{\bar{d}}{1-\tau_j}$ where $\bar{d} = \sum \tau_j d_{jj}$, and $\tilde{W}$ with $I - W$. In this case

$$\begin{aligned} \Delta \mathbf{e}_{hom} &= \theta \tilde{D}_{hom} \Delta \mathbf{p} \\ &= \theta \bar{d} I \Delta \mathbf{p} \\ &= \theta \bar{d} \Delta \mathbf{p} \end{aligned} \tag{46}$$

and so employment shares are just scaled versions of price changes. We term this the 'fully homogeneous' model.

While different valid counterfactuals to what follows could be constructed, expressions (45)–(46) have attractive properties in that they represent alternative simpler frameworks for labour supply and that they are consistent with the inner workings of the model.[53] We now provide detailed derivations and discussions of these issues.

E.3.1 Decomposition

Remark 4 shows that our matrix of elasticities can be written concisely as

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

where W is the matrix of stacked shares with j,kth element τ_k and $\otimes$ is the element-by-element product. C is a symmetric matrix with j,kth element $c_{jk} = Cov_\tau \left(\tilde{\pi}_{.,j}, \tilde{\pi}_{.,k}\right)$, for $j \neq k$, capturing the 'occupational-similarity' between sectors j and k, and j,jth element $c_{jj} = Var_\tau \left(\tilde{\pi}\right)$, capturing the dispersion of in-flows into sector j.

Element d_{jk} can therefore be expressed as

$$\begin{aligned} d_{jk} &= \mathbb{1}_{j=k} - \tau_k - \tau_k c_{jk} \\ &= \frac{d_{jj} \left(\mathbb{1}_{j=k} - \tau_k - \tau_k c_{jk}\right)}{1 - \tau_j - \tau_j c_{jj}} \end{aligned}$$

[53] For example, imposing homogenous cross-elasticities via $d_{jk} = \frac{d_{jj}}{N-1}$ for all $k \neq j$ yields very similar results to those shown below. However, eq. (45) based on such a D_{own} does not feature homogeneous cross-occupation similarities, neither does it become a diagonal matrix product.

$$= \frac{\tilde{d}_j \left(1-\tau_j\right)\left(\mathbb{1}_{j=k}-\tau_k-\tau_k c_{jk}\right)}{1-\tau_j-\tau_j c_{jj}}$$

where

$$\tilde{d}_j \equiv \frac{d_{jj}}{\left(1-\tau_j\right)} \tag{47}$$

is an occupation-size adjusted own-price elasticity. We can hence write D as

$$D = \tilde{D}_{diag}\tilde{W}$$

where $\tilde{D}_{diag}$ is a diagonal matrix with jth diagonal element $\tilde{d}_j$ and $\tilde{W}$ has j,kth element given by:

$$\tilde{w}_{jk} = \frac{\left(1-\tau_j\right)\left(\mathbb{1}_{j=k}-\tau_k-\tau_k c_{jk}\right)}{1-\tau_j-\tau_j c_{jj}}$$

E.3.2 Counterfactual Construction

We use formulation (44), i.e., $D = \tilde{D}_{diag}\tilde{W}$ with elements $d_{jk} = \tilde{d}_j\tilde{w}_{jk}$, to construct alternative sets of elasticities. Above we discussed constructing D_{own} by replacing $\tilde{W}$ with $I - W$. This results from replacing similarities c_{jk} of j with all k to be homogeneous and, since C is symmetric, this becomes $c_{jk} = c \leq 0\ \forall j \neq k$. The variances on C's main diagonal become $c_{jj} = -\frac{1-\tau_j}{\tau_j}c \geq 0$.

Checking the result, in the case of off-diagonal $\tilde{w}_{jk}\ \forall j \neq k$ we get

$$\tilde{w}_{jk} = \frac{\left(1-\tau_j\right)\left(-\tau_k-\tau_k c\right)}{1-\tau_j+(1-\tau_j)c} = \frac{-\tau_k\left(1-\tau_j\right)\left(1+c\right)}{\left(1-\tau_j\right)\left(1+c\right)} = -\tau_k$$

For on-diagonal elements $j = k$ we have

$$\tilde{w}_{jk} = \frac{\left(1-\tau_j\right)\left(1-\tau_j-\tau_j c_{jj}\right)}{1-\tau_j-\tau_j c_{jj}} = \frac{\left(1-\tau_j\right)\left(1-\tau_j+(1-\tau_j)c\right)}{\left(1-\tau_j\right)\left(1+c\right)} = \frac{\left(1-\tau_j\right)^2\left(1+c\right)}{\left(1-\tau_j\right)\left(1+c\right)} = 1-\tau_j$$

Therefore, in this counterfactual, we can write the weighting matrix as $\tilde{W} = (I - W)$. This ensures that D_{own} is a valid matrix of elasticities. The thought experiment here is neutralizing cross-occupation similarity in the $\tilde{W}$ component. Note however, that this does not imply eliminating all cross-occupation substitutability $\tau_j c$ (i.e., due to size differences) in $D_{own} = \tilde{D}_{diag}\left(I - W\right)$.

