Does Skill Heterogeneity Affect Aggregate Employment-Wage Comovements?*

John Grigsby
Department of Economics, Princeton University

Monday 21st April, 2025

Abstract

I study aggregate employment and wage movements when workers have heterogeneous skills for tasks which are subject to idiosyncratic labor demand shocks. Skill heterogeneity changes the elasticity of labor supply and thus the equilibrium response of the labor prices to demand shocks. The composition of labor demand shocks has consequences for aggregate employment and wages. Aggregate employment-wage comovements partly reflect reallocation of different workers across tasks and into employment, which representative agent economies would interpret as a labor supply shock or labor wedge. A calibrated frictionless model with realistic skill heterogeneity and labor demand shocks can replicate the employment collapse and mean wage increase of the Great Recession.

JEL Codes: E24, J24

KEYWORDS: Labor Supply, Labor Demand, Aggregation.

*I am greatly indebted to my thesis committee Erik Hurst, Ufuk Akcigit, Greg Kaplan and Robert Shimer for invaluable insight and support. In addition, I am grateful to Mark Aguiar, Mark Bils, Stéphane Bonhomme, Jonathon Hazell, Ali Hortaçsu, Thibaut Lamadon, Simon Mongey, Kevin Murphy, Jeremy Pearce, Richard Rogerson, Gustavo da Souza, Gianluca Violante, Liangjie Wu, and Yulia Zhestkova, as well as numerous seminar participants for their many helpful comments. A prior version of this paper was circulated under the title "Skill Heterogeneity and Aggregate Labor Market Dynamics."

Grigsby: Julis Romo Rabinowitz Building, Washington Rd, Princeton, NJ 08544, john.grigsby@princeton.edu

1. INTRODUCTION

Workers differ in their skills in multiple dimensions – the skills required of roofers are possessed by different individuals to the vastly different skills required of surgeons. Some recent evidence suggests that both inequality and the distribution of skills have changed over time.[1] Furthermore, the composition of labor demand shocks is not constant; for example, the Great Recession featured a collapse in demand for construction workers, while the Pandemic Recession saw demand for in-person service workers fall. Two natural questions therefore arise. How do economies with different skill distributions differ in their short-run response of aggregate employment and wages to labor demand shocks? And does the answer to this question depend on the exact composition of labor demand shocks hitting the economy?

To study these questions, I develop a model in which workers belong to one of a finite set of skill types, which differ in their skills for a variety of occupations. Workers choose their occupation as in a Roy model and have access to a home option. Workers are paid in proportion to their skill according to the price per unit of labor in the market for each occupation. Shocks to occupation-specific labor demand shift these prices, thereby reallocating workers into/out of employment and across occupations. Markets are competitive and the model's equilibrium is efficient.

The model offers three new implications relative to a representative agent benchmark. First, skill heterogeneity affects both the own- and cross-price elasticities of labor supply for a particular job-type. As skills become more specific, movements in labor prices in one occupation have less of an effect on labor supply elsewhere, as labor supply of one worker type is irrelevant to portions of the economy. That is, skill specificity dampens the cross-price elasticity of labor supply. What's more, the degree to which skill for an occupation is embodied within a few individuals increases the own-price elasticity of labor supply for that occupation, because the supply of total labor units is dictated by the decisions of just a few individuals. Skill heterogeneity thus influences the general equilibrium effects of granular labor demand shocks in a way that is not captured solely by a representative agent model, or indeed a model with vertically differentiated skills as in Chang and Kim (2007).

Second, movements in labor demand induce different relationships between aggregate employment and wages depending on the distribution of skills in the economy and the composition of labor demand shocks. Indeed, the model shows there always exists some combination of labor demand-driven movements in occupation-specific labor prices that lead aggregate employment and wages to move in opposite direc-

[1] See, for example, Erosa, Fuster, Kambourov, and Rogerson (2025).

tions so long as different workers earn different wages in different jobs. Therefore, the intuition that demand shocks cause prices and quantities to move in the same direction in frictionless economies is broken by aggregation, evoking Stockman (1983). As skills become more specific, workers are less likely to reallocate out of declining occupations into stable ones and instead leave the employed pool. Meanwhile, the larger the gap between the highest and lowest paid workers – that is, the greater the degree of inequality – the larger the potential for composition effects on the aggregate wage. Therefore, economies with different skill distributions or inequality will feature different responses of measured aggregate employment and wages to a given set of labor demand shocks. Furthermore, changing the composition of labor demand shocks will generate different responses of aggregate employment and wages, since they induce differential labor supply responses of workers with different skills.

Models which ignore worker and job heterogeneity would interpret these variable movements in employment and wages as stemming either from shifts in an aggregate labor supply curve (Beraja, Hurst, and Ospina, 2019), or from workers being off their labor supply curve due to some friction or "labor wedge" – a gap between workers' marginal rate of substitution and their marginal rate of transformation between consumption and leisure (Chari, Kehoe, and McGrattan, 2007; Brinca, Chari, Kehoe, and McGrattan, 2016). Therefore, the final implication of the model is that such "aggregate labor supply shocks" may be microfounded by aggregation of stable labor supply relationships in the face of heterogeneous labor demand shocks. In contrast to many models with other frictions that may cause labor wedge fluctuations – such as nominal rigidities – the model with skill heterogeneity is efficient.

These forces appear quantitatively important for understanding recent aggregate employment and wage movements. I apply the model to study the Great Recession, when real average wages increased by around 2% despite a crash in employment. I first non-parametrically estimate the multidimensional skill distribution by adapting the distributional framework of Bonhomme, Lamadon, and Manresa (2019). By observing the inter-occupation mobility patterns of workers and their wages before and after occupation switches, the econometrician can recover the distribution of types and all the occupation-specific skills for each type of worker and the mass of each type. Then, I shock the estimated model with a sequence of industry-specific revenue productivity shocks taken from the data.

Despite the model's stylized and frictionless nature, it replicates the negative comovement between employment and wages observed during the Great Recession. In the data, real average hourly earnings rose by 1.6% and hours fell by 8.8%, while the model generates wage increases of 2.5% and employment declines of 10%. This result suggests that labor demand shocks can generate negative comovements between employment and wages, even without frictions or exogenous labor supply shocks.

Both skill specificity and heterogeneity in average worker skills are important for this result. Skill specificity is necessary because it constrains the extent to which the decline in demand for construction workers in 2009 reduces the price of labor of other tasks. A model with worker fixed effects but perfectly transferable skills still generates declines in both aggregate employment and wages despite large composition effects.

I then use the model to assess how changes in the skill distribution and the composition of shocks affected the response of aggregate employment and wages during the Great Recession. I estimate that both the degree of skill specificity and the variance of workers' average skills have grown since the 1980s. Therefore, a given set of labor demand shocks will likely generate stronger composition effects and less pro-cyclical real wages going forward. Indeed, the model generates aggregate wage declines of 3% were the Great Recession's labor demand shocks to have occurred in the 1980s. Furthermore, the model suggests that aggregate wages would have fallen by approximately 6% were all sectors subject to the same aggregate shock, suggesting the Great Recession's unusually large impact on sectors which employ many manual laborers contributed to the decoupling of employment and wages in the aggregate.

In sum, this paper's principal contribution is to show three effects of skill heterogeneity. First, heterogeneity changes both own- and cross-price elasticities of labor supply, which affects the extent to which labor demand shocks translate into wage movements in general equilibrium. Second, it implies different combinations of labor demand shocks generate different comovements of aggregate employment and wages, owing in part to composition effects. Third, these different comovements manifest as fluctuations in labor wedge or workers' willingness to work when interpreted in a representative agent framework. These forces appear quantitatively important.

The analysis is subject to a few caveats. First, the model developed here does not feature any labor market frictions or inefficiencies, such in occupational mobility frictions, search, wage rigidities or market power, all of which have been shown to be important elsewhere.[2] This is to elucidate the economics of skill heterogeneity cleanly and to provide an efficient benchmark. The quantitative exercise should thus be interpreted as offering an indication that skill heterogeneity is an important consideration for aggregate labor market dynamics, but not as evidence that other frictions are unimportant. Second, the model features workers who belong to skill types which are fixed over time, and is therefore best-suited to studying labor market dynamics over

[2]See, for example, Grigsby, Hurst, and Yildirmaz (2021); Christiano, Eichenbaum, and Trabandt (2015); Pilossoph (2014) and Berger, Herkenhoff, and Mongey (2022)

the short-to-medium run when the distribution of skills is likely stable.

Literature Review – A seminal paper by Solon, Barsky, and Parker (1994) highlights the important role of composition effects in aggregate wage fluctuations by showing that wages were far more cyclical in a balanced panel of workers which removes compositional shifts in the PSID. This influential paper spawned a number of studies seeking to understand the cyclical selection patterns in the labor market (e.g. Gertler and Trigari (2009); Gertler, Huckfeldt, and Trigari (2020); Mueller (2017)) and their implications for macro aggregates (Chang, 2000; Patterson, 2023).[3]

Particularly related to my paper is Chang and Kim (2007), who argue that a model with incomplete markets and one-dimensional skill heterogeneity generates a low covariance between employment and aggregate productivity through movements in a labor wedge. Relative to this paper, my model features multidimensional skill heterogeneity so that skills can be specific. This generates qualitatively new implications – skill specificity affects the elasticity of labor supply, changes the nature of composition effects, and implies that the composition of labor demand shocks will appear as movements in the labor wedge. I additionally offer a method to estimate the skill distribution and suggest that it has changed over time.

Barlevy (2002) and Hagedorn and Manovskii (2013) consider variations of job-search models incorporating one-dimensional idiosyncratic productivities. These models more successfully replicate comovements of wages and employment over the business cycle through the combination of "cleansing" (low-match-quality workers are laid off in downturns) and "sullying" (low-quality jobs are created in a recession) effects. This mechanism is related to the composition effects studied here, but the focus is different: I elucidate the role of ex ante skill heterogeneity and emphasize the interplay of skill specificity and diffuse labor demand shocks for aggregate labor markets.

Some papers study the implications of occupational reallocation over the business cycle. Carrillo-Tudela and Visschers (2023) and Wiczer (2015) develop models with occupation-specific skills and frictional labor markets and argue that occupational mobility shapes the cyclicality of aggregate unemployment and unemployment durations. Carrillo-Tudela and Wiczer (2022) use a related model to argue that the cyclicality of earnings risk mostly arises from job- and occupation-switchers. Meanwhile, Şahin, Song, Topa, and Violante (2014) argue that mismatch between searchers and vacancy postings account for around one-third of Great Recession unemployment, and

[3] Bils (1985) argues that job-switchers' wages are very procyclical, while Grigsby et al. (2021) show that the excess cyclicality of job-switchers' wages is largely due to selection in the kinds of people who move over the business cycle and the sorts of jobs to which they move. Cajner, Crane, Decker, Grigsby, Hamins-Puertolas, Hurst, Kurz, and Yildirmaz (2020) show a large role for composition effects during the Pandemic recession.

Guvenen, Kuruscu, Tanaka, and Wiczer (2020) show that multidimensional skill mismatch appears to depress individual workers' wage growth. Baley, Figueiredo, and Ulbricht (2022) argue that skill mismatch grows in recessions. In contrast to these papers, my paper focuses on the comovement of aggregate employment and wages, and elucidates the role of the skill distribution in shaping labor supply elasticities. I show that skill heterogeneity implies that different combinations of labor demand shocks will generate different aggregate comovements between employment and wages, and thus manifest as a representative agent labor wedge.

A separate literature seeks to evaluate the importance of sectoral shocks to aggregate fluctuations. Lilien (1982) argues for an important role for sectoral shocks in aggregate fluctuations, while Abraham and Katz (1986) argue the opposite. Quah and Sargent (1993) use time series techniques to argue that sector-specific employment movements had common comovements for the period 1948-89. Forni and Reichlin (1998) use structural VARs to show that sector-specific shocks explain a large share of total output variance, but mainly generate high frequency dynamics. More recently, Pilossoph (2014) points out that since gross flows between sectors do not equal net flows between sectors, many canonical models overstate the importance of sectoral shocks for unemployment. Chodorow-Reich and Wieland (2019) show that reallocation across sectors has a large impact on unemployment during recessions, but little effect in expansions, and explain this in a model with sector-level downward nominal wage rigidity. Importantly, Garin, Pries, and Sims (2018) employ factor analysis on industrial production tables to argue that the importance of sectoral shocks has grown over time, which is confirmed by Foerster, Hornstein, Sarte, and Watson (2019).

This paper contributes to this literature in two ways. First, it highlights a new channel – imperfectly transferable human capital – for sector-specific shocks to translate into aggregate prices and quantities, and argues for its importance. Second, it suggests a reason as to why the role of sectoral shocks may have changed over time – whereas in the past, a declining sector may have had easily transferable skills to a growing sector, this may no longer be the case.

The observed acyclicality of real wages implies that large employment fluctuations must be rationalized in representative agent models either through a labor supply elasticity that is far larger than that implied by micro estimates,[4] shifts in preferences over labor supply (Hall, 1997), or as a wedge between a worker's optimal and realized labor supply decision (Chari et al., 2007). Indeed, Brinca et al. (2016) show that this "labor wedge" accounts for a large share of fluctuations during the Great Recession.

[4] An alternative approach to rationalizing the large estimated elasticity of aggregate labor supply is to appeal to differences between extensive and intensive margin elasticities (Rogerson and Wallenius, 2009; Chang, Kim, Kwon, and Rogerson, 2018).

There are many interpretations of this wedge, including changes in home production technology (Benhabib and Rogerson, 1991), or shifts in government spending (Christiano and Eichenbaum, 1992). Studies commonly suppose that workers are rationed due to some frictions in the labor market such as nominal wage rigidities (Christiano, Eichenbaum, and Evans, 2005; Christiano, Motto, and Rostagno, 2014; Hall, 2005), or search frictions.[5] Alternatively, Beraja et al. (2019) argue that standard models require large exogenous shocks to aggregate labor supply to rationalize the observed fluctuations in aggregate employment and wages. This paper provides a tractable microfoundation for these aggregate labor supply shocks or movements in the measured representative agent labor wedge through the differential response of heterogeneous workers in the face of diffuse labor demand shocks.

Skill heterogeneity has also been shown to matter for long-run phenomena. Gathmann and Schönberg (2010) study the transferability of human capital and find that job mobility is influenced by the similarity of tasks employed. Heckman and Sedlacek (1985); Böhm, von Gaudecker, and Schran (2024) and Roys and Taber (2022) develop Roy models featuring skill heterogeneity, of which the model in this paper is reminiscent. Heckman and Sedlacek (1985) shows that skill heterogeneity is important for the distribution of wages in the U.S. Böhm et al. (2024) argue that changes in skill prices have led to a surge in wage inequality. Roys and Taber (2022) use their model to study changes in the wage structure of low-skill men, finding that manual skills remain important for these workers. Erosa et al. (2025) argues that changes in the variance of within-occupation productivity, rather than skill prices, are behind increases in earnings inequality. Burstein, Morales, and Vogel (2019) use a Roy model to argue that computerization has accounted for much of the rise in the skill premium. This paper also employs a Roy model, but focuses on the implications of skill heterogeneity and industry shocks for shorter-run aggregate employment-wage comovements.

2. MODEL

2.1 Environment

Workers and Labor Supply – The economy consists of K distinct job types indexed by k and a unit mass of infinitely-lived workers who belong to one of J skill types indexed by j. The mass of workers who are type j is given by m_j. I think of job types as reflecting occupations, and so use the words “job” and “occupation” interchangeably.

[5] See Rogerson, Shimer, and Wright (2005) for a survey.

Workers have preferences over a numeraire consumption good c_t summarized by the weakly concave utility function $u(c_t)$ and discount future periods using a constant discount factor β. Workers also receive non-pecuniary benefits from their jobs, described below. Letting Ξ_{it} denote non-pecuniary benefits in period t, worker i's payoff is

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_{it}) + \Xi_{it} \right] \tag{1}$$

The J types of worker differ according to their skill in each occupation k. A worker of type j can supply γ_{jk} efficiency units of labor to occupation k. For notational simplicity, let Γ denote the matrix whose (j,k) element is γ_{jk}. There is one price per unit of human capital in occupation k in period t, which is denoted w_{kt}. A worker of type j therefore earns $\omega_{jkt} \equiv \gamma_{jk} w_{kt}$ if they work in occupation k. Worker skills are fixed over time; thus this model best reflects short-to-medium-run dynamics. Workers of type j receive non-labor income T_{jt} and consume their income each period. Let $c_{jkt} \equiv \gamma_{jk} w_{kt} + T_{jt}$ denote the consumption of a type j worker employed in occupation k in period t.

Workers' labor supply is indivisible and their only decision is their occupation choice, which they freely make each period. Each occupation provides some average non-pecuniary benefits ξ_k to workers which are fixed through time.[6] Workers may additionally choose to be non-employed, in which case they receive no wages but earn an inactivity benefit, which is normalized to 0. Given this normalization, the non-pecuniary benefits ξ_k may be thought of as the negative of non-employment benefits. In addition, each worker receives an idiosyncratic preference shock ζ_{ik} for each occupation. As a result, the occupation chosen by worker i in period t is given by

$$k_t(i) = \underset{k \in \{0,1,\dots,K\}}{\operatorname{argmax}} \left\{ u(\gamma_{j(i)k} w_{kt} + T_{jt}) + \xi_k + \zeta_{ikt} \right\} \tag{2}$$

where $k = 0$ represents the non-employed state. The preference shocks ζ_{ikt} are assumed to be distributed according to an i.i.d. mean zero type 1 extreme value distribution with standard deviation ν_j. The standard deviation ν_j determines the weight that workers place on pecuniary versus non-pecuniary benefits of working and therefore is a key determinant of the elasticity of labor supply. The expectation in household preferences (1) are taken over the realizations of idiosyncratic preferences ζ_{ikt} and skill prices w_{kt}, while $\Xi_{it} \equiv \xi_{k_t(i)} + \zeta_{ik_t(i)t}$.[7]

Final Goods Producers – There is a representative competitive firm which produces

[6] Sorkin (2018) shows that approximately 40% of workers receive a wage cut when switching employers and estimates that non-pecuniary benefits account for a large share of the variance of earnings.