Developing this discussion, we can derive the ensuing counterfactual vector of changes in employment shares, $\Delta\mathbf{e}_{own}$, as:

$$\begin{aligned}\Delta\mathbf{e}_{own} &= \theta D_{own}\Delta\mathbf{p} \\ &= \theta\tilde{D}_{diag}\left(I - W\right)\Delta\mathbf{p} \\ &= \theta\tilde{D}_{diag}\Delta\mathbf{p} + q\mathbf{1}\end{aligned}$$

where $\mathbf{1}$ is a vector of ones, q is a constant, and $q\mathbf{1}$ ensures that $\Delta\mathbf{e}_{own}$ has a weighted sum of zero. Alternatively, as in the main text, we can normalise $\Delta\mathbf{p}$ to have a weighted mean of zero without loss of generality, in which case $q = 0$.

E.3.3 Counterfactual Performance

The above counterfactual aims to remove heterogeneity in cross-price elasticities while, ideally, preserving steady-state employment shares and preserving heterogeneity in aggregate substitutability across occupations. It turns out that D_{own} in practice maintains well these desired properties of a counterfactual.

We investigate performance by using the definitions (5) in the main body, which we restate here:

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

This can be used to solve for an implied counterfactual transition matrix and counterfactual stationary employment shares. For computational reasons we performed this on a pairwise aggregation of our 120 occupations, yielding 60 occupations. In this case the correlation of the counterfactual and the baseline stationary shares is 0.97. Furthermore the diagonal elements of D_{own} match those of the empirical elements of D by construction, so aggregate substitutability is preserved. Finally we aim to remove heterogeneity in cross-price elasticities. To check this we computed matrix C as in equation (28), and equivalent counterfactual C_{own} for the new matrix D_{own}. The average off-diagonal element of C is -0.50 with a standard deviation of 0.89, while the average off-diagonal element of C_{own} is -0.54, with a standard deviation of 0.16, a reduction of around 80%. Even if heterogeneity in cross-occupational similarity is not totally eliminated in our counterfactual, therefore, it is substantially reduced.

F Extensions and Robustness: Supplementary Material

In this section, we present the details of the extensions and robustness checks of the main findings. We first extend the model to incorporate non-employment transitions. Second, we study changes in occupational prices and employment by sub-periods. Lastly, we introduce an alternative method for estimating changes in occupational prices.

F.1 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 particularly important if young workers' entry and old workers' exit from the labour market affect specific occupations' growth. The secular decline of German unemployment from the mid-2000s may also be relevant in this respect.

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 in order to model-consistently 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. Then consider eq. (11) with $N+M$ occupations, with M referring to non-employment sectors:

$$\Delta e_j \approx \theta \sum_{k=1}^{N+M} d_{jk} \Delta p_k = \theta \sum_{k=1}^{N} d_{jk} \Delta p_k + \sum_{m=N+1}^{N+M} (\theta \Delta p_m) d_{jm} \tag{48}$$

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 d_{jm} are control variables for the occupation j's predicted 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'. This product cannot be disentangled, as Δp_m is unobserved, but other than that the model is again identified.

In what follows, we show the results from these estimations with $M = 3$ 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. A limitation of the records from unemployment insurance is that we cannot observe the exact reasons for individuals entering or leaving the dataset (e.g. health shock, discouraged worker, emigration, self-employment, military service or becoming a civil servant). Outside the age range for labour market entry or retirement, these are all treated as out of the labour force for our purposes. As shown in Table F.1 (for both the restricted and the unrestricted models) and Table F.2, the R-squared is consistently higher in all specifications as more of the heterogeneity in employment growth can be explained when allowing for occupations' different elasticities with respect to non-employment states. Importantly, the estimated role of own- and cross-price effects turn out similar to the main results (both OLS and IV estimates). In unreported results, we verify the main results do not change when further separating part-time work and work with benefit receipt from 'out of the labour force', or when merging the three states into one single non-employment sector.

Could refer to column 2 of Table B.5. Could also create an analogous figure to Figure 2b by median d_{jj}^{NE}.

Table F.1: Accounting for Non-Employment Transitions (I)

Three-Type Decomposition		Dependent Variable: $\Delta\mathbf{e}$				
		Unrestricted Model			Restricted Model	
		(1)	(2)	(3)	(4)	(5)
fixed relationship:	$\overline{d}_{diag}\Delta p_j$	2.41***	2.55***	3.70***		
		(0.41)	(0.40)	(0.74)	2.49***	
heter. own effect:	$(d_{jj} - \overline{d}_{diag})\Delta p_j$		1.41**	3.27***	(0.39)	4.06***
			(0.63)	(0.99)		(0.68)
total cross effect:	$\sum_{j\neq k} d_{jk}\Delta p_k$			2.83**		
				(1.19)		
elast wrt unemp:	d_{jN+1}	−0.59***	−0.58***	−0.54***	−0.52***	−0.47***
		(0.17)	(0.16)	(0.14)	(0.14)	(0.13)
elast wrt olf:	d_{jN+2}	0.25	0.18	0.24	0.07	0.22
		(0.20)	(0.17)	(0.19)	(0.12)	(0.20)
elast wrt entry/exit:	d_{jN+3}	0.04	0.04	0.03	0.05	0.04
		(0.19)	(0.19)	(0.19)	(0.19)	(0.19)
R-squared		0.421	0.438	0.470	0.426	0.463
Number of occupations		120	120	120	120	120

Notes: Regressor in column (4) is $d_{jj}\Delta \ln p_j$. In column (5), the regressor is $\sum_j d_{jk}\Delta \ln p_k$, i.e. corresponding to the full model. We consider $M = 3$ different non-employment sectors: unemployment 'unemp', out of the labour force 'olf' (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. Standard errors in parentheses; *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Observations weighted by occupation j's initial employment size. Period 1985–2010.

Table F.2: Accounting for Non-Employment Transitions (II)

Three-Type Decomposition (Restricted Model)		Dependent Variable: $\Delta \mathbf{e}$					
		(1)		(2)		(3)	
		OLS	IV	OLS	IV	OLS	IV
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	2.41***	1.70***				
		(0.41)	(0.49)				
own effect:	$d_{jj}\Delta p_j$			2.49***	1.75***		
				(0.39)	(0.49)		
own & cross effect:	$\sum_{k=1}^{N} d_{jk}\Delta p_k$					4.06***	4.48***
						(0.68)	(1.24)
elast wrt unemp:	d_{jN+1}	−0.59***	−0.37*	−0.52***	−0.32*	−0.47***	−0.54**
		(0.17)	(0.19)	(0.14)	(0.17)	(0.13)	(0.23)
elast wrt olf:	d_{jN+2}	0.25	0.03	0.07	−0.11	0.22	0.29
		(0.20)	(0.27)	(0.12)	(0.18)	(0.20)	(0.29)
elast wrt entry/exit:	d_{jN+3}	0.04	0.07	0.05	0.08	0.04	0.03
		(0.19)	(0.19)	(0.19)	(0.20)	(0.19)	(0.18)
Observations		120	120	120	120	120	120
R-squared		0.421	-	0.426	-	0.463	-
F-stat 1st Stage		-	102	-	53	-	36

Notes: OLS and instrumental variable two-stage least squares (IV-2SLS) estimation results of the restricted model (48) controlling for non-employment transitions in matrix D of dimension $N+M$. We consider $M = 3$ different non-employment sectors: unemployment 'unemp', out of the labour force 'olf' (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. For the IV, in columns (1)-(2), the instrument is $D_{diag}\left(\frac{\theta}{\sigma}D+I\right)^{-1}(I-W)\mathbf{r}$. In column (3), the instrument is $D\left(\frac{\theta}{\sigma}D+I\right)^{-1}(I-W)\mathbf{r}$. Standard errors are in parentheses; *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Observations weighted by j's initial employment size. Period 1985–2010.

F.2 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 x 5 sub-periods) is used to estimate an extended version of eq. (11):

$$\Delta e_{jt} = \alpha + \theta d_{jj}\Delta p_{jt} + \theta \sum_{k\neq j} d_{jk}\Delta p_{kt} + \delta_t\ (+\gamma_j) + \varepsilon_{jt} \tag{49}$$

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).[54] 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 F.1 plots prices against employment growth for the pooled sample of 600 occupation-sub-periods (panel a) as well as separately for each sub-period (panel b), analogous to the main text Table 2. The previous finding is strengthened in the sense that each regression slope for above-median own-price elastic occupations is flatter than any slope for below-median own-price inelastic occupations. From Table F.3 and Table F.4, we see 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. Note that we can only do the OLS for this as our instrument does not vary by period. In sum, estimation in a series of shorter intervals shows that the role of occupational price elasticities persists, with some evidence that even acceleration/deceleration of price growth in different sub-periods is translated into employment growth according to these elasticities.