[7] If utility is linear, one can think of the occupation choice decision as representing the decision a worker makes each hour. Under this conception, there would be a unit mass of *hours* and m_j would represent the share of available hours that are accounted for by workers of type j.

numeraire using the output from a finite set of S intermediate goods sectors as inputs to a constant elasticity of substitution (CES) production function. That is, the output of the final good is given by

$$Y_t = \left(\sum_{s=1}^{S} \mu_{st} \hat{y}_{st}^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \tag{3}$$

where $\hat{y}_{st}$ the demand for sector s's output from the final goods producer in period t and μ_{st} is a shifter of demand for sector s. The demand curve for sector s's output is

$$p_{st} = \mu_{st} \left(\frac{Y_t}{\hat{y}_{st}} \right)^{\frac{1}{\eta}} \tag{4}$$

Intermediate Goods Firms – Each intermediate sector s is populated by a representative competitive firm. The firm hires labor in the K occupations to produce output $y_{st} = z_{st} F^{(s)}(l_{s1t}, l_{s2t}, \ldots, l_{sKt})$, where z_{st} denotes the productivity (TFP) of sector s in period t which follow a common knowledge stochastic process, l_{skt} is the quantity of occupation k services hired by sector s in t, and $F^{(s)}(\cdot)$ is a sector-specific production function which is increasing and concave in each of its arguments. Firms take as given all prices. Therefore, intermediate firms in period t solve

$$\pi_{st} = \max_{\{l_{s1t}, l_{s2t}, \ldots, l_{sKt}\}} \quad p_{st} z_{st} F^{(s)}(l_{s1t}, l_{s2t}, \ldots, l_{sKt}) - \sum_{k=1}^{K} w_{kt} l_{skt} \tag{5}$$

Total profits in the economy is the sum of all sectors' profits: $\Pi_t := \sum_{s=1}^{S} \pi_{st}$. Profits are redistributed equally between all workers, so that $T_{jt} = \Pi_t$ for all j. Conditional on the choice of occupation, workers are indifferent between sectors.

Note that μ_{st} and z_{st} affect the marginal revenue product of labor in sector s isomorphically: μ_{st} affects it through prices while z_{st} affects the quantity marginal product of labor. Therefore, I normalize $\mu_{st} = 1$ for all sectors and interpret shocks to industries' revenue TFP as reflecting a combination of productivity and demand shocks.

2.2 Aggregation and Equilibrium

Let $E_{jk}(\mathbf{w}_t)$ denote the probability that a worker of type j chooses occupation k given the occupation price vector $\mathbf{w}_t \equiv \{w_{kt}\}_{k=1}^{K}$. Movements in $\mathbf{w}_t$ induce workers of different types to reallocate across occupations and to non-employment. The distributional assumptions on ζ_{ikt} are standard in the discrete choice literature and generate a tractable form for the occupational choice probabilities of workers:

$$E_{jk}(\mathbf{w}_t) = \frac{\exp\left(\frac{u(\gamma_{jk}w_{kt}+T_{jt})+\xi_k}{\nu_j}\right)}{\sum_{k'=0}^{K}\exp\left(\frac{u(\gamma_{jk'}w_{k't}+T_{jt})+\xi_{k'}}{\nu_j}\right)} \tag{6}$$

Type j employment rates are $\tilde{E}_j = \sum_k E_{jk}$. Aggregating workers' individual decision problems yields occupation-level labor supply curves. Employment in occupation k is

$$E_k(\mathbf{w}_t) = \sum_{j=1}^{J} m_j E_{jk}(\mathbf{w}_t) \tag{7}$$

This $E_k(\mathbf{w}_t)$ schedule returns the measure of workers in each occupation for any set of labor prices. This employment is measurable in the data. Because worker types differ in their skill endowments, the labor supply curve in each occupation is given by the human-capital-weighted employment in each occupation:

$$L_k(\mathbf{w}_t) = \sum_{j=1}^{J} m_j E_{jk}(\mathbf{w}_t)\gamma_{jk} \tag{8}$$

Summing over each occupation yields the aggregate employment and labor units supplied, which depend on the vector of occupation prices $\mathbf{w}_t$:

$$\bar{E}(\mathbf{w}_t) = \sum_{k=1}^{K} E_k(\mathbf{w}_t), \qquad L(\mathbf{w}_t) = \sum_{k=1}^{K} L_k(\mathbf{w}_t) \tag{9}$$

When $\mathbf{w}_t$ moves, it may induce separation between $E_k(\mathbf{w}_t)$ and $L_k(\mathbf{w}_t)$ depending on which workers respond to the labor price changes. This may change the mean human capital of employed workers. Define the mean human capital units supplied by workers employed in a given occupation k to be the ratio of labor units to employment:

$$\bar{\gamma}_k(\mathbf{w}_t) \equiv \frac{L_k(\mathbf{w}_t)}{E_k(\mathbf{w}_t)} \tag{10}$$

Since workers are remunerated according to their human capital levels, movements in $\bar{\gamma}_k(\mathbf{w}_t)$ can shift mean earnings while leaving employment unaffected. As discussed below, this selection force can induce all manner of relationships between aggregate employment and measured wages. One may express the measured aggregate wage as the employment-share-weighted average occupational wage, or the weighted average of worker type mean wages:

$$\bar{\omega}(\mathbf{w}_t) = \sum_{k=1}^{K} w_{kt}\bar{\gamma}_k(\mathbf{w}_t)\left(\frac{E_k(\mathbf{w}_t)}{E(\mathbf{w}_t)}\right) = \sum_j \left(\frac{m_j\tilde{E}_j(\mathbf{w}_t)}{\bar{E}(\mathbf{w}_t)}\right)\tilde{\omega}_j(\mathbf{w}_t) \tag{11}$$

where the symbol ω represents take-home pay and $\tilde{\omega}_j$ is the mean wage of employed type j workers. Note that take-home pay, which increases with worker skill, is distinct from the skill price w. The mean wage of type j workers is given by

$$\tilde{\omega}_j(\mathbf{w}_t) = \sum_{k=1}^{K} \frac{E_{jk}(\mathbf{w}_t)}{\tilde{E}_j(\mathbf{w}_t)}\gamma_{jk}w_{kt} \tag{12}$$

This equation shows that skill influences workers' wage in two ways. The first is direct: workers with high γ_{jk} earn higher wages from working in occupation k by virtue of being more productive in that occupation. In addition, there is an allocation effect operating through $E_{jk}(\mathbf{w}_t)$: workers with higher γ_{jk} *relative* to $\gamma_{jk'}$ are more likely to work in occupation k.

Equilibrium Definition – A competitive equilibrium is a sequence of set of output prices $\mathbf{p}_t = \{p_{st}\}_{s=1}^{S}$, occupation prices $\mathbf{w}_t = \{w_{kt}\}_{k=1}^{K}$, and decision rules $\{E_{jk}(\mathbf{w}_t)\}_{j,k}, \mathbf{l} = \{l_{sk}(\mathbf{w}_t|\mathbf{p}_t, z_{st})\}_{s,k}, \hat{\mathbf{y}} = \{\hat{y}_s(\mathbf{p}_t)\}_s$ such that, in every period t, given sectoral productivities $\mathbf{z}_t = \{z_{1t}, \dots, z_{St}\}$,

1. Workers' occupation choice decisions $E_{jk}(\mathbf{w}_t)$ solve (2) and labor supply is (8),
2. Occupation demand functions $\{l_{sk}(\mathbf{w}_t|\mathbf{p}_t, \mathbf{z}_t)\}_{s,k}$ solve the intermediate sectors' firm's problem (5),
3. The demand for each sector's output from the final goods producer $\hat{y}_s(\mathbf{p}_t)$ is equal to the supply of that sector's output $z_{st}F^{(s)}(\mathbf{l}_s(\mathbf{w}_t|\mathbf{p}_t, \mathbf{z}_t))$,
4. The final goods market clears; that is, aggregate output equals total income: $Y_t = C_t = \bar{\omega}_t\bar{E}_t + \Pi_t$, where Π_t is consistent with the firm problem (5), while $\bar{\omega}_t$ and $\bar{E}_t$ are consistent with aggregation (9) and (11),
5. Occupation-specific labor markets clear

$$L_k(\mathbf{w}_t) = \sum_{s=1}^{S} l_{sk}(\mathbf{w}_t|\mathbf{p}_t, z_{s_t}) \qquad \text{for all } k, t \tag{13}$$

The approach to characterizing equilibrium is detailed in Appendix B.[8]

[8] Since no agent makes dynamic decisions, equilibrium is insensitive to the specification of expectations.

2.3 Discussion

The matrix Γ permits rich heterogeneity in the skill distribution, both vertically and horizontally. The level of γ_{jk} determines the absolute advantage of type j workers in performing occupation k. Workers with a high mean γ_{jk} are generally skilled. Meanwhile, the ratio of γ_{jk} to $\gamma_{jk'}$ measures the comparative advantage of type j workers in k relative to k'. Workers with less variance in their skill vector will generally have transferable skills, as the return to working is similar across all occupations.

This structure nests common paradigms for skill heterogeneity. If $\gamma_{jk} = \gamma_k$ for all j, then every worker type has the same skill vector. This is akin to a representative worker framework. Alternatively, if $\gamma_{jk} = \gamma_j$ for all k, then workers are vertically differentiated – although some workers are high-skill (have high γ_j), no worker has comparative advantage in any particular occupation. This is the worker fixed effect specification of, for example, Abowd, Margolis, and Kramarz (1999) or Chang and Kim (2007). Finally, workers have perfectly specific human capital if Γ is a diagonal matrix: they are able to supply productive labor to their occupation of skill, but not to any other occupation.

Γ is a useful reduced form for a much larger array of traits that individuals may possess (Welch, 1969). Lazear (2009) argued that specific human capital may be thought of in a "skill-weights" framework. In Lazear's set up, jobs are characterized by the weights that they place on a discrete mix of skills in determining worker output. Workers with high ability levels in the skills required by a particular job may be thought to have job-specific human capital. Following this idea, recent papers have developed measures of skill remoteness between occupations using surveys of the skills required to perform the tasks of an occupation, such as O*NET in the US (Guvenen et al., 2020) or the German Qualification and Career Survey (QCS) (Gathmann and Schönberg, 2010; Geel and Backes-Gellner, 2009). A consistent finding of this literature is that workers who move to more remote occupations realize larger wage declines (Poletaev and Robinson, 2008; Nedelkoska, Neffke, and Wiederhold, 2015), while Cortes and Gallipoli (2018) estimates a gravity equation of worker flows to claim that task-independent occupation-specific factors account for most of the variation in transition costs between occupations.[9] The Γ matrix captures many of these features, but has the benefit of providing cardinal units for skills. Estimating the Γ matrix, the mass of each type and the parameters determining the non-pecuniary benefits of job choice therefore permits a detailed estimation of flexible skill substitution patterns.

[9] Macaluso (2019) argues that this skill remoteness can have local labor market effects in a search model of mismatch unemployment.

The model is intentionally parsimonious to focus on the role of skill and shock heterogeneity for aggregate labor markets. In particular, it abstracts from three important dynamic considerations which are worth further discussion. The first is savings. Savings change the mapping between earnings and consumption, which in turn changes workers' valuations of labor income.[10] Assuming workers are hand-to-mouth implements a strong form of market incompleteness. However, my conclusions hold even if workers are risk-neutral, in which case the time path of consumption does not affect worker welfare and the hand-to-mouth assumption is of little consequence.

The second dynamic consideration from which I abstract are fixed costs of occupation adjustment, such as foregone professional networks.[11] Such fixed costs slow reallocation between occupation and make the response of aggregate employment and wages to shocks history-dependent. While interesting, this history dependence complicates the below analysis substantially. I therefore abstract from them for clarity. Finally, I abstract from human capital accumulation; thus my results reflect the effect of labor demand shocks given a pre-existing skill distribution, and are best-suited to considering shorter-run phenomena.[12]

Since the model is frictionless, it can be solved period-by-period. The model is thus best-suited to understanding net reallocation of workers between occupations, rather than gross flows.[13] Nevertheless, net reallocation is what determines aggregate employment and wages, as the accounting identities (9) and (11) make clear.

These simplifications carry two significant benefits. First, they grant great tractability and transparency, which allows clean derivations of the effects of labor demand shocks on skill prices and aggregate employment and wages as discussed below. Second, they permit flexible estimation of the Γ matrix which, as we will soon see, is crucial for the equilibrium behavior of measured employment and wages.

3. THE ROLE OF SKILL HETEROGENEITY

This section examines how skill heterogeneity shapes the response of aggregate labor markets to labor demand shocks. First, I study how skill heterogeneity affects the shape of labor supply curves, and thus the response of occupational prices to labor demand shocks in general equilibrium. I then show that the distribution of skill het-

[10] Although imperfect, one can loosely think of the utility function $u(\gamma_{jk} w_{kt} + T_{jt})$ as representing the indirect utility of a worker receiving income $\gamma_{jk} w_{kt} + T_{jt}$.

[11] See Traiberman (2019) and Dvorkin (2014) for models with dynamic occupation choice.

[12] Wasmer (2006) and Lazear (2009) develop models endogenizing general and specific skill.

[13] See Pilossoph (2014) for a discussion of gross flows in a model with industry shocks.

erogeneity shifts the measured movement of aggregate employment and wages in response to given movements in skill prices. Finally, I show how skill heterogeneity offers a frictionless microfoundation for labor wedges or labor supply shocks in representative agent economies. I suppress dependence on time t throughout this section.

3.1 Labor Supply Curves and General Equilibrium Spillovers

Consider a shock which shifts the TFP in a collection of sectors, $\boldsymbol{d\ln z} \equiv (d\ln z_1, \ldots, d\ln z_S)$. Such a shock shifts the labor demand for the occupations principally employed in the shocked sectors. One can derive how such a shock would affect skill prices and the allocation of workers by differentiating the labor market clearing conditions for each occupation (13) with respect to shocks to each sector s. Doing so for each market clearing condition, stacking and re-arranging yields the following expression for the response of skill prices:

$$\boldsymbol{d\ln w} = \left(\boldsymbol{\varepsilon^S} - \boldsymbol{\varepsilon^D}\right)^{-1}\left(\frac{\boldsymbol{\partial\ln L^D}}{\boldsymbol{\partial\ln z}} + \Theta\frac{\boldsymbol{d\ln p}}{\boldsymbol{d\ln z}}\right)\boldsymbol{d\ln z} \tag{14}$$

where $\boldsymbol{d\ln w} \equiv (d\ln w_1, \ldots, d\ln w_K)'$ is a $K\times 1$ vector of skill price responses, $\boldsymbol{\varepsilon^S}$ is a $K\times K$ matrix whose (k,k') element is the elasticity of labor supply in k to a change in skill price in k', $\varepsilon^S_{k,k'} = \frac{\partial\ln L_k}{\partial\ln w_{k'}}$, and $\boldsymbol{\varepsilon^D}$ is a $K\times K$ matrix whose (k,k') element is the elasticity of labor demand in k to a change in skill price in k', $\varepsilon^D_{k,k'} = \frac{\partial\ln\sum_{s'} l_{s'k}}{\partial\ln w_{k'}}$. Finally, $\frac{\boldsymbol{\partial\ln L^D}}{\boldsymbol{\partial\ln z}}$ is a $K\times S$ matrix of direct effects on labor demand for each occupation from the shock to sector s, whose (k,s) entry is $\frac{\partial\ln\sum_{s'} l_{s'k}}{\partial\ln z_s}$, $\frac{\boldsymbol{d\ln p}}{\boldsymbol{d\ln z}}$ is a $S\times S$ matrix whose (s,s') entry is the elasticity of equilibrium output prices of sector s to $z_{s'}$, and Θ is a $K\times S$ matrix whose (k,s) element is the elasticity of labor demand in k to a change in the price of sector s's output: $\Theta_{k,s} = \frac{\partial\ln\sum_{s'} l_{s'k}}{\partial\ln p_s}$.

Equation (14) shows that the first-order effect of an arbitrary set of productivity shocks depends on the matrix of labor supply and demand responses to the shock. In the case with one worker type, one occupation and one sector, this expression simply states that the movement in prices is equal to the size of the outward shift in demand divided by the sum of the absolute values of supply and demand elasticities. As is standard, the larger the total supply and demand elasticities, the smaller the movement in prices. Equation (14) is similar, but accounts for cross-occupation spillovers in both labor supply and labor demand that arise as firms substitute away from high-priced occupations while workers substitute towards such occupations. The total outward shift in demand is given by both the direct effect from increased productivity, and the effect that arises from shifts in equilibrium sectoral output prices, for which the elasticity of substitution across sectors η is key.