Table F.3: Full Model Pooled Sub-Periods (OLS)

	Dependent Variable: $\Delta \mathbf{e}$				
Three-Type Decomposition	Unrestricted Model			Restricted Model	
	(1)	(2)	(3)	(4)	(5)
fixed relationship: $\bar{d}_{diag}\Delta p_j$	1.69***	1.93***	3.90***		
	(0.29)	(0.29)	(0.67)	2.41***	
heter. own effect: $(d_{jj} - \bar{d}_{diag})\Delta p_j$		1.57***	4.04***	(0.34)	4.43***
		(0.30)	(0.83)		(0.70)
total cross effect: $\sum_{j \neq k} d_{jk}\Delta p_k$			3.69***		
			(1.09)		
R-squared	0.419	0.438	0.492	0.456	0.486
Period fe	yes	yes	yes	yes	yes
Cluster Std. Errors	yes	yes	yes	yes	yes
Number of occupations	600	600	600	600	600

Notes: Results on pooled panel sample containing 600 observations (120 occupations x 5 sub-periods). Sub-periods are: 1985-1990, 1990-1995, 1995-2000, 2000-2005, and 2005-2010. The regressor in column (4) is $d_{jj}\Delta \ln p_j$. In column (5), the regressor is $\sum_j d_{jk}\Delta \ln p_k$, i.e. corresponding to the full model. Standard errors clustered at the occupation level in parentheses; all coefficients shown are significant at the 1% level. 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).

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

Figure F.1: Stayers' Wage Growth and Employment Changes (by Own-Price Elasticity Median Split)

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

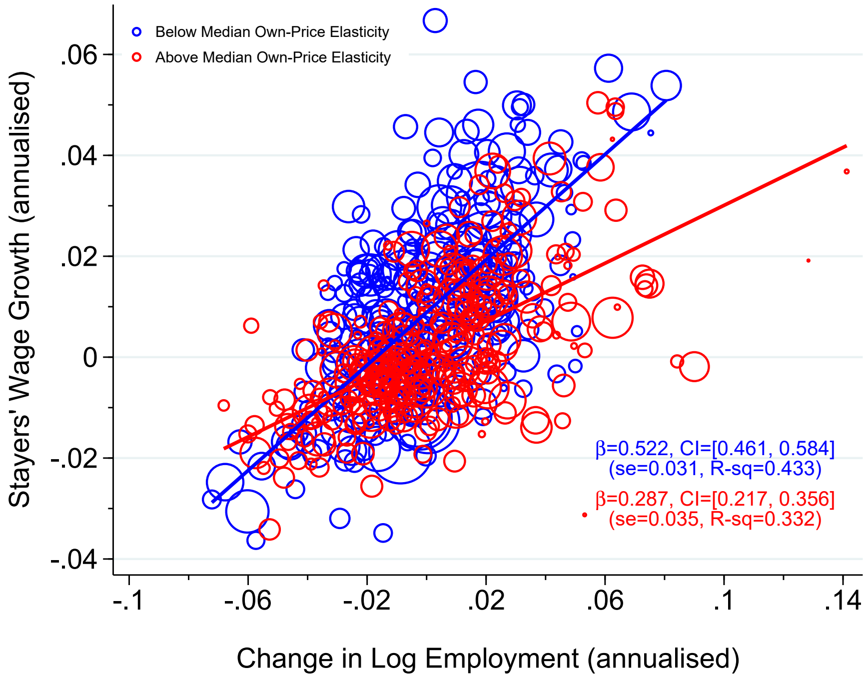

(b) By Sub-Periods.