The distribution of skills is crucial for the nature of labor supply spillovers captured by the matrix $\boldsymbol{\varepsilon}^S$. Appendix B shows that the (k, k') element of $\boldsymbol{\varepsilon}^S$ can be expressed as:

$$\boldsymbol{\varepsilon}^S(k, k') = \frac{1}{L_k} \sum_j m_j \left(\frac{u'(c_{jk})}{\nu_j} \right) \gamma_{jk} \gamma_{jk'} w_{k'} \left(\mathbf{1}\left\{k' = k\right\} - E_{jk'} \right) E_{jk} \tag{15}$$

Labor supply responses are partially governed by the standard deviation of the idiosyncratic preference shocks ν_j and workers' marginal valuation of income $u'(c_{jk})$. Note, however, that skills play an important role as well. Labor supply spillovers are larger between two sectors which have a higher covariance of skills, as can be seen by the presence of the $\gamma_{jk}\gamma_{jk'}$ term. That is, more transferable skills implies larger cross-occupation labor supply spillovers. Indeed, in the limit where workers have perfectly specific skills – when every worker type j has $\gamma_{jk} = 0$ for all but one occupation – cross-occupation labor supply elasticities are locally zero.

Intuitively, when the price of labor falls in one occupation, workers seek employment elsewhere. This shifts out the labor supply curve in occupations unaffected by the initial shock. As a concrete example, if demand for construction workers falls, workers previously employed in construction may seek employment in manufacturing, thereby exerting downward pressure on the price of labor in manufacturing, but may be less likely to become nurses,[14] and so exert less downward pressure on the price of labor in nursing.

In addition to affecting cross-occupation labor supply spillovers, skill heterogeneity additionally impacts own-price labor supply elasticities. Specializing (15) to the case where $k = k'$, we have that own-price labor supply responses are

$$\boldsymbol{\varepsilon}^S(k, k) = \frac{1}{L_k} \sum_j m_j \left(\frac{u'(c_{jk})}{\nu_j} \right) \gamma_{jk}^2 w_k \left(1 - E_{jk}\right) E_{jk}. \tag{16}$$

Observe that γ_{jk} enters quadratically into this expression. Jensen's inequality implies, therefore, that own-price labor supply elasticities will tend to be larger if occupational skills are more concentrated among a few workers. Intuitively, this is because the labor supply for that occupation is very sensitive to the decisions of those few workers. This is mediated by the $(1 - E_{jk})E_{jk}$ term: elasticities will be small if all workers are already employed in that occupation (as there are few remaining workers to attract) or

[14] Appendix A.1 presents reduced form evidence that such a force exists using an exogenous decline in the demand for mining labor between 2014 and 2016. Beaudry, Green, and Sand (2012) additionally provides evidence of a similar force operating through outside options and bargaining.

if no workers are employed in that occupation (as that implies the occupation is very unattractive relative to alternatives).

To build intuition for these labor supply functions, let us consider a simplified version of the model in which there are only two types of workers who occupy an equal share of the population and two types of job: $K = J = 2$. For simplicity, suppose further that workers share a common fundamental labor supply elasticity $\nu_j^{-1} = 2$ and have linear utility $u(c) = c$ so that the $m_j u'(c_{jk})/\nu_j$ terms drop out of equations (15) and (16). Consider three specifications for the skill matrix Γ, defined as follows

$$\Gamma^{(RA)} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad \Gamma^{(CA)} = \begin{pmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{pmatrix} \qquad \Gamma^{(AA)} = \begin{pmatrix} 1.5 & 1.5 \\ 0.5 & 0.5 \end{pmatrix} \tag{17}$$

$\Gamma^{(RA)}$ is a representative agent skill matrix in which all worker types have one unit of human capital that can be equally well applied to both occupations. In contrast, $\Gamma^{(CA)}$ is a comparative advantage skill matrix which admits some skill specificity – type 1 workers can supply 1.5 units of human capital to occupation 1 but only 0.5 units to occupation 2, and vice versa for type 2 workers. Finally, $\Gamma^{(AA)}$ is an absolute advantage or "worker fixed effect" skill matrix: type 1 workers have 1.5 units of human capital and type 2 workers have 0.5 units of human capital, but that human capital is equally well applied in each occupation.

One can plot the labor supply curve in occupation 1 given a price of labor in occupation 2. Panel A of Figure 1 plots this curve for the three skill matrices assuming $w_2 = 1$. The figure shows upward-sloping labor supply curves for each of our skill matrices. The labor supply curve under a representative agent skill matrix resembles that of the absolute advantage skill distribution. This is because there are two offsetting effects. First is the quadratic in γ_j: high-type workers respond strongly to the increase in the price of labor in occupation 1. This pushes towards a higher own-price elasticity. However, low-type workers have low employment E_{j1} and low skills and so respond very little. These two forces roughly offset for much of the range of w_1.

By contrast, the economy with a comparative advantage skill distribution exhibits a different labor supply curve due to skill specificity. Type 1 workers are very responsive to increases in w_1. As a result, for low levels in the price of labor, the labor supply curve is very elastic. Eventually, however, almost all type 1 workers are employed in occupation 1, at which point the labor supply curve becomes very inelastic, as type 2 workers do not respond much to increases in w_1. The degree of skill specificity therefore affects the behavior of skill prices in response to labor demand shocks.

Panel A assumes that w_2 were equal to 1. Suppose now that the price of labor in oc-

FIGURE 1: General Equilibrium Labor Supply Spillovers

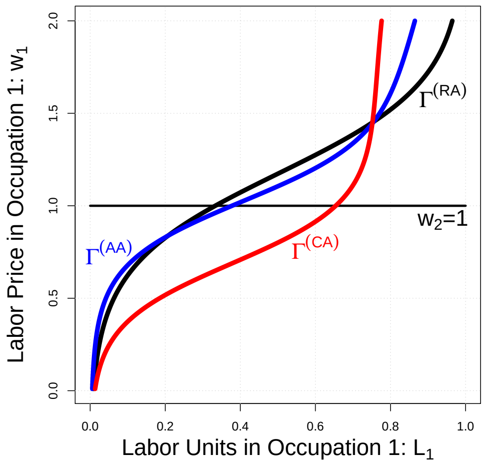

PANEL A: LABOR SUPPLY CURVES IN OCC. 1

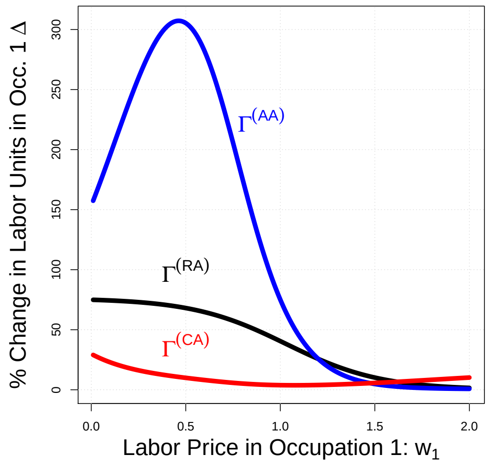

PANEL B: OUTWARD SHIFT IN OCC. 1 LABOR SUPPLY CURVE FROM DECLINE IN w_2 (Δ)

Notes: Figure shows occupation 1's labor supply given exogenously specified prices of labor. Panel A plots the labor supply curve in occupation 1 if w_2 is fixed at 1. Panel B plots the percentage horizontal shift in occupation 1's labor supply curve when w_2 falls to 0.5, as described in equation (18). The black line has a representative agent skill matrix, the blue line has a worker fixed effect skill matrix, and the red line has a comparative advantage skill matrix, as defined in equation (17).

cupation 2 exogenously moved to $w_2 = 0.5$. This simulates a large negative demand shock to occupation 2. Because of this decline, the relative value of working in occupation 1 increases, which leads to an outward shift of occupation 1's labor supply curve. One can quantify the magnitude of this shift by measuring the percentage horizontal movement in the labor supply curve for every level of w_1:

$$\Delta(w_1) = \frac{L_1(w_1|w_2 = 1) - L_1(w_1|w_2 = 0.5)}{L_1(w_1|w_2 = 1)} \tag{18}$$

This function $\Delta(w_1)$ is plotted in Panel B of Figure 1 for the three skill matrices. The more specific are skills, the less impact will a shock to the price of labor in occupation 2 have on the labor supply curve of occupation 1, as suggested by equation (15). This is captured by the fact that the red curve representing the comparative advantage skill distribution is substantially below the black and blue curves, which both have perfectly transferable skills for workers between the two occupations. The blue curve representing the absolute advantage skill matrix has the largest shift for low levels of w_1. This is because the workers who move from occupation 2 to occupation 1 are principally high-type workers. Indeed, equation (15) specialized to the worker fixed effect specification is features the expectation of γ_j^2; this convexity in γ_j gives rise to the larger cross-price elasticities under the absolute advantage skill matrix than in the representative agent skill matrix because of Jensen's inequality.

This section shows how skill heterogeneity affects the equilibrium response of skill prices to sectoral labor demand shocks. Skill specificity mutes cross-occupation labor supply spillovers, while heterogeneity in skill levels pushes towards more elastic labor supply curves, as small movements of the most productive workers have large effects on total labor supply to each occupation. Next, I study how the reallocations induced by these equilibrium skill price movements manifest in measured aggregate employment and wages.

3.2 Aggregate Employment and Wage Responses to Shocks

Differentiating equation (9), one can show that the first-order response of aggregate employment to an arbitrary set of shocks $\mathbf{z}$ (see Appendix B) may be written:

$$\frac{d \ln \bar{E}}{d \ln \mathbf{z}} = \sum_j \underbrace{\left(\frac{m_j \tilde{E}_j}{\bar{E}}\right)}_{\substack{\text{Employment}\\\text{Share of } j}} \underbrace{(1 - \tilde{E}_j)}_{\substack{\text{Pool of}\\\text{Workers who}\\\text{Respond}}} \underbrace{\tilde{\omega}_j}_{\substack{\text{Income}\\\text{Effect}}} \sum_k \underbrace{\left(\frac{u'(c_{jk})}{\nu_j}\right)}_{\substack{\text{LS}\\\text{Elasticity}}} \underbrace{\left(\frac{\omega_{jk} E_{jk}}{\tilde{\omega}_j \tilde{E}_j}\right)}_{\substack{\text{Share of } k\\\text{in } j\text{'s earnings}}} \underbrace{\frac{d \ln w_k}{d \ln \mathbf{z}}}_{\substack{\text{Occupational}\\\text{Price Movement}}} \tag{19}$$

Shocks can affect employment through their impact on the price of labor in each of the occupations w_k. The employment of a worker type j responds more to shocks which influence the price of labor in occupations in which j earns a large share of their income, which tend to be occupations in which j has skills. Likewise, employment responds more the more labor supply responds to wages, which is governed by $u'(c_{jk})/\nu_j$. If workers earn high wages on average, then they are further from the outside option of non-employment and thus may be more likely to work in response to shocks. Finally, employment responds more to the shock if there are more non-employed workers that can be drawn into employment.

One can additionally write the first-order response of aggregate wages as

$$\begin{aligned}\frac{d\ln\bar{\omega}}{d\ln \mathbf{z}} = \sum_j \underbrace{\frac{m_j\tilde{\omega}_j\tilde{E}_j}{\bar{\omega}\bar{E}}}_{\substack{\text{Share of } j\\ \text{in Agg. Earnings}}} \sum_k \underbrace{\left(\frac{\omega_{jk}E_{jk}}{\tilde{\omega}_j\tilde{E}_j}\right)}_{\substack{\text{Share of } k\\ \text{in } j\text{'s earnings}}} \Bigg[\overbrace{1}^{\text{Direct}} &+ \overbrace{\underbrace{\left(\frac{u'(c_{jk})}{\nu_j}\right)}_{\substack{\text{LS}\\ \text{Elasticity}}} \underbrace{(\omega_{jk}-\tilde{\omega}_j)}_{\substack{\text{Cross-Occ.}\\ \text{Wage Change}}}}^{\text{Reallocation}} \\ &+ \overbrace{\underbrace{\left(\frac{u'(c_{jk})}{\nu_j}\right)}_{\substack{\text{LS}\\ \text{Elasticity}}} \underbrace{(1-\tilde{E}_j)}_{\substack{\text{Pool of}\\ \text{Workers who}\\ \text{Respond}}} \underbrace{(\tilde{\omega}_j-\bar{\omega})}_{\substack{\text{Relative}\\ \text{Wages}}}}^{\text{Composition}} \Bigg] \frac{d\ln w_k}{d\ln \mathbf{z}}\end{aligned} \tag{20}$$

The shock has three first-order effects on aggregate wages. First is the direct effect: if the shock increases the price of labor for workers' existing jobs, aggregate wages rise *ceteris paribus*. This effect exists even if there is only one type of worker and one type of job. In addition, there is a reallocation effect, which arises if workers move to different jobs in response to the shock. This is more positive if (i) employment in relatively high-wage jobs is elastic to the shock and (ii) the destination constitutes a large share of pre-shock earnings. Finally, there is a composition effect, which is positive if workers who are elastic to the shock have high average wages relative to the aggregate wage. Unsurprisingly, the reallocation and composition effects are larger the more elastic is labor supply (the higher the marginal value of income is relative to the variance of preference shocks). More interestingly, the composition effect is more positive if the occupations being shocked tend to constitute a large share of income for low-wage workers. This is why the interaction of skill (and thus earnings) heterogeneity and the composition of shocks is critical for understanding aggregate wage movements.

To gain intuition for when the composition effect is large or small, consider a version of the model in which $u(c) = \ln c$ and $T_{jt} = 0$. The composition effect then becomes

$$\sum_j \left(\frac{m_j\tilde{E}_j}{\bar{E}}\right)\left(\frac{\tilde{\omega}_j}{\bar{\omega}} - 1\right) \cdot \underbrace{(1-\tilde{E}_j)\left(\frac{1}{\nu_j}\right)\sum_k \left(\frac{E_{jk}}{\tilde{E}_j}\right)\frac{d\ln w_k}{d\ln \mathbf{z}}}_{\frac{d\ln \tilde{E}_j}{d\ln \mathbf{z}}}$$

This expression illustrates five determinants of the composition effect. First, unsurprisingly, the composition effect will be negative if the employment of high wage workers responds by less than low-wage workers. Second, as inequality rises – the gap between low and high wage workers grows on average – the strength of this composition effect will grow. Third, stronger occupational sorting increases the occupational share term $E_{jk}/\tilde{E}_j$. Fourth, more dispersion in employment rates across types leads to greater heterogeneity in employment responses through the $1-\tilde{E}_j$ term. Last, more uneven shocks to labor demand across occupations makes the occupational price term $d\ln w_k/d\ln \mathbf{z}$ larger. As discussed above, the degree of skill specificity is crucial for this occupational price term.

An implication of (20) is that there will always exist some combination of labor price movements in this model that generate aggregate wages and employment to move in opposite directions, so long as workers have different earnings either from one another or in different jobs. Therefore, aggregation breaks the traditional intuition that labor demand shocks cause employment and wages to positively covary (Stockman, 1983). This is further explored in Appendix A.2.

Skill heterogeneity therefore has two effects on measured aggregate employment and wages. The first is that both skill specificity and differences in absolute advantage change the nature of occupation-specific labor supply elasticities relative to a representative agent benchmark, which govern the general equilibrium responses of skill prices to shocks. The second is that, for a given movement in skill prices, economies with different skill distributions will feature different movements in both measured employment and wages owing to the reallocation of heterogeneous workers both between occupations and out of employment. Rises in inequality give scope for larger composition effects. I provide some examples of how the aggregate relationship between employment and wages changes as one changes the skill distribution in Appendix A.3. Next, I explore the implications of these differential shifts in aggregate employment and wages for inference in representative agent economies.

3.3 Representative Agent Labor Supply Shocks and Labor Wedges

Representative agent models have found an important role for either the "labor wedge" (Brinca et al., 2016) or labor supply shocks (Beraja et al., 2019) for explaining recent wage dynamics. The role of such wedges or shocks is inferred by examining shifts in the measured relationship between aggregate employment and wages. The above analysis hints that this relationship may depend both on workers' skills and on the nature of labor demand shocks hitting the economy. This section argues that skill heterogeneity and a changing mix of labor demand shocks offers a frictionless microfoundation of labor wedge movements and labor supply shocks.

The labor wedge represents a representative agent being off their labor supply curve due to frictions such as search or wage rigidity, while labor supply shocks are usually represented by a preference shifter on labor. To fix ideas, consider the following representative agent household problem. The household chooses consumption c_t, savings b_{t+1} and labor E_t to maximize

$$\max_{c_t, E_t, b_{t+1}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t) - \varphi_t h(E_t)\right]$$

subject to budget constraints

$$c_t + b_{t+1} = (1+r)b_t + \omega_t E_t (1 - \tau_{lt}) + T_t$$

This is a simplified version of the benchmark model of Chari et al. (2007) with competitive markets and labor supply shocks specified as in Beraja et al. (2019). Taking the first order conditions of this problem with respect to consumption and labor yields the following optimality condition:

$$\frac{\varphi_t h'(E_t)}{u'(c_t)} = (1 - \tau_{lt})\omega_t \tag{21}$$

This states that the household's marginal rate of substitution between consumption and labor equals their wage times one minus the labor wedge τ_{lt}.[15] Observe that both the labor wedge τ_{lt} and the preference shifter φ_t^{-1} enter this expression isomorphically. Therefore, given assumptions on utility functions and data on wages, employment, and consumption, one can infer the size of the representative labor wedge or labor supply shock. Let $\hat{\tau}_t$ and $\hat{\psi}_t$ be the estimates of the labor wedge or preference shifter that would be inferred by inverting the above equations given data on consumption, employment, aggregate wages and assumptions about the utility function. For instance, in the case with log separable utility between labor and consumption as in Chari et al. (2007) so that $u(c) = \ln c$ and $h(E) = \ln(1-E)$,[16] we can infer the product of labor supply shocks and labor wedge as:

$$\varphi_t^{-1}(1 - \tau_{lt}) = \frac{c_t}{E_t \omega_t} \tag{22}$$

Taking derivatives of this expression with respect to some shock $\mathbf{z}$ and rearranging, we arrive at:

$$\frac{d \ln\left(\varphi_t^{-1}(1 - \tau_{lt})\right)}{d\mathbf{z}} = \frac{d \ln(1 - E_t)}{d\mathbf{z}} - \frac{d \ln \omega_t}{d\mathbf{z}} + \frac{d \ln c_t}{d\mathbf{z}}$$

[15]Traditionally, the labor wedge τ_t is defined as the wedge between households' marginal rate of substitution and the marginal rate of transformation between labor and consumption. Equation (21) holds under competitive markets.

[16]Beraja et al. (2019) also assume log separable utility, but with habit formation in consumption.

That is, the inferred labor wedge or labor supply shock will move in response to $d\mathbf{z}$ if measured labor responds differently in response to measured wage movements than predicted by the model. Interpreting the data through this lens will ascribe any movement in non-employment $1 - E$ that is not matched by either a movement in consumption c or aggregate wages ω to the labor wedge or a labor supply shock – that is, people must either change their value of work or not be on their labor supply curve.

In contrast, this paper's model offers a frictionless rationalization of the data: since, as discussed above, different combinations of occupation-specific shocks can induce different comovements between aggregate wages and employment, changing the composition of labor price shocks will manifest as movements in the labor wedge or exogenous labor supply shocks. To see this, note one can trace out an aggregate employment-wage schedule as the price of occupational labor services changes using the model. Specifically, one can vary the vector of labor prices $\mathbf{w}$ to generate occupation choices according to equation (2). One can then plot the relationship between $\bar{E}(\mathbf{w})$ and $\bar{\omega}(\mathbf{w})$ as implied by equations (9) and (11), respectively. To illustrate how shock composition appears as a measured labor wedge or labor supply shock, I do this under the comparative advantage skill matrix defined above for both an aggregate and idiosyncratic shock to labor prices.[17]

Aggregate Shock – Suppose that price movements are such that $w_1 = w_2 = w$: both occupations have an equal price. The results of this exercise is represented by the red lines in Figure 2. Panel A plots the mean human capital level of employed workers $\bar{\gamma}$ as we vary the price of labor in both occupations w. Panel B plots the implied relationship between aggregate employment and wages.

As Panel A shows, the mean human capital level $\bar{\gamma}$ is increasing in the price per unit of labor w. To see why, consider the case in which $w = 0$. When the price of labor is 0, there is no gain for workers to sort into an occupation in which they have comparative advantage: they will earn nothing regardless of which occupation they choose. Thus there is no sorting. As the price of labor rises, so too do the gains from working in one's highest-skill occupation. Thus workers sort more and more as the price of labor increases, leading to the increasing relationship between $\bar{\gamma}$ and w. As a result, the aggregate relationship between employment and wages is simply and upward-sloping curve. This monotonically increasing human capital level then generates a standard upward-sloping labor supply relationship between employment and wages: as wages rise, more workers enter the employed pool.

[17] For this exercise, the variance of the idiosyncratic preference shocks ν_j is 0.25, while the fixed non-pecuniary benefit is set to -1 across both occupations. Workers are risk neutral for illustration. I show additional results for alternative specifications of the Γ matrix in Appendix A.3.

FIGURE 2: Aggregate Employment-Wage Schedule As Vary Occupation Prices

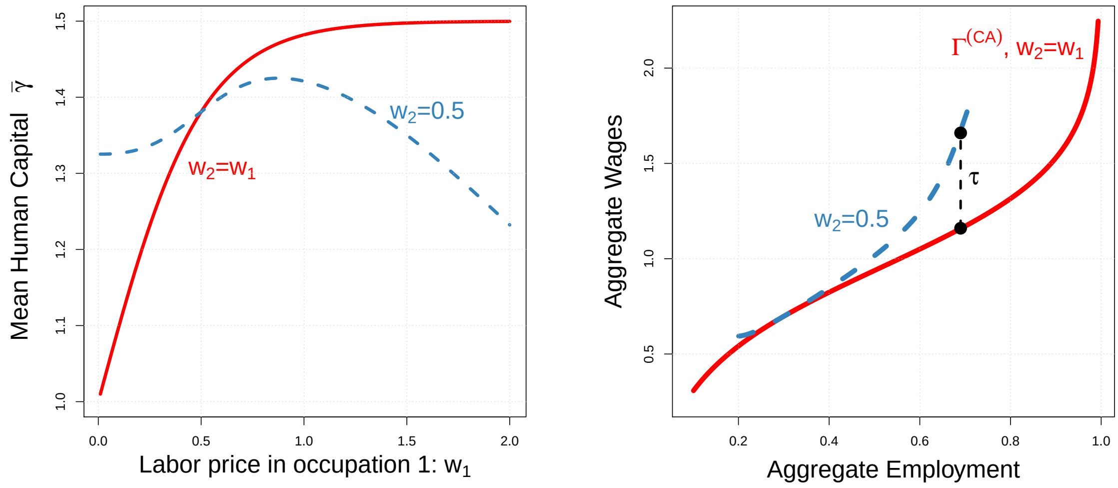

PANEL A: MEAN HUMAN CAPITAL

PANEL B: AGG. EMPLOYMENT-WAGE RELATIONSHIP

Notes: Figure presents the behavior of the aggregate employment and wages after movements in occupational labor prices w_k in a two occupation, two-type labor supply model. Panel A plots the mean human capital of employed workers $\bar{\gamma}$ against the prevailing price of labor, while Panel B plot the implied relationship between aggregate employment and wages. The red solid line is the curve when $w_1 = w_2$, and Γ exhibits comparative advantage. The blue dashed line when w_2 is fixed to 0.5, and Γ exhibits comparative advantage. Γ matrices defined as in equation (17).

Idiosyncratic Shock – Now consider the opposite extreme in which w_2 were fixed at 0.5, while only w_1 varies. This is represented by the blue dashed lines in Figure 2. Fixing the wage in occupation 2 steepens and shifts the aggregate relationship between employment and wages inward. This is because inducing type 2 workers to enter the labor force requires large movements in the price of occupation 1. For high values of w_1, the majority of type 1 workers are employed, and type 2 workers are only marginally responsive to the movements in the price of labor.

If one were to assume a representative agent labor supply curve, one might estimate that curve to be given by, for instance, the red line in Panel B. Realizing a point of data on the blue dashed line would therefore be rationalized either as evidence that workers' supply curve has shifted, or that workers are off their frictionless labor supply curve. This wedge between the realized data and the estimated labor supply curve is depicted on the figure by τ and may be interpreted as a labor wedge, or an exogenous shock to labor supply, when in reality all workers frictionlessly choose their labor from a stable labor supply curve. Therefore, a shifting composition of labor demand shocks offers a frictionless microfoundation the labor wedge or labor supply shocks through its effects on the composition of employment.

3.4 Summary and Discussion

This section shows that skill heterogeneity has three important effects on the equilibrium response of aggregate labor markets to labor demand shocks. First, skill heterogeneity changes the equilibrium response of skill prices to shocks by shifting the equilibrium elasticities of labor supply. As skills become more specific, spillovers across occupations weaken; this is governed by the covariance of γ_{jk} across k. Meanwhile, a Jensen inequality term suggests that an occupation's own-price elasticity of labor supply will be larger if its skill is concentrated within a few individuals; this is governed by the expectation of γ_{jk}^2 taken across j and is thus related to the variance of skills across types.

Second, skill distributions effect the measured response of aggregate employment and wages to shocks. This is because workers with different skills respond differently to any given set of labor price movements. Thus, one will observe large movements in employment if skill prices move in occupations which tend to employ many workers on the margin of being employed or not. Meanwhile, composition effects make it so that measured aggregate wage movements decouple from aggregate employment movements, leading to potentially negative comovements between aggregate employment and wages, even without frictions or labor supply shocks.

Finally, the different movements in aggregate employment and wages that arise from skill heterogeneity and diffuse labor demand shocks will manifest as either an aggregate labor supply shock or labor wedge if interpreted in a representative agent economy. Therefore, periods with different compositions of demand shocks or changes in the skill distribution over time will appear as shifts in the representative agent labor wedge or labor supply curve.

The model presented here abstracts from frictions and dynamics. These considerations are unlikely to change the qualititative results presented here, but likely affect the rate at which these shifts occur. For instance, both real and nominal wage rigidities intuitively operate by changing the extent to which real skill prices change in response to demand shocks. But conditional on the price of skill that workers face, the expressions for the elasticity of supply remain the same.[18] Therefore, the conceptual point regarding the importance of skill specificity for labor supply spillovers and the role of composition effects remain the same. Similarly, if workers face fixed costs to switch occupations which render the occupation choice decision dynamic, this would affect the elasticities of labor supply quantitatively, but the qualitative insights regarding the

[18] If frictions lead to worker rationing, there may further be a wedge between the skill price paid by firms and the shadow price that workers take into account when making labor supply decisions.

role of the skill distribution will continue to hold.

It remains to assess whether these mechanisms are quantitatively important. To that end, I now study the implications of the model for Great Recession labor markets.

4. Great Recession Labor Markets

I first describe my approach to estimating the labor supply side of the model, before describing the calibration of the labor demand side of the model. Then I study the model's implications for measured employment and wages during the Great Recession. Finally, I study whether changes in the skill distribution or composition of labor demand shocks might have mattered for this period. The goal is to argue that these features are important for understanding this period's unusual negative comovement between employment and wages, rather than to rule out the importance of other frictions such as sticky wages.

4.1 Estimation of Labor Supply Parameters

I pursue an annual calibration and assume workers have linear, risk-neutral utility: $u(c) = c$.[19] The identification and estimation of the skill distribution follows closely the distributional framework developed by Bonhomme et al. (2019). Estimation takes a random-effects maximum likelihood approach. Let $k_t(i)$ be the occupation choice of individual i in period t. I assume that individual wages in period t are observed with multiplicative measurement error ϵ_{it}, which has type-occupation-specific parametric distribution $\Psi(\epsilon_{it}|k_t(i), j(i), \theta_\epsilon)$ with unit mean, summarized by the parameter vector θ_ϵ.[20] Log observed wages $\ln \omega_{it}$ are then $\ln \omega_{it} = \ln \gamma_{j(i)k_t(i)} + \ln w_{k_t(i)t} + \ln \epsilon_{it}$.

This model of earnings is similar to that of Bonhomme et al. (2019), with two primary differences. First, while Bonhomme et al. (2019) study firm and worker sorting, I study the sorting of workers to occupations and assign an economic meaning to the wage differences of two workers employed in the same occupation - namely, occupational skill. Second, while Bonhomme et al. (2019) treat the probability that workers switch between each firm type as additional parameters to be estimated, I impose a Roy model of occupational choice so that workers select into jobs. This is consistent with the general equilibrium model presented above, permitting counterfactual ana-

[19] I choose a linear utility function to minimize the impact of abstracting from savings in the model. Note, however, that this choice implies workers' labor supply elasticities are determined only by ν_j.

[20] The disturbance in wages ϵ_{it} may be interpreted as measurement error, or unit mean multiplicative productivity shocks realized after a worker has chosen her occupation.

lyses, and improves the power of my estimation routine by utilizing both wage *and* occupation choice information to estimate the skill vector of each type, rather than just wage information as in Bonhomme et al. (2019).

To fix notation, let $\mathbb{P}_{kk'}(j) = Pr\{k_t(i) = k, k_{t+1}(i) = k'|j(i) = j\}$ denote the probability a type j worker chooses occupations k and k' in subsequent periods. Let the history of realizations of a random variable Z up to period t be given by $Z^t = \{Z_{i1}, \ldots, Z_{it}\}$. Following Bonhomme et al. (2019), I maintain Assumption 1:

Assumption 1. ***Identification Assumptions***

1. *(Mobility Determinants) - The realization of idiosyncratic preference shocks ζ_{ikt} is independent of the history of measurement error in a worker's wage ϵ_i^{t-1}, conditional on the worker's type $j(i)$ and their history of occupation choices $k^{t-1}(i)$.*
2. *(Serial independence) - The realization of period t's measurement error for worker i ϵ_{it} is independent of the history of disturbances ϵ_i^{t-1} and the period t preference shock ζ_{ikt} conditional on the worker's current occupation choice $k_t(i)$ and type $j(i)$.*
3. *(Full Rank) - There exist finite sets of M values for ω_t and ω_{t+1} such that, for all $r \in \{1, \ldots, R\}$, the matrices $A(k_r, \tilde{k}_r)$ and $A(k_{r+1}, \tilde{k}_r)$ have rank J where $A(k, k')$ has $(\hat{\omega}_1, \hat{\omega}_2)$ element*

$$Pr\{\omega_{it} \leq \hat{\omega}_1, \omega_{it+1} \leq \hat{\omega}_2 | k_t(i) = k, k_{t+1}(i) = k', m_{it+1} = 1\}$$

Assumption 1 implies that the parameters of the labor supply model are identified and may be estimated through maximum likelihood, as derived below. It has three pieces.[21] The first is that measurement error in the wage is uncorrelated with workers occupation choice, conditional on their type and history of occupations. This amounts to a timing assumption – a worker may have decided to pursue a career in occupation k based on a forecast of the wage their type can expect to earn in that occupation, but they do not know their precise wage draw. The assumption ensures random mobility *conditional* on workers' type-by-occupation expected earnings in each occupation.

The second piece of Assumption 1 requires that the measurement error is serially independent, conditional on a worker's type and occupation choice.[22]

[21] Bonhomme et al. (2019) requires an additional assumption that any two occupations belong to a connecting cycle for every type of worker. This is similar to the graph connectedness assumption of Abowd et al. (1999) and is implied in my model since ζ_{ikt} are distributed according to a Type 1 Extreme Value distribution.

[22] Bonhomme et al. (2019) show that first-order Markov processes for wages may be accommodated with four-period panel data. The identification problem in two-period panels is that if wages are persistently high for a given individual, one is unable to identify whether that is because they have high γ for their job, or because idiosyncratic wage draws are highly persistent. This distinction likely has little meaning for the purposes of studying composition effects – all that matters is whether high or low wage workers switch jobs.

Finally, the third item in Assumption 1 is a rank condition that will be satisfied if all worker types draw from different distributions for each occupation. That is, it must be the case that worker types are meaningfully different.

To construct the likelihood function, consider the likelihood of observing a single worker i who chooses occupation k in period 1 and k' in period 2, realizing wages ω_{i1} and ω_{i2} in periods 1, and 2, respectively. The estimation seeks to recover the distribution of wages by worker type and occupation. Let the parameters of the model be given by θ, which includes $m_j, \gamma_{jk}w_{kt}, \xi_k$ and the parameters governing the idiosyncratic taste shocks ζ_{ikt} and measurement error θ_ϵ. Let $\psi(\omega|k, j, \theta)$ be the density of idiosyncratic wages implied for a type j worker in occupation k. Unemployed workers' wage density has mass 1 and does not affect the likelihood function. The likelihood of observing this worker may be written as

$$l_i(k, k', \omega_{i1}, \omega_{i2}|\theta) = \sum_{j=1}^{J} m_j \underbrace{\mathbb{P}_{kk'}(j|\theta)\psi(\omega_{i1}|k_1(i) = k, j(i) = j, \theta)\psi(\omega_{i2}|k_2(i) = k', j(i) = j, \theta)}_{l_{ij}}$$

where $\mathbb{P}_{kk'}(j|\theta)$ is the probability that a worker chooses occupation k in period 1 followed by k' in period 2. If we knew the worker's type, the likelihood of observing their occupation choices and wages is given by the probability that their type made their occupation choices, multiplied by the probability of observing the two wage draws. This likelihood is denoted l_{ij}. The multiplication of densities and choice probabilities results from the independence assumption between ζ_{ikt} and the measurement error in wages, conditional on occupation choices and worker type. The overall likelihood of observing that individual, therefore, integrates over the likelihood for each of unobserved type that the worker could be.

Aggregating over all individuals yields the full log-likelihood of the data:

$$\mathcal{L}(\theta) = \sum_i \sum_{k=0}^{K} \sum_{k'=0}^{K} \mathbf{1}\{k_1(i) = k, k_2(i) = k'\} \ln l_i(k, k', \omega_{i1}, \omega_{i2}|\theta) \tag{23}$$

The formal argument for identification closely follows that of Bonhomme et al. (2019).[23] The output is an estimated mean and variance of the wage distributions for each type j worker employed in each occupation k in every period t, a mass of each worker type m_j, and parameters governing non-pecuniary benefits. The mean of the wage distribution of a type j worker employed in occupation k is $\gamma_{jk}w_{kt}$. To separate skill prices w_{kt} and skill levels γ_{jk}, I normalize the average human capital in each occupation to

[23] Assumption 1 may be relaxed at the expense of greater data requirements. Unfortunately, the set of large, representative, panel datasets containing information on occupation and wages is small, requiring the use of panel data with just two periods in my application.

1: $\sum_j m_j \gamma_{jk} = 1$.[24] Note that since labor demand only affects labor supply decisions through their impact on w_{kt}, which is common to all workers, labor demand shocks are effectively absorbed into this occupation $\times$ time fixed effect w_{kt}.

Intuitively, identification stems from occupation switchers. Consider the wage change of an individual of type j who moves from occupation k to k' between $t-1$ and t[25]

$$\ln \omega_{it} - \ln \omega_{it-1} = \underbrace{(\ln \gamma_{jk'} - \ln \gamma_{jk})}_{\text{Relative Skills}} + \underbrace{(\ln w_{k't} - \ln w_{kt-1})}_{\Delta \text{ Skill Price}} + \underbrace{(\ln \epsilon_{it} - \ln \epsilon_{it-1})}_{\text{Measurement Error}}$$

Workers can realize a wage increase on an occupation switch because (i) they have relatively high-skills in the destination occupation, (ii) they go to a relatively high price occupation or (iii) they realize high measurement error in the destination occupation. Importantly, there is one price in each occupation in any given period t that applies to all workers. Therefore, the endogenous price of labor in occupation k in period t (w_{kt}) is absorbed into an occupation-by-time fixed effect about which the marginal distribution of wages is highly informative. Since the ϵ_{it} are i.i.d. across job switches, the only systematic determinant of wage changes for occupation switchers after controlling for occupation-by-time fixed effects is the worker's skill for k relative to k'. The distribution of wage changes for workers switching from occupation k to k' therefore informs the distribution of relative skills in the economy. In addition, the frequency of moves from occupation k to k' further pin down the relationship between γ_{jk} and $\gamma_{jk'}$. Finally, normalizing the γ_{jk} to have unit mean within an occupation converts the distribution of relative skills into a distribution of skill levels.