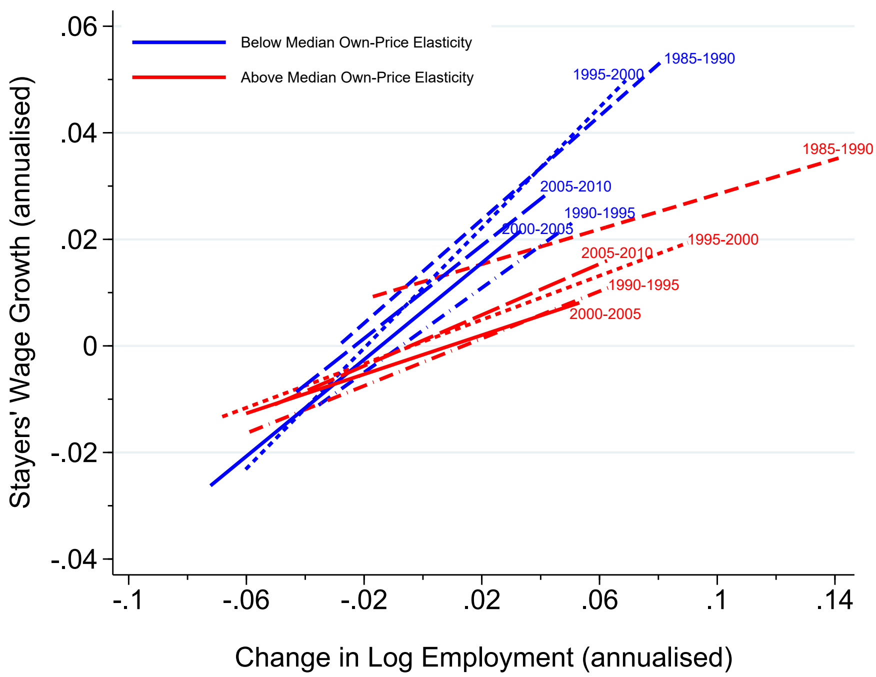

Notes: Figure shows scatter plots of occupations' change in the log of employment (x-axis) and stayers' wage growth (y-axis) for the pooled sample of 600 occupation-sub-periods (panel (a)) as well as separately for each sub-period (panel (b)). Sub-periods are: 1985-1990, 1990-1995, 1995-2000, 2000-2005, and 2005-2010. The graphs depict colour codes and linear regression lines by occupations below (blue, inelastic) and above (red, elastic) the median predicted elasticity of labour supply with respect to own price (d_{jj}). β refers to the slope coefficient, CI stands for the 95% confidence interval, se refers to standard error, and $R\text{-}sq$ stands for the R squared of the regression.

Table F.4: Full Model Pooled Sub-Periods (OLS-IV)

Three-Type Decomposition (Restricted Model)	Dependent Variable: $\Delta\mathbf{e}$								
	(1)			(2)			(3)		
	OLS	OLS	IV	OLS	OLS	IV	OLS	OLS	IV
	period fe	period & occ fe	period fe	period fe	period & occ fe	period fe	period fe	period & occ fe	period fe
fixed relationship: $\bar{d}_{diag}\Delta p_j$	1.69***	2.89***	1.09***						
	(0.29)	(0.40)	(0.41)						
own effect: $d_{jj}\Delta p_j$				1.93***	2.81***	1.18***			
				(0.29)	(0.40)	(0.44)			
own & cross effect: $\sum_{k=1}^{N} d_{jk}\Delta p_k$							4.01***	3.18***	4.17***
							(0.56)	(0.51)	(1.32)
Observations	600	600	600	600	600	600	600	600	600
R-squared	0.419	0.796	-	0.437	0.799	-	0.491	0.791	-
F-stat 1st Stage	-	-	124	-	-	111	-	-	13

Notes: OLS and instrumental variable two-stage least squares (IV-2SLS) estimation results of the restricted pooled model (49). The pooled panel sample contains 600 observations (120 occupations x 5 sub-periods). Sub-periods are: 1985-1990, 1990-1995, 1995-2000, 2000-2005, and 2005-2010. For the IV, in columns (1)-(2), the instrument is $D_{diag}\left(\frac{\theta}{\sigma}D + I\right)^{-1}(I - W)\mathbf{r}$. In column (3), the instrument is $D\left(\frac{\theta}{\sigma}D + I\right)^{-1}(I - W)\mathbf{r}$. Standard errors clustered at the occupation level in parentheses; all coefficients shown are significant at the 1% level. 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).

F.3 Alternative Price Estimation

The main results in Section 4 use the annual wage growth of occupation stayers as the main estimate of an occupation's changing log price or wage rate per efficiency unit of skill. 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 estimated by

$$\ln w_t(\omega) = \sum_j Z_{jt}(\omega)\varphi_{jt} + \sum_j Z_{jt}(\omega)X_t(\omega)\zeta_j + \sum_j Z_{jt}\kappa_j(\omega) + \mu_t(\omega) \quad (50)$$

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.[55] 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 (Δp_j).