The parameters governing the non-pecuniary benefits are principally affected by occupation choices and flows. The likelihood that a worker chooses low expected utility jobs is determined by the variance ν of the idiosyncratic taste shocks, which I assume is the same across worker types in the estimation. The level of employment in the economy informs the level of the fixed non-pecuniary benefits ξ_k. Meanwhile, the relative value of ξ_k to $\xi_{k'}$ allows the model to match the fact that many high wage occupations constitute small shares of overall employment. In this way, the ξ_k reflect not just the utility benefits of working in occupation k, but the broader compensating differentials earned by workers in each occupation.[26]

To maximize the likelihood function, I assume that measurement error in wages fol-

[24]This is without loss of generality. Were one to double the number of units of human capital that every worker possesses in an occupation, the equilibrium price of labor would halve.

[25]In the absence of overt switching costs, the model considered above may be solved as sequence of static economies. Occupation switching therefore occurs in the model due to shifts in skill prices w_{kt} or fresh draws of the idiosyncratic taste shocks ζ_{ikt}.

[26]Engineering, for instance, may have a low ξ_k not because engineering is an unpleasant occupation, but rather because the annualized cost of acquiring and maintaining engineering knowledge is high.

lows a log-normal distribution which is type-occupation specific, following Bonhomme et al. (2019), and that taste shocks are independent through time. This assumption is strong, as it generates close to random mobility. Stickiness in occupation choices therefore principally loads into small variance in ζ_{ikt}, which is in turn related to the micro elasticity of labor supply. Nevertheless, as discussed below, the estimated value of this variance generates a micro labor supply elasticity which closely aligns with those found in the literature.

The estimated Γ matrix reflects reasons why workers may see large wage changes when they change occupations and why workers may be more or less likely to switch between any two occupations. Therefore, it reflects both skill differences, but also various adjustment frictions excluded by my model, such as licensing requirements, differential bargaining power across occupations or information frictions. The estimated Γ matrix thus reflects workers' ability to be remunerated in each occupation, rather than pure skill. This may be reasonable for assessing aggregate wage dynamics, but is worth bearing in mind as a caveat.

Data and Implementation – Identification requires that every unobserved worker type forms a connecting cycle across occupations. As a result, using the full set of detailed Standardized Occupation Classification (SOC) codes is infeasible. To circumvent this challenge, I classify occupations into groups with similar skill requirements using a *k*-means algorithm. To do so, I first rank SOC occupations according to the share of workers with at least some college education using data from the Current Population Survey (CPS).[27] I then split occupations into terciles of educational attainment to rank occupations according to their general skill requirement.

Next, I cluster occupations within each education tercile according to the skill content required by the occupation. To do this, I employ data from O*NET, which surveys thousands of occupation holders about the level of skill and knowledge required to perform their job.[28] Skills include both hard skills, such as mathematics and science, and soft skills, such as critical thinking and social perceptiveness. Knowledge categories include specific occupational knowledge such as Personnel and Human Resources and Foreign Languages. Respondents rank the level of knowledge required for their job on a scale from 1 to 7, where examples are provided for select numeric values. Within each education tercile, I cluster occupations into five groups according to their required level of knowledge and skills using a *k*-means algorithm.

[27]Throughout, I harmonize occupation codes to follow the 2010 Census occupation coding provided by IPUMS, and use the crosswalk to detailed SOC codes from census. Note that clustering occupation codes has the added benefit of reducing the influence of measurement error in self-reported occupations. More details are in Appendix C.

[28]Gathmann and Schönberg (2010) build indices of skill relatedness using these surveys.

A full picture of the $K = 15$ clustered occupations is provided in Appendix C. Clusters are ordered according to their mean annual income in the period 2002-2006, as implied by data from the Bureau of Labor Statistics' Occupational Employment Statistics (OES). The occupation clustering is intuitive, with similar occupations being paired into the same cluster. Within each cluster, there remains a variety of occupations. For instance, the medical cluster pairs nurses together with surgeons. It is natural that these occupations might be clustered together within a broader medical clustering. However, surgeons are generally thought to be higher skill workers than are nurses. This would be captured by the γ_{jk}.

With the occupation clusters in hand, I turn to the estimation of the Γ matrix. I assume that the number of types J is equal to 8.[29] I use the March Supplement of the CPS going back to 1984 focusing on workers aged between 21 and 60 years old and link individuals over time to generate two-year panels; more details are provided in Appendix C. Worker earnings ω_{it} are measured by the total labor income of workers over the prior year, deflated by the CPI-U.[30] I drop workers who report earning less than \$1,000 in a year fearing that measurement error is large for these workers. Although the CPS surveys a relatively large sample, I estimate the model on data aggregating the period immediately before the Great Recession (2002-2006) to minimize sampling noise.

Estimation Results – Appendix Table A2 reports the transpose of the estimated matrix Γ, along with the mass of each type of worker m_j for the period 2002-2006. Each column reports the γ_{jk} vector for a given worker type j, while each row reports the γ_{jk} entry for a given occupation k. Worker types are ordered according to the mean of their γ_{jk} vector, reported in the row labeled $\mathbb{E}_k[\gamma_{jk}]$. In addition, the final column reports the non-pecuniary benefit of each occupation ξ_k, while the final row reports the variance of each column vector.

The table shows, for instance, that a type 1 worker supplies 0.81 units of human capital to routine occupations (cashiers, security guards etc.), but only 0.05 units of human capital to skilled business services occupations (such as financial analysts or management consultants). Recall the γ_{jk} are normalized to have unit mean (weighted by worker type shares) within each occupation. Thus, γ_{jk} may be interpreted as the

[29]Reducing the number of types to 5 or fewer removes the model's ability to match the aggregate data as there is little scope for composition effects with few types. Increasing the number of types renders estimation more noisy, but does not drastically change the macro implications of the model. Bonhomme et al. (2019), on which the estimation is based, allowed for the equivalent of $K = 10$ and $J = 6$.

[30]The model has no scope for hours to vary. As a result, hours-induced earnings fluctuations will appear as differences in workers' human capital levels γ. Additionally, I do not residualize earnings against observable characteristics, such as worker age or education, preferring instead to interpret predictable earnings differences from these observables as reflecting differences in human capital.

amount of human capital a type j worker has in occupation k relative to a mean worker in the economy.

The estimation is an excellent fit in-sample. The details of the model fit are provided in Appendix A and only briefly summarized here. The correlation between the estimated mean and variance of occupational wage distributions with those of the data is between 0.99 and 1. Similarly, the employment shares implied by the model match the data almost exactly. The model also accurately predicts the share of flows of from occupation k that go to any other occupation k'. At the (k, k') level, the correlation of occupation flows predicted by the model to those in the data is 0.84. That is, the model replicates the net flows between occupations, which is key for aggregate employment and wages. However, the model overpredicts gross occupational flows due to the i.i.d. assumption on the idiosyncratic preference shocks ζ_{ikt}. The estimation also produces intuitive patterns of skill relatedness. For instance, workers who are skilled in routine occupations, such as stock clerks, are also skilled as manual laborers (Figure A9).

4.2 Calibrating Labor Demand

Table 1 summarizes the model's calibration. I calibrate intermediate sectors to represent 3-digit NAICS sectors and assume that the elasticity of substitution η between intermediate sectors in the production of the final good is 4 following Broda and Weinstein (2006).[31] I further assume that the production function within sector s is Cobb-Douglas with returns to scale x_s and output elasticity with respect to occupation k given by $\alpha_{sk} x_s$. The Cobb-Douglas structure of production guarantees that the degree of diminishing returns in sector s, x_s, will be equal to labor's share of value added in sector s, while α_{sk} will be the share of sector s's wage bill that is accounted for by occupation k.[32] Hence x_s is chosen to match the BEA's estimate of the labor share of production in each sector, while the α_{sk} is chosen to match the share of the wage bill in each of the 15 occupation clusters in the BLS' Occupation Employment Statistics data series. These quantities are fixed to the average share in each sector over the period

[31] Broda and Weinstein (2006) estimate the mean elasticity of substitution across 3-digit SITC products, rather than sectors. Changing the elasticity of substitution across sectors would have the effect of reducing the dispersion of labor demand shocks for each occupation, as a shock to a particular sector is partially capitalized into the price of that sector's output. That is, it would affect the $\boldsymbol{d} \ln \boldsymbol{p} / \boldsymbol{d} \ln \boldsymbol{z}$s in equation (14).

[32] Imposing a CES production function would necessitate estimation the elasticity of substitution across occupations at the sector level, which is outside of the scope of this paper. A CES production function could increase or decrease cross-sector labor spillovers if the elasticity of substitution is greater than or less than 1, respectively. Intuitively, suppose there is a decline in the TFP in the construction sector. This reduces the price of manual laborers. If the elasticity of substitution across occupations is high in the manufacturing sector, this reduced price will induce the manufacturing sector to absorb some of these displaced laborers, substituting away from other occupations. In effect, the elasticity of substitution between occupations in production affects the ϵ^D in equation (14)

TABLE 1: Calibration Overview

PARAMETER	DESCRIPTION	SOURCE
Structural Estimation		
γ_{jk}	Effective Labor supply of type j	Maximum Likelihood
σ_{jk}	Variance of idiosyncratic Wage Draw	Maximum Likelihood
m_j	Share of workers who are type j	Maximum Likelihood
ξ_k	Compensating Differential of Occ k	Maximum Likelihood
ν	S.D. of T1EV shocks	Maximum Likelihood
External Calibration		
S	Number of Sectors	57 (# 3-Digit NAICS)
J	Number of types	8 (Imposed)
K	Number of occupations	15 (Imposed)
η	Elast. of Subs. Between Sectors	4 (Broda and Weinstein, 2006)
$F^{(s)}(l_{s1},\ldots,l_{sK})$	Sector s production function	$F^{(s)}(\mathbf{l}_s) = \left(\prod_{k=1}^{K} l_{sk}^{\alpha_{sk}}\right)^{x_s}$
x_s	Labor Share of Sector s	BEA Labor Share
α_{sk}	Share of Occupation k in Sector s	OES Share in Wage Bill
z_{st}	TFP series for sector s	Adjusted VA/Worker

2002-2006.

Estimating Sector-Level TFP Series – I estimate TFP using a Solow residual approach. A challenge arises from selection on unobservable quality in labor inputs. Through the lens of my model, the problem arises because $\bar{\gamma}_{kt}$ fluctuates. Specifically, let n^s_{kt} denote the number of workers employed in occupation k in sector s. Because workers are indifferent over sectors conditional on their occupation, the total labor units employed in occupation k in sector s are $l_{skt} = \bar{\gamma}_{kt} n^s_{kt}$. This implies that the TFP of sector s in period t may be estimated using the equation[33]

$$\text{(24)}\quad \ln z_{st} = \ln \text{Value Added}_{st} - x_s \sum_{k=1}^{K} \alpha_{sk} \ln(\bar{\gamma}_{kt} n^s_{kt}) - (1 - x_s) \ln(\text{Non-Labor Input}).$$

The employment in each sector in each occupation, n^s_{kt}, is observed in the data. However $\bar{\gamma}_{kt}$ is not directly observed. To calculate $\bar{\gamma}_{kt}$, I estimate the labor supply parameters – $\Gamma, \xi_k, m_j, \sigma_{jk}$ and ν, and the mean of the wage distribution for each type-occupation pair – in two-year rolling windows using the CPS every year from 1990 through to 2014. Equations (2) and (10) then yield an estimate of $\bar{\gamma}_{kt}$.

I employ sectoral value added and non-labor input data from the BLS' KLEMS Mul-

[33] Note that since neither workers nor firms make dynamic decisions, one does not need to take a stand on whether these shocks are anticipated.

tifactor Productivity Series.[34] These are annual data going back to 1987.[35] To calculate the employment of each occupation in each sector, I combine data from the Quarterly Census of Employment and Wages (QCEW) with data from the CPS. The QCEW provides the total employment by sector using administrative data derived from tax records. Using the CPS, I calculate the share of employment in each 3-digit NAICS sector that is accounted for by each of the 15 occupation clusters. Combining these gives an estimate of the total number of employees in each sector-occupation pair every year. Finally, I use equation (24) to estimate sectoral revenue TFP series.

The adjustment for $\bar{\gamma}_{kt}$ is meaningful. Appendix Table A3 describes the annual percentage changes in implied total factor productivity for the largest sectors in the 1990-91 and 2008-2009 recession. Whereas the BLS series shows no drop in productivity in the Construction sector in 2009, the series adjusted for human capital selection shows a 6 percentage point decline. In some sectors, however, the adjustment has little bite. For example, in the hospital and residential care facilities sector, both series show a 1.3% increase in productivity from 2008 to 2009. The fact that selection is unimportant in this sector is intuitive given the specialized nature of medical care. Aggregating sectoral TFP series according to their 2008 shares of aggregate value added, the adjusted TFP series shows a decline in aggregate productivity of 5.9%, compared to a 4.2% decline in the unadjusted BLS series.

4.3 Great Recession Aggregate Labor Market Dynamics

Figure 3 plots the aggregate labor market dynamics implied by the model when the skill distribution estimated from 2002-06 is subject to the selection-adjusted sectoral TFP realizations from 2008-2012. The black lines plot the evolution of average weekly earnings deflated by the CPI-U (solid line) and aggregate hours (dashed line) in the Current Employment Statistics. The blue solid line plots the evolution of real wages, while the green dashed line plots the evolution of employment, both relative to the pre-recession peak of March 2008, in the model. The model replicates the increase in average wages in 2009 with flat or declining wages in the following years, as well as an employment collapse and subsequent slow recovery. This is not guaranteed – the labor supply parameters and output elasticities in production are all estimated using pre-recession data.

The mechanism for this negative aggregate comovement is that the price of labor falls

[34] These value added data include sectoral prices. Thus the z_{st} arising from this calculation may reflect shocks both to productivity and to the demand for sector s, μ_{st}.

[35] The Pandemic recession featured an extremely sharp crash and recovery, which makes it difficult to study with annual data. For this reason, I focus on the Great Recession.

FIGURE 3: Aggregate Wage and Employment Responses, 2008-2012

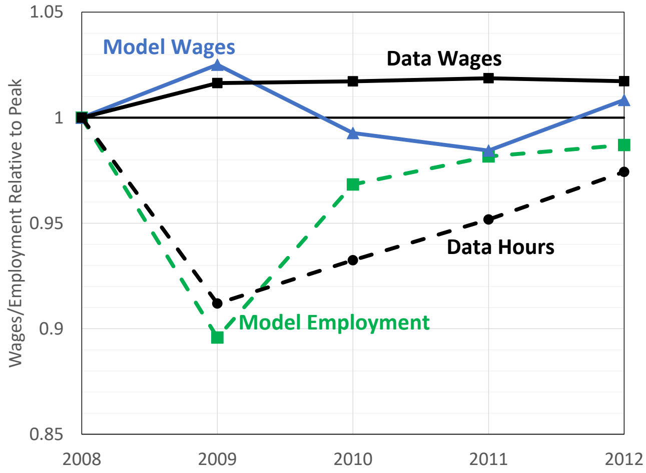

Notes: Figure plots the model-implied aggregate behavior of wages and employment in response to the calibrated sectoral TFP series around the Great Recession. Wages and employment normalized to be 1 in 2008. Black lines plot aggregate hours and average weekly earnings of all private employees, deflated by the CPI-U, from the Current Employment Statistics. Labor supply parameters estimated using data from 2002-2006.

in some occupations, inducing composition effects. This is true in the data as well. The correlation between log average occupation wages in 2008 and the change in occupational employment shares between 2008 and 2009 is 0.52 in the model and 0.44 in the Occupational Employment Statistics (OES) – low wage jobs saw larger employment declines in both model and data. Meanwhile, regressing occupational average wages from the OES against model-implied wages yields a coefficient of 0.81 (SE: 0.44) in 2008 and 0.83 (SE: 0.28) in 2009, neither of which are significantly different from 1. That is, the model successfully predicts occupational wage dynamics and the strength of the relationship between employment changes and wage levels.

Indeed, the model also predicts occupation-specific employment dynamics around the Great Recession. Regressing log employment in the data on occupation fixed effects and log employment in the model for the 15 occupation groups between 2007 and 2013 gives an R^2 of 0.83 with occupation fixed effects. That is, the model accounts for 83% of the variation in occupational employment movements around the Great Recession, at the level of these 15 occupations. Put differently, the model replicates net flows between occupations during this period.[36]

Following the Great Recession, the model generates a sluggish recovery in employment which is qualitatively consistent with the data. In the data, real wages did not increase much after 2009, which the model also predicts; however, the model actually

[36]Note however that the model overpredicts gross flows between occupations, though this has no impact on macro aggregates in this model featuring frictionless reallocation.

TABLE 2: Wage and Employment Changes During Great Recession

Specification	Wage Change (1)	Employment Change (2)
Data	+1.6%	-8.8%
Model: Calibrated	+2.5%	-10.4%
Model: $\gamma_{jk} = 1 \, \forall j, k$	-3.5%	-1.0%
Model: Only Comparative Advantage	-1.6%	-2.1%
Model: Only Absolute Advantage	-1.3%	-7.9%
Model: No Home Sector	-2.7%	0.0%

Notes: Table reports the wage (column 1) and employment (column 2) change between 2008 and 2009 in the data and a variety of model calibrations. Wages in the data correspond to average weekly earnings in the Current Employment Statistics, deflated by the CPI-U. The "Model: Calibrated" uses the skill distribution estimated in the CPS from 2002-2006. The model with only comparative advantage divides each worker type's skill vector by its mean so that all workers have the same average human capital. The model with only absolute advantage sets each worker type's skill vector to be a constant equal to its estimated mean. The "No Home Sector" row reports estimates from a model in which there is no home sector.

predicts a wage decline for a few years following the Great Recession. These patterns arise in the model because the workers re-entering the employed pool during the recovery are on average lower skill than those already employed; they thus have lower wages and are less attached to employment. This suggests worker heterogeneity may also be a partial contributor to the slow recovery of the labor market.

Table 2 shows the key ingredients for this result. Each row of the table represents the movement of aggregate labor market variables between 2008 and 2009 either in the data or a particular calibration of the model. Column 1 shows the implied change in real wages, column 2 shows the change in employment.

The table shows that real weekly earnings rose by 1.6% between 2008 and 2009, while employment fell by 8.8% in the data.[37] The calibrated model reveals a wage increase of 2.5% and employment decline of 10.4% over the same period. The composition effects generated by the model are sufficiently strong to generate the negative correlation between employment and wages observed in the data. This is noteworthy: exogenous shocks to labor demand generate negative comovements between measured employment and wages in the aggregate even without frictions.

The final five rows of the table illustrate the necessity of each ingredient of the model to generate these patterns. Each row selectively removes one element of the model and re-estimates the equilibrium response to the change in sectoral TFP between 2008 and 2009. The third row considers the case with no labor supply heterogeneity: every

[37] The data numbers consider year over year changes from March 2008 to March 2009.

worker has one unit of human capital that they can supply to any occupation. In this model, real wages decline by 3.5% while employment falls by just 1%. Without skill heterogeneity, there is no scope for reallocation and composition effects to buttress measured wages. As a result, the economy behaves as though a representative agent faced a decline in demand for their labor: both prices and quantities fall. The labor demand shock trades along the representative agent's inelastic labor supply curve. Indeed, the implied elasticity of labor supply in this representative agent model – the ratio of employment changes to wage changes – is 0.27, roughly in the range of micro labor supply elasticities found in the literature surveyed by Chetty, Guren, Manoli, and Weber (2011). Thus, shifting composition can also rationalize the disconnect between estimated micro and macro labor supply elasticities.

Both horizontal and vertical differentiation between workers are important. The fourth row of Table 2 removes absolute advantage from the economy, but maintains comparative advantage. To construct this counterfactual, I suppose that each worker type has the same mean γ_j, but maintains the estimated pattern of comparative advantage. That is, I construct a counterfactual Γ matrix by dividing each column of Table A2 by its mean. When faced with the decline in TFP between 2008 and 2009, this Γ matrix yields wage and employment declines of 1.6% and 2.1%, respectively. Composition effects are weak in the absence of absolute advantage since all worker types earn similar wages, so employment and wages continue to move together.

A worker fixed effect model also does not generate negative comovements between employment and wages. To construct this counterfactual, presented in the fifth row of the table, I assume that all worker types' vector of skills is a constant equal to the mean of their estimated γ_{jk} vector. In this model, employment falls by 7.9% while real wages fall by 1.3%. Here, there remains a meaningful composition effect: low-skill workers leave employment more than high-skill workers in response to the negative demand shocks, putting upward pressure on measured wages. However, workers' labor supply is relevant to all possible pursuits. Therefore, as foreshadowed in Section 3.1, the negative demand shocks to routine construction and manufacturing jobs of the Great Recession exert strong downward pressure on the price of, for instance, medical labor. These spillovers more than overcomes the selection force generated in the pure absolute advantage model, thereby preserving a positive comovement between employment and wages in the aggregate.

This shows that both absolute advantage and skill specificity are important to generate negative comovements between aggregate employment and wages in the face of labor demand shocks in equilibrium. Absolute advantage gives scope for composition effects, while comparative advantage limits the general equilibrium spillover effects that exert pressure on the price of labor elsewhere in the economy.

The "Model: No Home Sector" row of Table 2 removes the home option so that all workers are forced to work. In this case, there is no composition effect. Because all workers must work, labor is supplied inelastically and wages decline by 2.7%.

This section shows that worker and demand shock heterogeneity can be quantitatively important to understand the employment and wage dynamics of the Great Recession. An appropriately calibrated frictionless model with stable primitive labor supply parameters and only labor demand shocks successfully replicates the observed negative comovement between aggregate employment and wages over this period. Both skill specificity and absolute advantage are necessary for this result in equilibrium.

4.4 Counterfactuals

In this section, I study what led to the unusual negative comovements between employment and wages in the Great Recession through a series of counterfactuals.

4.4.1 Changes in Labor Supply

I begin by considering the importance of changes in the skill distribution. Figure 4 displays the predicted change in aggregate wages (Panel A) and employment (Panel B) relative to 2008 under the estimated model (blue solid line), and in two counterfactual economies. First, I re-estimate the labor supply side of the model using data from immediately prior to the 1990-91 recession: the years 1984-89.[38] The black dash-dot line reports the evolution of employment and wages in a model in which the skill distribution of 1984-89 were subjected to the TFP shocks of the Great Recession. That is, it studies the impact of the Great Recession's labor demand shocks were they to occur 20 years earlier. This counterfactual exercise shows that real wages would have declined by approximately 3% with employment falling approximately 2% were the Great Recession to occur with the skill distribution of the 1980s. This stands in stark contrast to the estimated model which predicts rising wages. Thus, changes in the nature of labor supply were important to generate the wage and employment patterns observed recently. I explore these changes below.

To begin, consider the effect of unilateral increases in the price of each occupation w_k relative to the estimated equilibrium labor prices as of 2007. Increasing these prices induces movements out of non-employment. One can calculate the implied labor supply

[38] I report the estimated labor supply parameters for the 1984-1989 period in Appendix Table A4.

FIGURE 4: Predicted Wage and Employment Dynamics in Great Recession under Counterfactual Skill Distributions and Labor Demand Shocks

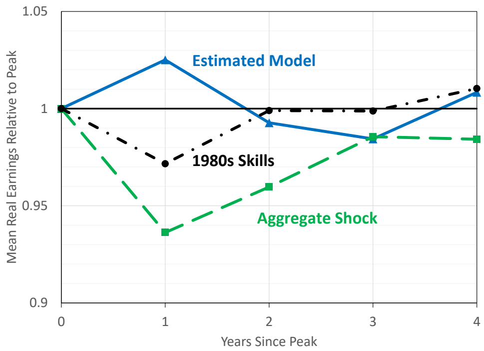

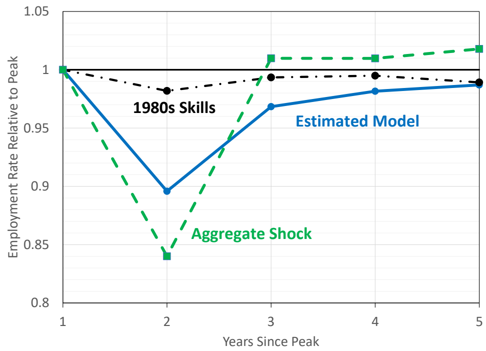

PANEL A: WAGES PANEL B: EMPLOYMENT

Notes: Figure reports the model-implied behavior of aggregate wages (Panel A) and employment (Panel B) under counterfactual skill distributions and sectoral shocks. The blue solid line reports the behavior of the estimated model around the Great Recession. The black dash-dot line plots a counterfactual in which the skill distribution is as estimated during the 1984-89 period. The gray dashed line shows an aggregate shock counterfactual in which all sectors saw the same movement in exogenous TFP.

elasticity of non-employment to the price of each occupation as $d \ln(1 - \bar{E}) / d \ln w_k$. Figure 5 plots these implied elasticities for each occupation. The black bars plot the elasticities for the 1984-1989 period, while the gray bars plot the elasticities for the 2002-2006 period.

The figure shows two noteworthy features. First, there is substantial variation in the elasticity of non-employment to changes in occupation prices. As a result, recessions and expansions that differ according to the sectoral (and thus occupational) composition of labor demand shocks will generate different aggregate wage-employment comovements, which appear as representative agent labor wedges as discussed in sections 3.2 and 3.3. Low-wage occupations generally have higher non-employment elasticities than do high wage occupations, such as engineering. Routine occupations have the highest non-employment elasticity, reflecting that many low wage workers move from routine occupations to non-employment and vice versa.

Second, the figure shows that non-employment elasticities of labor supply have generally risen through time. Whereas the mean elasticity of non-employment to changes in the price of occupation-specific labor was -0.12 in 1984, that fell to -0.33 in 2002-2006. As a result, for any given change in the price of labor, one might expect to see larger fluctuations in employment in the mid-2000s relative to the late 1980s.[39]

[39]Note that these estimates of occupation-specific non-employment elasticities are roughly in line with the micro labor supply elasticities estimated elsewhere (Chetty et al., 2011).

FIGURE 5: Estimated Labor Supply Elasticities for Each Occupation Over Time

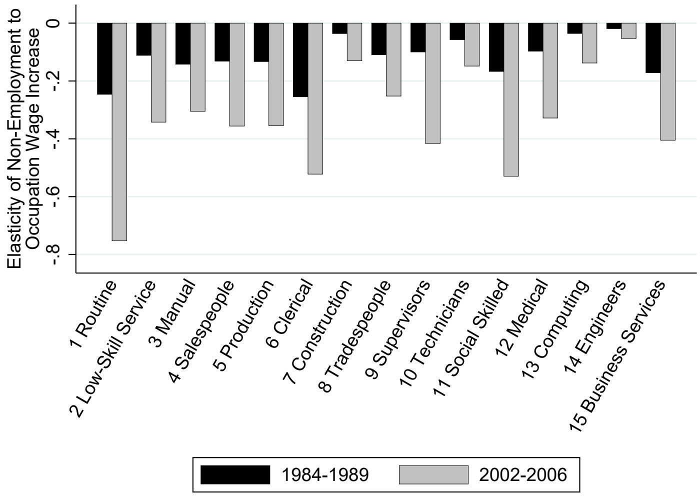

Notes: Figure reports the estimated model-implied elasticity of non-employment to a change in the price of each occupation's price of labor w_k. The estimation procedure is outlined in Section 4.1, and carried out separately in the CPS March Supplement for the periods 1984-1989 (black bars) and 2002-2006 (gray bars). Elasticities are defined by calculating the percentage change in non-employment rates in response to a unilateral 1% change in the price of labor relative to the 2007 equilibrium price in each occupation.

This change in the elasticity of labor supply primarily results from changes in the skill distribution and the standard deviation of idiosyncratic preference shocks ν. This standard deviation is estimated to have declined from 0.60 to 0.29 so that workers have become more responsive to changes in labor prices in their occupation choice. This stems from the decline in the average wage change of occupation switchers in the data, from 6.4% in the 1984-89 period to 4.5% from 2002-06.

Changes in the Γ matrix are explored in depth in Section A.7 and outlined here. First, skills have become more heterogeneous between workers. In the period before the 1991 recession, top workers supplied 4.66 units of human capital to the market in an average occupation. By contrast, the lowest type workers only supplied 0.44 units of human capital, roughly one-tenth that of the highest types. The cross-type range of skills has increased, with the best workers in the 2002-2006 period supplying 7.54 units of human capital on average, compared with 0.55 for type 1 workers. The variance in average skills across workers has also grown by 10.4% in this period – the gap in earnings potential of the most and least skilled workers has grown over time. This increases the scope for strong composition effects through rising inequality and because low-wage workers may be more on the margin of non-employment. It additionally increases own-price labor supply elasticities, as discussed in section 3.1.

The second relevant change in the skill distribution is that skills have become more specific. A natural measure of the specificity of a worker type j's skills is the variance of their vector of γ_{jk}s. If this variance is high, it indicates that the worker is much better at some tasks than others and thus cannot easily transport their skills across jobs. The cross-type average of this "within-type variance" grew by nearly 50% between the 1980s and 2000s. Growing skill specificity limits cross-occupation labor supply spillovers and contributes to increased non-employment elasticities as workers are unable to reallocate to other productive pursuits.

Heuristically, these results arise from two moments in the data. The increase in within-type variance owes to an increase in the variance of log earnings changes on occupation switches from 0.37 in the late 1980s to 0.52 in the mid 2000s. The higher is this variance, the more one infers that individual workers' skills are better tailored to particular applications. Meanwhile, the increase in cross-type variance is inferred from a rise in the within-occupation variance in wages, as this moment reflects the degree to which workers differ in their skill within each occupation. Indeed, a regression of log earnings on occupation cluster fixed effects has an R^2 of 0.19 in the 2002-06 period, but just 0.16 in the 1984-89 period.

The results presented in this section suggest that a representative agent framework with interchangeable skills has become a worse approximation of reality. The gap between the most and least skilled workers in the economy has risen. Comparative advantage has similarly risen: workers have become more specialized over the last twenty years. There may be many reasons for these changes, such as changes in education policy, occupational licensing or technological change in the task composition of occupations. Understanding the source of these changes is outside of the scope of this paper, but is fertile ground for future research. The primitive occupation-specific labor supply elasticities have also increased over time, thereby leading to larger employment declines and smaller wage declines in the face of a given set of labor demand shocks. As a result, one might expect larger cyclical fluctuations in employment and larger composition effects on wage cyclicality in the future. Indeed, Appendix A.6 provides evidence that composition effects have grown more important over the last 60 years and demonstrates their importance during the Great Recession in particular.

4.4.2 Labor Demand Shock Changes

The majority of industries receiving large negative shocks in the Great Recession primarily employ manual laborers. With the exception of the insurance industry, the largest declines in labor demand were concentrated in manufacturing, transportation, construction, and mining. This is highlighted by Figure 6, which plots the percentage

FIGURE 6: Industry TFP shocks by Share of Wage Bill in Manual Occupations

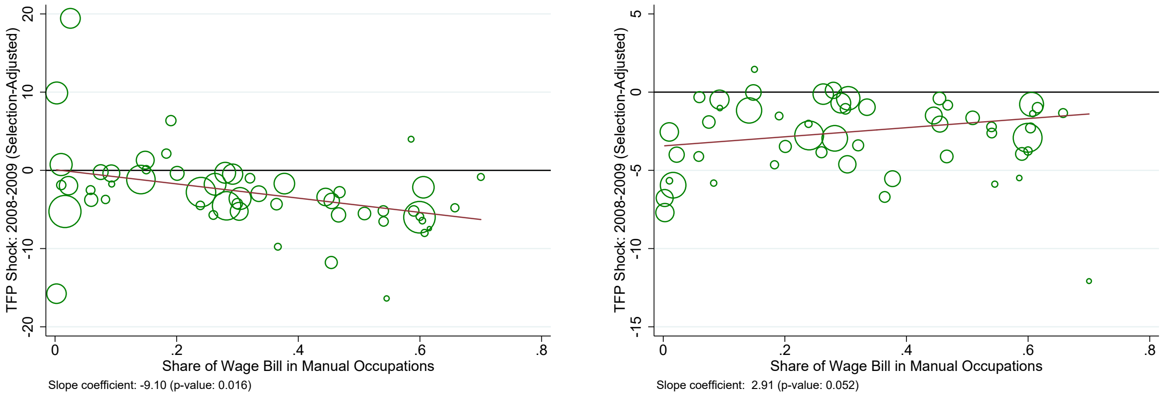

PANEL A: 2008-2009 PANEL B: 1990-91

Notes: Figure shows the percentage change in selection-adjusted TFP between 2008 and 2009 (Panel A) and between 1990-91 (Panel B) by the share of an industry's wage bill that accrues to manual occupations. Manual occupations are defined to be routine ($k = 1$), manual ($k = 3$), production ($k = 5$), construction ($k = 7$), and tradespeople ($k = 8$) occupations. Each dot is a different 3-digit NAICS sector, and its size is proportional to the value added share of that sector in the immediate pre-recession year.

change in selection-adjusted TFP between 2008 and 2009 (Panel A) and, as a point of comparison, during the 1990-91 recession (Panel B) against the share of an industry's wage bill that accrues to manual occupations.[40] Each dot is a different 3-digit NAICS sector, and its size is proportional to the value added share of that sector in the immediate pre-recession year. The figure shows that nearly all sectors mostly employing manual laborers saw large declines in labor demand during the Great Recession, but this was not true during the 1990-91 recession.

The green dashed lines in Figure 4 report a counterfactual evolution of aggregate employment and wages were there no sectoral heterogeneity in TFP shocks around the Great Recession. Specifically, it assumes that all sectors had declines in TFP of 5.9%: the average decline of TFP observed in the data, weighted by pre-recession sectoral value added. Under this counterfactual set of labor demand shocks, the model implies that real wages would have declined by approximately 6% with large employment declines. The sectoral shocks during the Great Recession ensured that workers who have skills in manual occupations both received a large negative demand shock for their skills and had little scope to apply their human capital elsewhere. As a result, they left the employed pool and, due to skill specificity, exerted limited downward pressure on the price of labor in other occupations. Since these workers tend to be low skill (i.e. have low γ_{jk}s), this generated a large composition effect with limited spillovers to the rest of the economy. Both the unique nature of labor demand shocks and the shifting structure of labor supply contributed to the negative comovement

[40] Manual occupations are defined to be routine ($k = 1$), manual ($k = 3$), production ($k = 5$), construction ($k = 7$), and tradespeople ($k = 8$) occupations

between employment and wages observed during the Great Recession.