The results using prices á la Cortes (2016) turn out similar to our main results. The main figures of the paper are replicated using these alternative prices in Figure F.2. The main regression results (Table F.5-F.6-F.7-F.8), including those when accounting for non-employment transitions, turn out very similar. Our findings hence remain consistent and robust to this alternative price estimation.

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

Table F.5: Prices á la Cortes (2016) and Employment: Full Model (OLS)

Three-Type Decomposition		Dependent Variable: $\Delta \mathbf{e}$				
		Unrestricted Model			Restricted Model	
		(1)	(2)	(3)	(4)	(5)
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	2.23***	2.70***	4.46***	2.70***	5.18***
		(0.45)	(0.49)	(1.30)	(0.49)	(1.15)
heter. own effect:	$(d_{jj} - \bar{d}_{diag})\Delta p_j$		2.25***	4.65***		
			(0.67)	(1.73)		
total cross effect:	$\sum_{j \neq k} d_{jk}\Delta p_k$			3.23*		
				(1.81)		
R-squared		0.287	0.340	0.371	0.337	0.350
Number of occupations		120	120	120	120	120

Notes: Results using prices á la Cortes (2016). The regressor in column (4) is $d_{jj}\Delta \ln p_j$. In column (5), the regressor is $\sum_j d_{jk}\Delta \ln p_k$, i.e. corresponding to the full model. Standard errors in parentheses; *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Observations weighted by j's initial employment size. Period 1985–2010.

Table F.6: Prices á la Cortes (2016) and Employment: Full Model (OLS-IV)

Three-Type Decomposition (Restricted Model)		Dependent Variable: $\Delta \mathbf{e}$					
		(1)		(2)		(3)	
		OLS	IV	OLS	IV	OLS	IV
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	2.23***	1.71***				
		(0.45)	(0.50)				
own effect:	$d_{jj}\Delta p_j$			2.70***	2.05***		
				(0.49)	(0.57)		
own & cross effect:	$\sum_{k=1}^{N} d_{jk}\Delta p_k$					5.18***	6.48***
						(1.15)	(2.12)
Number of Occupations		120	120	120	120	120	120
R-squared		0.287	-	0.337	-	0.350	-
F-stat 1st Stage		-	91	-	69	-	10

Notes: OLS and instrumental variable two-stage least squares (IV-2SLS) estimation results of the restricted model using prices á la Cortes (2016). In columns (1)-(2), the instrument is $D_{diag}\left(\frac{\theta}{\sigma}D + I\right)^{-1}(I - W)\mathbf{r}$. In column (3), the instrument is $D\left(\frac{\theta}{\sigma}D + I\right)^{-1}(I - W)\mathbf{r}$. Standard errors are in parentheses; all coefficients shown are significant at the 1% level. Observations weighted by j's initial employment size. Period 1985–2010.