5. CONCLUSION

This paper argues that skill heterogeneity and specificity have important effects on the aggregate effects of labor demand shocks. The general equilibrium effects of diffuse labor demand shocks are determined by the distribution of skills: skill specificity reduce cross-price labor supply elasticities, while differences in average levels of skill across workers make labor supply most responsive to the choices of the most productive. Shifts in the price of labor of different tasks induce workers to reallocate across occupations or out of the employed pool. As a result, shocks to labor demand generate composition effects on the aggregate wage, which may overturn the traditional intuition that demand shocks induce employment and wages to positively comove. Such composition effects are larger as wage inequality rises. Aggregation thus offers a frictionless microfoundation for movements in a representative agent labor wedge, exogenous representative agent labor supply shocks, or the large gap between estimated micro and macro labor supply elasticities.

This mechanism appears quantitatively important in equilibrium. I estimate the multidimensional distribution on skills by using two-period panel data on workers' occupation choices and earnings. Feeding a sequence of realistic shocks to sectoral labor demand through the estimated model replicates the large decline in employment from 2008 to 2009 and the roughly 2% rise in real wages over this period, as well as net flows between job types. Although the model is frictionless and the only shock is to labor demand, endogenous compositional shifts are sufficiently strong to cause employment and wages to move in opposite directions. Skill specificity and differences in average skill across workers are both necessary for this result. The model suggests that the strong composition forces of this period arose both due both to changes in the skill distribution – skills have become more specific and the gap between lowest and highest skill workers has grown – and from the particular composition of labor demand shocks. Future research should seek to understand the source of these changes and extend the model presented here to incorporate various dynamic considerations.

REFERENCES

ABOWD, J. M., D. N. MARGOLIS, AND F. KRAMARZ (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333.

ABRAHAM, K. G. AND L. F. KATZ (1986): "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy*, 94, 507–522.

BALEY, I., A. FIGUEIREDO, AND R. ULBRICHT (2022): "Mismatch Cycles," *Journal of Political Economy*, 130, 2943–2984.

BARLEVY, G. (2002): "The Sullying Effect of Recessions," *Review of Economic Studies*, 69, 65–96.

BEAUDRY, P., D. A. GREEN, AND B. SAND (2012): "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory," *Econometrica*, 80, 1063–1104.

BENHABIB, J. AND R. ROGERSON (1991): "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy*, 99, 1166–1187.

BERAJA, M., E. HURST, AND J. OSPINA (2019): "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 87.

BERGER, D., K. HERKENHOFF, AND S. MONGEY (2022): "Labor Market Powe," *American Economic Review*, 112, 1147–1193.

BILS, M. J. (1985): "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy*, 93, 666–689.

BÖHM, M. J., H.-M. VON GAUDECKER, AND F. SCHRAN (2024): "Occupation Growth, Skill Prices, and Wage Inequality," *Journal of Labor Economics*, 42.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 87, 699–738.

BRINCA, P., V. V. CHARI, P. J. KEHOE, AND E. R. MCGRATTAN (2016): "Accounting for Business Cycles," *NBER Working Paper Series*, 22663.

BRODA, C. AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 121, 541–585.

BURSTEIN, A., E. MORALES, AND J. VOGEL (2019): "Changes in Between-Group Inequality: Computers, Occupations, and International Trade," *American Economic Journal: Macroeconomics*, 11, 345–400.

CAJNER, T., L. D. CRANE, R. A. DECKER, J. GRIGSBY, A. HAMINS-PUERTOLAS, E. HURST, C. KURZ, AND A. YILDIRMAZ (2020): "The U.S. Labor Market during the Beginning of the Pandemic Recession," *Brookings Papers on Economic Activity*, Spring.

CARRILLO-TUDELA, C. AND L. VISSCHERS (2023): "Unemployment and Endogenous Reallocation Over the Business Cycle," *Econometrica*, 91, 1119–1153.

CARRILLO-TUDELA, C. AND D. WICZER (2022): "Cyclical Earnings, Career and Employment Transitions," *IZA Working Paper Series*, 15603.

CHANG, Y. (2000): "Wages, business cycles, and comparative advantage," *Journal of Monetary Economics*, 46.

CHANG, Y. AND S.-B. KIM (2007): "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations," *American Economic Review*, 97.

CHANG, Y., S.-B. KIM, K. KWON, AND R. ROGERSON (2018): "Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins," .

CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2007): "Business Cycle Accounting," *Econometrica*, 75, 781–836.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2011): "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review Papers and Proceedings*, 101, 471–75.

CHODOROW-REICH, G. AND J. WIELAND (2019): "Secular Labor Reallocation and Business Cycles," *Journal of Political Economy*, 128.

CHRISTIANO, L. J. AND M. EICHENBAUM (1992): "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," *American Economic Review*, 82, 430–450.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

CHRISTIANO, L. J., M. S. EICHENBAUM, AND M. TRABANDT (2015): "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7, 110–167.

CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

CORTES, G. M. AND G. GALLIPOLI (2018): "The Costs of Occupational Mobility: An Aggregate Analysis," *Journal of the European Economic Association*, 16, 275–315.

DVORKIN, M. (2014): "Sectoral Shocks, Reallocation and Unemployment in Competitive Labor Markets," .

EROSA, A., L. FUSTER, G. KAMBOUROV, AND R. ROGERSON (2025): "Labor Market Polarization and Inequality: A Roy Model Perspective," Working Paper 33687, National Bureau of Economic Research.

FLOOD, S. M. AND J. PACAS (2008): "Using the Annual Social and Economic Supplement as Part of a Current Population Survey Panel," *Census Bureau Working Paper*, 1–36.

FOERSTER, A., A. HORNSTEIN, P.-D. SARTE, AND M. W. WATSON (2019): "Aggregate Implications of Changing Sectoral Trends," *NBER Working Paper*, 25867, 1–55.

FORNI, M. AND L. REICHLIN (1998): "Let's Get Real: Factor to Analytical Approach to Disaggregated Business Cycle Dynamics," *The Review of Economic Studies*, 65, 453–473.

GARIN, J., M. PRIES, AND E. SIMS (2018): "The Relative Importance of Aggregate and Sectoral Shocks and the Changing Nature of Economic Fluctuations," *American Economic Journal: Macroeconomics*, 10, 119–148.

GATHMANN, C. AND U. SCHÖNBERG (2010): "How General Is Human Capital? A Task-Based Approach," *Journal of Labor Economics*, 28, 1–49.

GEEL, R. AND U. BACKES-GELLNER (2009): "Occupational Mobility Within and Between Skill Clusters: An Empirical Analysis Based on the Skill-Weights Approach," *Leading House Working Paper*, 47.

GERTLER, M., C. HUCKFELDT, AND A. TRIGARI (2020): "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," *The Review of Economic Studies*, 87, 1876–1914.

GERTLER, M. AND A. TRIGARI (2009): "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 117, 38–86.

GRIGSBY, J., E. HURST, AND A. YILDIRMAZ (2021): "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, 111, 428–71.

GUVENEN, F., B. KURUSCU, S. TANAKA, AND D. G. WICZER (2020): "Multidimensional Skill Mismatch," *American Economic Journal: Macroeconomics*, 12, 210–244.

HAGEDORN, M. AND I. MANOVSKII (2013): "Job Selection and Wages over the Business Cycle," *American Economic Review*, 103, 771–803.

HALL, R. E. (1997): "Macroeconomic Fluctuations and the Allocation of Time," *Journal of Labor Economics*, 15.

——— (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *The American Economic Review*, 95, 50–65.

HECKMAN, J. J. AND G. SEDLACEK (1985): "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market," *Journal of Political Economy*, 93.

LAZEAR, E. P. (2009): "Firm-Specific Human Capital: A Skill-Weights Approach," *Journal of Political Economy*, 117, 914–940.

LILIEN, D. M. (1982): "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy*, 90, 777–793.

MACALUSO, C. (2019): "Skill Remoteness and Post-Layoff Labor Market Outcomes," *Working Paper*.

MUELLER, A. I. (2017): "Separations, sorting, and cyclical unemployment," *American Economic Review*, 107, 2081–2107.

NEDELKOSKA, L., F. NEFFKE, AND S. WIEDERHOLD (2015): "Skill Mismatch and the Costs of Job Displacement," *CESifo Area Conference on the Economics of Education Papers*.

PATTERSON, C. (2023): "The Matching Multiplier and the Amplification of Recessions," *American Economic Review*, 113, 982–1012.

PILOSSOPH, L. (2014): "Sectoral Shocks and Move Unemployment," .

POLETAEV, M. AND C. ROBINSON (2008): "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000," *Journal of Labor Economics*, 26.

QUAH, D. AND T. J. SARGENT (1993): "A Dynamic Index Model for Large Cross Sections," in *Business Cycles, Indicators, and Forecasting*, ed. by J. H. Stock and M. W. Watson, University of Chicago press, January, chap. 7, 285–310.

ROGERSON, R., R. SHIMER, AND R. WRIGHT (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43.

ROGERSON, R. AND J. WALLENIUS (2009): "Micro and Macro Elasticities in a Life Cycle Model with Taxes," *Journal of Economic Theory*, 144.

ROYS, N. A. AND C. R. TABER (2022): "Skill Prices, Occupations, and Changes in the Wage Structure for Low Skilled Men," .

SOLON, G., R. BARSKY, AND J. A. PARKER (1994): "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *The Quarterly Journal of Economics*, 109, 1–25.

SORKIN, I. (2018): "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 133, 1331–1393.

STOCKMAN, A. C. (1983): "Aggregation Bias and the Cyclical Behavior of Real Wages," .

TRAIBERMAN, S. (2019): "Occupations and Import Competition: Evidence from Denmark," *American Economic Review*, 109, 4261–4301.

WASMER, E. (2006): "General versus Specific Skills in Labor Markets with Search Frictions and Firing Costs," *American Economic Review*, 96, 811–831.

WELCH, F. (1969): "Linear Synthesis of Skill Distribution," *The Journal of Human Resources*, 4, 311–327.

WICZER, D. (2015): "Long-Term Unemployment: Attached and Mismatched?" Tech. rep.

ŞAHIN, A., J. SONG, G. TOPA, AND G. L. VIOLANTE (2014): "Mismatch unemployment," *American Economic Review*, 104, 3529–3564.

Appendix: For Online Publication

Appendix A. Additional Analyses

This appendix presents additional results that are supplementary to those contained in the main text. and wages. First, I present reduced form evidence for labor supply spillovers of the sort hypothesized in Section 3.1. Next, I study when labor price movements can generate negative comovements between aggregate employment Then, I present the estimated Γ matrices over the period 2002-06 and 1984-89, as well as the selection-corrected industry TFP shocks used in the model application of Section 4.3. Then I present the in-sample fit of the estimated model for occupation-switching, employment shares and wages. After that, I present reduced form evidence that composition effects have been strengthening over time and were particularly strong during the Great Recession. Finally, I present the changes in the estimated skill distribution between 1984-89 and 2002-06.

A.1 Reduced Form Evidence for Labor Supply Spillovers

In this section, I test the model's implication that a negative shock to a sector s will induce positive labor supply spillovers to sectors with skills related to s. To do so, I exploit the sudden precipitous decline in labor demand in the Mining and Utilities sectors between 2014 and 2016. Towards the end of 2014, the Chinese government, fearing the formation of a credit bubble, implemented contractionary monetary policies. Concurrently, the booming American economy prompted the Federal Reserve to raise interest rates slowly, strengthening the dollar in the process. This further put pressure on many emerging economies, whose firms had many debt obligations denominated in dollars. The result of the Chinese expansion and strengthening dollar was a steep decline in emerging markets' demand for commodities, leading to a sharp drop in prices. Crude oil fell from $106 per barrel at the end of 2014, to just over $30 per barrel in early 2016, while prices for aluminum, copper, tin, and other commodities similarly fell. The end result was an aggregate decline in mining employment of over 30% in the span of just 2 years.[41] The time series of aggregate mining and utilities employment is shown in Panel A of Figure A1. That the decline in employment was restricted

[41]See Irwin, Neil. "The Most Important Least-Noticed Economic Event of the Decade." *New York Times*, Sept. 29, 2018. https://www.nytimes.com/2018/09/29/upshot/mini-recession-2016-little-known-big-impact.html (Accessed 10/19/2021) for an excellent synopsis of this economic period.

FIGURE A1: Shock to Mining Employment

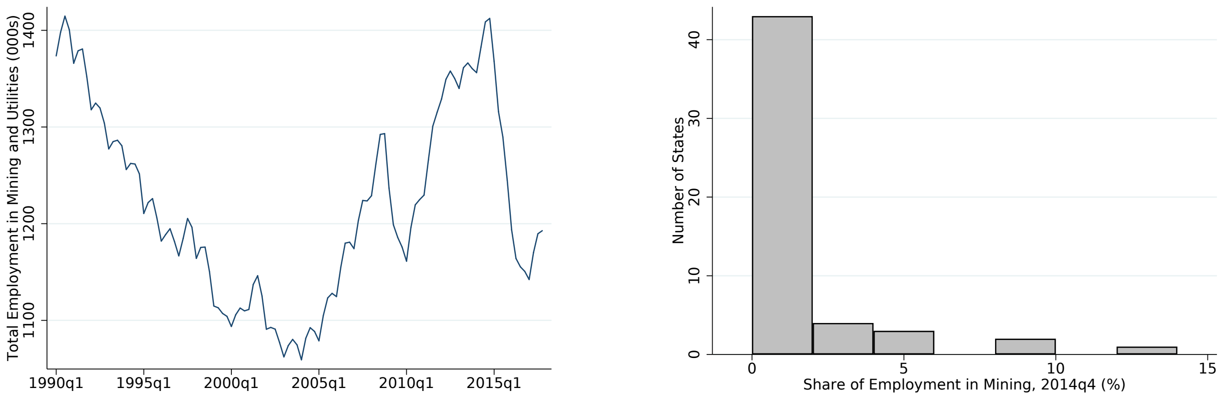

PANEL A: AGGREGATE MINING EMPLOYMENT

PANEL B: HISTOGRAM OF STATE-LEVEL SHARE OF 2014 EMPLOYMENT IN MINING

Notes: Figure plots the time series of aggregate employment in mining sector, and a histogram of the share of total state-level employment in mining as of the fourth quarter of 2014. Data come from the QCEW.

to mining and utilities merits emphasis – this period was one of rapid expansion of employment in the US, with both aggregate employment and mean wages rising.

This mining shock had heterogeneous impacts on local communities. For some states—such as West Virginia, Texas and North Dakota—mining is a significant share of employment, while for others, such as Massachusetts and Florida, mining is a relatively small share of employment. As a result, this aggregate mining shock generates a larger labor demand shock in states like Texas than it did in Florida, providing a laboratory to study the impact of a sectoral decline on related sectors.

Let λ_r^{MINING} be the share of region r's employment that is in mining as of the fourth quarter of 2014, and let $\Delta \ln E_{MINING,-r}$ denote the percent change in mining employment in all states other than r between the fourth quarter of 2014 and the fourth quarter of 2016. Define the predicted employment loss from mining in a region r to be $\sigma_r = |\lambda_r \Delta \ln E_{MINING,-r}|$: that is, the interaction of the national employment change in mining with the pre-existing share of employment in state r. If this negative labor demand shock to mining constitutes a positive labor supply shock to sectors with related skills, then one would expect the share of non-mining employment to rise in sectors more related to mining, while the wages of those sectors would fall relative to unrelated sectors. These patterns should be more concentrated in states with a higher pre-existing mining share of employment.

To test these hypotheses, I construct a reduced-form measure of the skill distance between sectors using the O*NET survey data, similar to the approach of Gathmann and Schönberg (2010). Given the responses to this survey, one can construct vectors of skills required for each occupation. To aggregate to sector-level skill vectors, let χ_{sk} be the share of employees in sector s who are employed in occupation k in the CPS and let $h_m(k)$ be the level of skill m required for occupation k according to O*NET. Define the level of skill m required by sector s to be the weighted average of $h_m(k)$, where the weights are the shares of s's employment in occupation $k : \chi_{sk}$. That is,

$$\bar{h}_m(s) = \sum_k \chi_{sk} h_m(k).$$

One can interpret this measure to be the expected skill vector a worker would require in sector s if one were to randomly sample workers in that sector.

I then define the skill distance between two sectors $d(s, s')$ to be the directed Euclidean distance between their skill vectors.[42] It should be noted that these survey measures do not provide a *cardinal* measure of skill relatedness. The goal here is to provide model-free reduced form evidence of cross-sector labor supply spillovers from skill transferability, rather than to estimate the exact size of these spillovers.

I combine these skill distance measures with data from the Quarterly Census of Employment and Wages (QCEW), which provides information on the average weekly earnings and employment levels at the sector-state level for every quarter back to 1975. I restrict attention to the set of tradable 3-digit NAICS sectors which have skills which are highly related or unrelated to mining. Sectors with highly related skills are defined as those in the bottom quartile of skill distance to mining – $d(s, Mining)$ is small – while those with unrelated skills are in the top quartile of skill distance. Restricting attention to tradable sectors isolates local labor supply effects by abstracting from movements in local labor demand resulting from the decline in mining. I estimate the following regression at the region-sector level

$$\Delta \ln y_{sr} = \alpha \cdot \mathbf{1}\{s \text{ is Related}\} + \eta \cdot \sigma_r + \beta \mathbf{1}\{s \text{ is Related}\} \cdot \sigma_r + \epsilon_{sr} \tag{A1}$$

where ΔZ is an operator which takes the difference in the variable Z between the fourth quarters of 2016 and 2014. I do this for two dependent variables y: real average weekly wages from the QCEW and the share of non-mining employment in region r that is in sector s. The hypothesis is that sectors that have more related skills to mining should experience a relatively large positive labor supply shock, particularly in states with large mining shares. This should increase employment and reduce wages in those

[42] This may be microfounded by supposing there is a cost $c(h', h)$ of acquiring skill level h' given that a worker is already at skill level h. Construct the distance between k and k' as $d(k, k') = G(\sum_m c(h_m(k'), h_m(k)))$ for $h_m(k)$ the level of skill m required by k, and G some function. If $c(h, h') = \max\{0, h' - h\}^2$, and $G(x) = \sqrt{x}$, then $d(k, k')$ is a directed Euclidean distance.