Figure F.2: Occupational Prices á la Cortes (2016) and Employment

(a) Overall Relationship

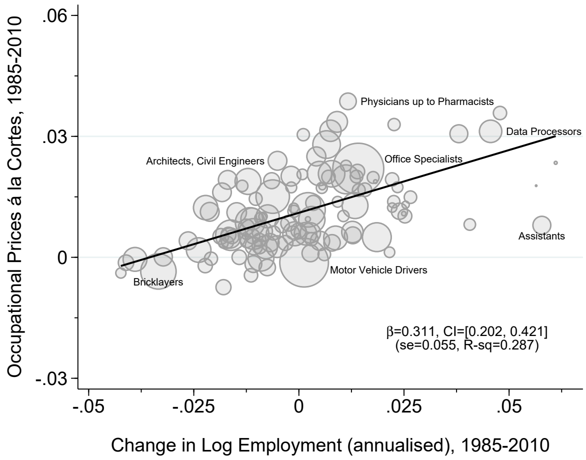

(b) By d_{jj} Median Split

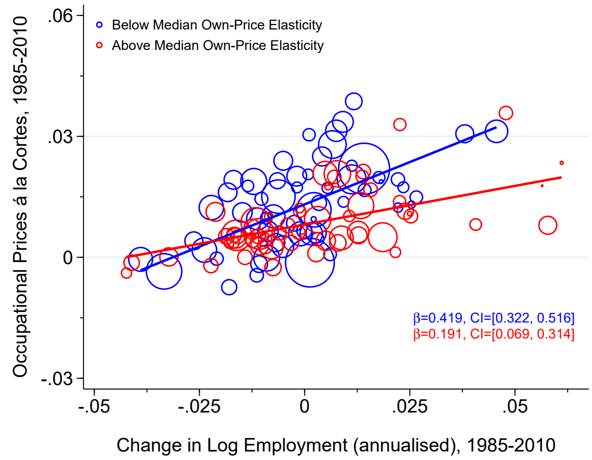

(c) By d_{jj} Quartiles

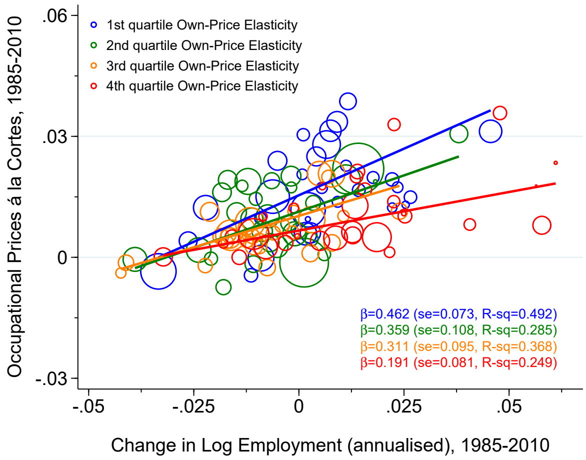

(d) IVfigure (TBA)

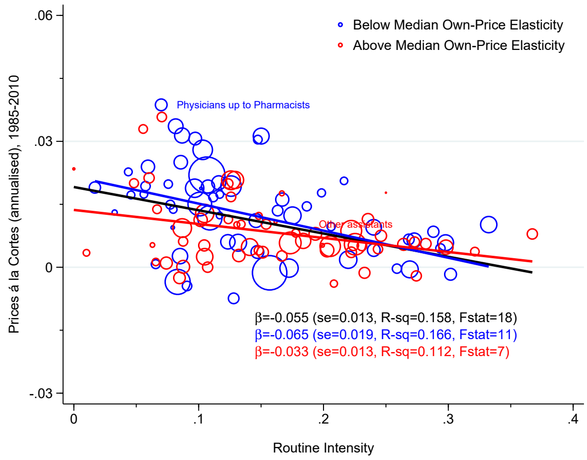

(e) Pooled Sub-Periods. By d_{jj} median split

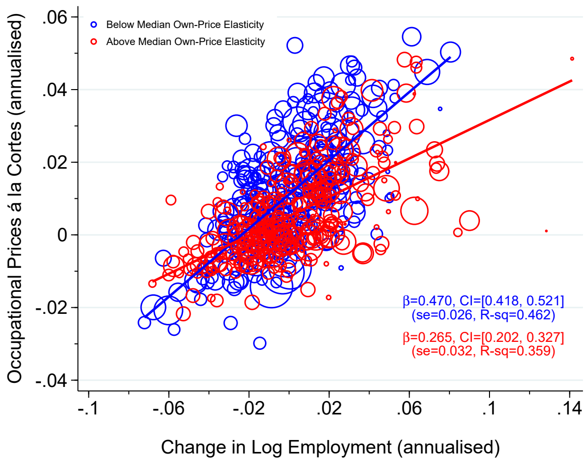

(f) Sub-Periods. By d_{jj} median split

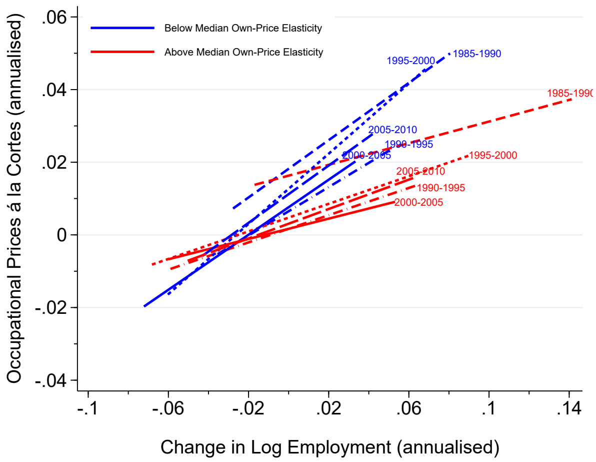

Notes: Panels A-C show the scatter of occupations' change in log of total employment (x-axis) and task prices á la Cortes (2016) (y-axis) during 1985–2010. Panel (a) shows the overall regression line. Panel (b) shows colour codes by occupations below (blue, inelastic) and above (red, elastic) the median predicted elasticity of labour supply with respect to own price d_{jj}. Panel (c) shows colour codes by occupations in the lowest (blue), second (green), third (orange), and highest (red) quartile of the predicted elasticity of labour supply with respect to own price d_{jj}. Panel (d) shows first-stage regressions of price changes (á la Cortes (2016)) on each occupation's routine task intensity. Panel (e) shows the overall regression line for the pooled 600 occupations x sub-periods case. Finally, panel (f) split by d_{jj} median the pooled occupations x sub-periods sample.

Table F.7: Prices á la Cortes (2016). Accounting for Non-Employment (OLS-IV)

Three-Type Decomposition (Restricted Model)		Dependent Variable: $\Delta \mathbf{e}$					
		(1)		(2)		(3)	
		OLS	IV	OLS	IV	OLS	IV
fixed relationship:	$\bar{d}_{diag} \Delta p_j$	3.05^{***}	2.25^{***}				
		(0.63)	(0.68)				
own effect:	$d_{jj} \Delta p_j$			3.00^{***}	2.29^{***}		
				(0.57)	(0.68)		
own & cross effect:	$\sum_{k=1}^{N} d_{jk} \Delta p_k$					4.76^{***}	5.45^{***}
						(1.10)	(1.78)
elast wrt unemp:	d_{jN+1}	-0.49^{***}	-0.32	-0.30^{**}	-0.19	-0.29^{**}	-0.36
		(0.18)	(0.20)	(0.13)	(0.16)	(0.14)	(0.24)
elast wrt olf:	d_{jN+2}	0.08	-0.08	-0.09	-0.19	0.06	0.14
		(0.19)	(0.25)	(0.12)	(0.17)	(0.20)	(0.28)
elast wrt entry/exit:	d_{jN+3}	0.06	0.09	0.13	0.13	0.12	0.12
		(0.20)	(0.20)	(0.21)	(0.21)	(0.20)	(0.20)
Observations		120	120	120	120	120	120
R-squared		0.386	-	0.401	-	0.402	-
F-stat 1st Stage		-	79	-	28	-	23

Notes: Results using prices á la Cortes (2016). OLS and instrumental variable two-stage least squares (IV-2SLS) estimation results of the restricted model (48) controlling for non-employment transitions in matrix D of dimension $N + M$. We consider $M = 3$ different non-employment sectors: unemployment 'unemp', out of the labour force 'olf' (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. For the IV, in columns (1)-(2), the instrument is $D_{diag} \left(\frac{\theta}{\sigma} D + I \right)^{-1} (I - W)\mathbf{r}$. In column (3), the instrument is $D \left(\frac{\theta}{\sigma} D + I \right)^{-1} (I - W)\mathbf{r}$. Standard errors are in parentheses; *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$. Observations weighted by j's initial employment size. Period 1985–2010.

Table F.8: Prices á la Cortes (2016). Full Model Pooled Sub-Periods (OLS-IV)

Three-Type Decomposition (Restricted Model)		Dependent Variable: $\Delta\mathbf{e}$								
		(1)			(2)			(3)		
		OLS	OLS	IV	OLS	OLS	IV	OLS	OLS	IV
		period fe	period & occ fe	period fe	period fe	period & occ fe	period fe	period fe	period & occ fe	period fe
fixed relationship:	$\bar{d}_{diag}\Delta p_j$	2.06***	2.85***	1.33***						
		(0.35)	(0.42)	(0.50)						
own effect:	$d_{jj}\Delta p_j$				2.41***	2.78***	1.51***			
					(0.34)	(0.42)	(0.54)			
own & cross effect:	$\sum_{k=1}^{N} d_{jk}\Delta p_k$							4.43***	3.24***	4.92***
								(0.70)	(0.55)	(1.56)
Observations		600	600	600	600	600	600	600	600	600
R-squared		0.426	0.789	-	0.456	0.792	-	0.486	0.785	-
F-stat 1st Stage		-	-	110	-	-	85	-	-	11

Notes: Results using prices á la Cortes (2016). OLS and instrumental variable two-stage least squares (IV-2SLS) estimation results of the restricted pooled model (49). The pooled panel sample contains 600 observations (120 occupations x 5 sub-periods). Sub-periods are: 1985-1990, 1990-1995, 1995-2000, 2000-2005, and 2005-2010. For the IV, in columns (1)-(2), the instrument is $D_{diag}\left(\frac{\theta}{\sigma}D+I\right)^{-1}(I-W)\mathbf{r}$. In column (3), the instrument is $D\left(\frac{\theta}{\sigma}D+I\right)^{-1}(I-W)\mathbf{r}$. Standard errors clustered at the occupation level in parentheses; all coefficients shown are significant at the 1% level. 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).