TABLE A1: Response of Sectors to Mining Shocks

	Change in Emp. Share		Change in Log Mean Wage	
	(1)	(2)	(3)	(4)
Related Skills $\times$ Mining Decline	0.040***	0.040***	-0.041***	-0.041***
	(0.006)	(0.006)	(0.008)	(0.008)
Trend Control	N	Y	N	Y
Observations	784	742	727	716
Mean of Dep. Var.	-0.014	-0.015	0.001	0.000
S.D. of Dep. Var.	0.080	0.082	0.087	0.085

Notes: Table reports coefficients estimated from equation (A1). Sectors with related skills are defined to be those sectors in the bottom quartile of skill distance with Mining sectors. Only tradable sectors in the top and bottom quartile of skill distance included. Standard errors clustered at 3-digit NAICS sector code level reported in parentheses.

sectors, suggesting that $\beta > 0$ for employment, and $\beta < 0$ for wages.

The results are presented in Table A1. Columns 2 and 4 control for state-sector-specific trends (i.e. long run growth between 1990 and 2014), while columns 1 and 3 do not. The table shows that sectors with skills related to mining experienced larger declines in wages and increases in employment relative to sectors with unrelated skills in states which had large pre-existing mining shares. This suggests that the decline in mining from 2014-2016 led to a disproportionate positive labor supply shock for related sectors relative to unrelated sectors. A one standard deviation increase in the size of the regional exposure to the mining decline is associated with an increase of 4 percentage points in the share of workers employed in sectors with skills related to mining, relative to sectors with unrelated skills. This is coupled with a relative decline in log wages of approximately 4% in these sectors. These patterns are consistent with positive labor supply spillovers from the mining sectors to sectors most related to mining.

A.2 Negative Aggregate Employment-Wage Comovements

This section studies when aggregate employment and wages may move in opposite directions. Figure A2 partitions the shock space (dw_1, dw_2) for the skill matrices defined in (17) in a two-type, two-occupation version of the labor supply block of the model. In each panel, the dashed black line plots the combination of labor price movements that induce no aggregate employment change, while the solid black line plots the combination of labor price movements that induce no aggregate wage change, calculating by solving for zero changes using equations (19) and (20), respectively. Panel A assumes a representative agent skill matrix and that the pre-shock labor prices are

FIGURE A2: Splitting Regions of Shock Space by Direction of Aggregate Wage and Employment Movements in Economy with Two Job Types

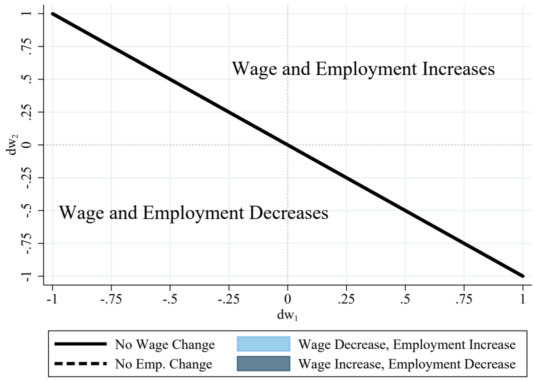

PANEL A: $\Gamma^{(RA)}$: REPRESENTATIVE AGENT, $w_1 = w_2 = 1$

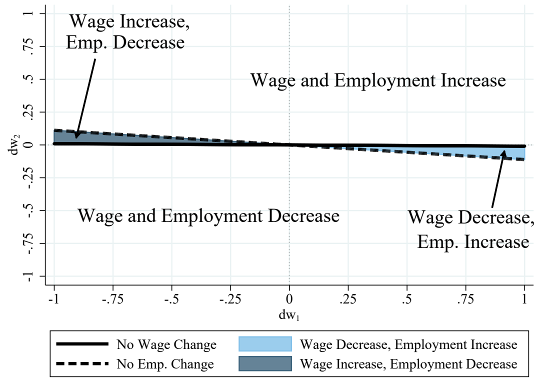

PANEL B: $\Gamma^{(RA)}$: REPRESENTATIVE AGENT, $w_1 = 0.75$, $w_2 = 1.5$

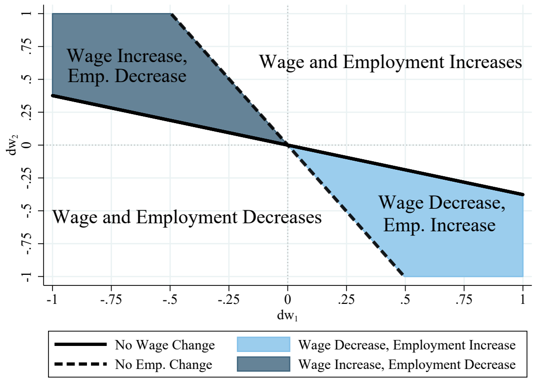

PANEL C: $\Gamma^{(CA)}$: COMPARATIVE ADVANTAGE, $w_1 = 0.75$, $w_2 = 1.5$

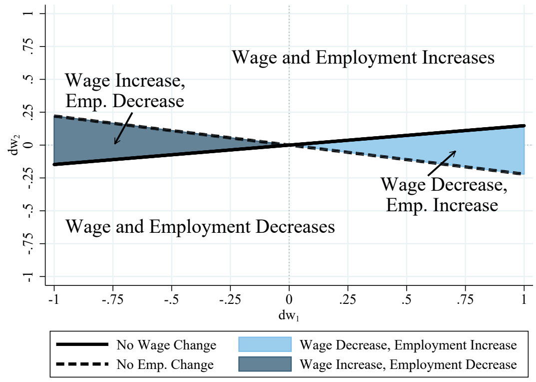

PANEL D: $\Gamma^{(AA)}$: ABSOLUTE ADVANTAGE, $w_1 = 0.75$, $w_2 = 1.5$

Notes: Figure divides space of labor price shocks into four regions for a variety of initial labor prices (w_1, w_2) and skill matrices Γ in a labor supply model with $K = J = 2$. The solid and dashed black lines are a set of labor price shocks that induce no aggregate wage and employment change, respectively. Panel A considers shocks from initial labor prices $w_1 = w_2 = 1$, while Panels B-D start from initial labor prices $w_1 = 0.75, w_2 = 1.5$. The skill matrices Γ are defined in equation (17). Text describes additional calibration.

equal to one another: $w_1 = w_2 = 1$. In this case, both of these lines have a slope of minus 1 passing through the origin: in order for there to be no employment change, it must be the case that a positive shock to w_1 is offset by a negative shock to w_2 of exactly the same size. Indeed, in this specification, equations (19) and (20) imply that the economy behaves as if there is just one shock whose size is the sum of the shocks to w_1 and w_2. As a result, the exact composition of shocks does not matter; furthermore, all shock combinations will induce employment and wages to move in the same direction. These lines always overlap if the direct effect is the only force operating on aggregate wages.

Panels B through D introduce some price dispersion by assuming that $w_1 = 0.75$ and $w_2 = 1.5$ before the shock. In this case, the solid and dashed lines diverge in all three skill matrices. This is due to the reallocation and composition effects. Such divergences are important because shock combinations which lie between the zero-employment and zero-wage change schedules induce aggregate employment and wages to move in opposite directions.

Panel B persists with the representative agent skill matrix. Even when workers have identical skills, price dispersion opens up the possibility that aggregate employment and wages move in different directions. Small movements in w_1 induce very little change in employment as occupation 1 is already far less attractive than occupation 2. They can, however, induce changes in wages. Movements in w_1 induce some set of marginal workers to reallocate from job 2 to job 1, putting downward pressure on average wages. As a result, when shocks to w_1 (which have some effect on wages but no effect on employment) are far larger than offsetting shocks to w_2, employment and wages can negatively comove.

The region of employment-wage divergence is substantially larger in Panel C, which considers the comparative advantage skill matrix. This is because type 1 workers tend to be low paid, as they sort towards the low-price job in which they have skills. What's more, type 1 workers are far more responsive to movements in the price of labor in job 1 than are type 2 workers. Thus, when w_1 increases, employment growth predominantly stems from low-wage workers. If a w_1 increase is coupled with a moderate decline in w_2, aggregate employment and wages may diverge.

Finally, Panel D shows that employment and wages can diverge in the presence of absolute advantage, even if all labor price shocks are positive. In a worker fixed effect model, increases in labor prices may disproportionately affect the employment of low-wage workers. This composition effect can be stronger than the direct effect on wages, inducing aggregate wage declines even if all labor prices increase.

A.3 Aggregate Wage-Employment Relationships Under Different Γ

This section explores relationship between employment and wages in the face of aggregate shocks for a variety of skill matrices. Figure A3 is constructed analogously to Figure 2, considering only aggregate shocks to skill prices (so that $w_1 = w_2 = w$ always) and the three skill matrices defined as in equation (17). The red lines report the plots for the comparative advantage skill matrix, which is discussed in section 3.3.

The black line shows the case in which there is a representative agent skill matrix.

FIGURE A3: Aggregate Employment-Wage Schedule As Vary Γ and Occupation Prices

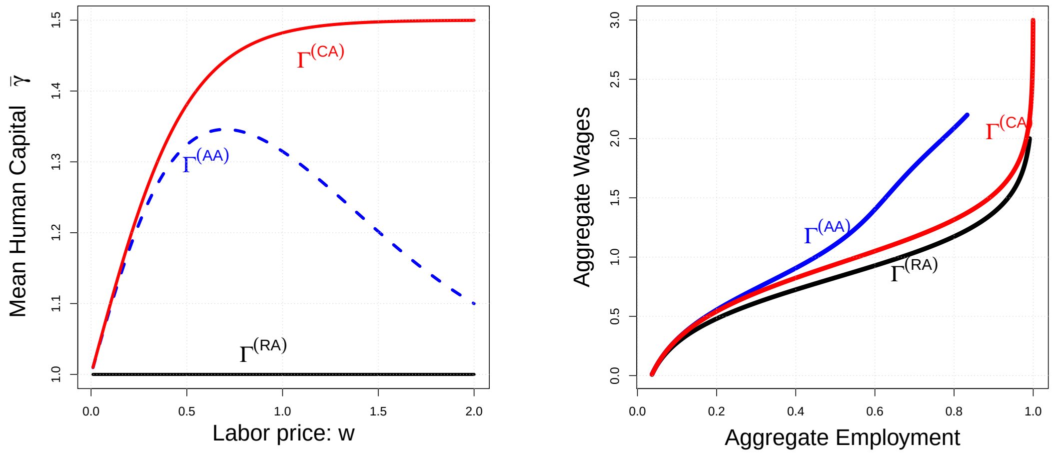

PANEL A: MEAN HUMAN CAPITAL $w_1 = w_2 = w$

PANEL B: AGG. EMPLOYMENT-WAGE SCHEDULE; $w_1 = w_2 = w$

Notes: Figure presents the behavior of the aggregate employment and wages after movements in occupational labor prices w_k in a two occupation, two-type labor supply model. Both panels plot the implied movements when the price of labor in occupation 1 is constrained to equal the price in occupation 2. Panel A plots the mean human capital of employed workers $\bar{\gamma}$ against the prevailing price of labor, while Panel B plots the implied relationship between aggregate employment and wages. The solid black line is the representative agent curve with $\gamma_{jk} = 1$ for all j and k, while the blue line reports the curve when Γ has worker fixed effects. The red solid line is the curve when Γ exhibits comparative advantage. Γ matrices defined as in equation (17), while each worker type constitutes half of the market, $u(c) = c$ and $\nu_j = 0.5$

Panel A shows that as we vary the price of labor w, there is no selection in the set of workers employed: all employed workers can only supply one unit of labor, regardless of the price of labor. This produces a familiar upward-sloping relationship between aggregate employment and wages in Panel B, as would be the case in representative agent models of labor supply.

When there is absolute advantage (the blue line), the aggregate wage-employment schedule becomes relatively steep at low levels of employment. This is due to the composition effect. When the price of labor is 0, absolute advantage does not affect allocations, as both low and high-type workers are equally unlikely to work. As the price of labor increases, high-type workers disproportionately enter the labor force, leading to growing positive selection at low levels of w. This leads to higher wages than observed in the representative agent economy for low levels of employment. Eventually, most of the high-type workers are employed, so additional increases in the price of labor w have a larger marginal impact the employment of low-type workers and average human capital falls. At this point, high-type workers are mostly inframarginal to the increases in w, but still receive wage increases. As a result, increases in the price of labor generate little increase in employment for a given wage movement, yielding a steep relationship between wages and employment. As seen above, this could generate backward-bending aggregate relationships between employment and wages if increases in the price of labor induced large enough inflows of low-type workers.

A.4 Estimated Γ matrices and selection-corrected TFP series

This section reports the estimated skill matrices and selection-corrected TFP series that are used for the calibrated model of Section 4. Table A2 shows the estimated Γ matrix from the period 2002-2006. Table A3 shows the estimated selection-adjusted TFP series and the raw BLS KLEMS project TFP series. Table A4 reports the estimated Γ matrix from the period 1984-89.

TABLE A2: Estimated Γ, m_j and ξ_k, 2002-2006 CPS

	Worker type j								
Occupation k	1	2	3	4	5	6	7	8	ξ_k
1 - Routine	0.806	0.739	0.699	0.910	1.585	0.382	3.853	13.710	-2.01
2 - Low-Skill Service	0.040	0.777	0.704	1.010	1.672	2.889	4.063	3.806	-2.12
3 - Manual	1.180	0.046	0.869	1.187	2.002	0.293	1.448	16.644	-2.45
4 - Sales	0.036	0.778	0.674	0.980	1.564	2.774	3.800	12.819	-2.31
5 - Production	1.028	0.602	0.739	0.959	1.684	0.896	3.816	1.057	-2.74
6 - Clerical	0.034	0.798	0.656	1.019	1.565	2.773	3.735	12.375	-2.31
7 - Construction	1.059	0.377	0.773	0.989	1.871	0.699	4.268	1.577	-2.92
8 - Tradespeople	1.064	0.035	0.769	1.039	1.781	2.740	3.853	1.929	-2.99
9 - Supervisors	0.669	0.732	0.629	0.891	1.438	2.452	3.269	10.511	-2.68
10 - Technicians	0.865	0.718	0.627	0.943	1.539	2.381	3.197	1.206	-3.24
11 - Social Skilled	0.031	0.858	0.767	1.036	1.574	2.657	3.567	3.530	-2.88
12 - Medical	0.028	0.920	0.771	1.085	1.588	2.652	1.142	10.787	-3.33
13 - Computing	0.659	0.766	0.660	0.926	1.434	2.331	2.929	9.223	-3.53
14 - Engineers	0.731	0.905	0.719	0.125	1.677	2.662	3.392	3.601	-3.90
15 - Business Services	0.053	0.844	0.711	1.030	1.552	2.570	3.314	10.273	-3.17
m_j	0.143	0.223	0.288	0.120	0.154	0.045	0.023	0.004	–
$\mathbb{E}_k[\gamma_{jk}]$	0.552	0.660	0.718	0.942	1.635	2.077	3.310	7.537	–
$Var_k(\gamma_{jk})$	0.211	0.080	0.004	0.057	0.024	0.926	0.797	28.329	–

Notes: Table reports the estimated matrix of skills Γ, mass of worker types m_j for the period 2002-2006. A cell (k, j) in the matrix reports the estimated units of human capital that a worker of type j supplies to occupation k on average. The final column reports the net non-pecuniary benefits of each occupation ξ_k. The final three rows report the mass of each worker type, the mean of each type's skill vector (column of the Γ matrix), variance of each type's skill vector. Estimation procedure laid out in Section 4.1, and carried out using data from 2002-2006 in the CPS.

TABLE A3: TFP Series: Annual Percentage Changes in the Raw BLS Multifactor Productivity Series Versus Series Adjusted for Human Capital Selection

NAICS		1990-1991		2008-2009	
Code	Sector Title	BLS Raw	Adjusted	BLS Raw	Adjusted
211	Oil and gas extraction	0.9	-0.3	22.6	-3.8
212	Mining, except oil and gas	-0.0	-2.2	-5.9	-5.2
221	Utilities	-1.4	-0.7	3.8	-0.5
230-238	Construction	-0.5	-2.9	0.0	-6.0
311-312	Food and beverage and tobacco products	-0.8	-0.8	0.7	-2.2
315-316	Apparel and leather and allied products	4.2	-1.0	-19.7	-7.4
322	Paper products	0.1	-4.0	3.5	-5.2
323	Printing and related support activities	-0.6	-3.8	-3.4	-5.7
324	Petroleum and coal products	3.3	0.1	-6.4	-0.3
325	Chemical products	-1.9	-0.1	-1.4	-1.8
326	Plastics and rubber products	1.3	-0.4	3.3	-11.8
331	Primary metals	-0.6	-1.6	1.0	-5.5
332	Fabricated metal products	-1.8	-2.0	-7.5	-3.9
333	Machinery	-5.5	-1.0	-4.0	-3.0
334	Computer and electronic products	3.8	-0.5	3.4	-0.4
335	Electrical equipment/appliances/components	-3.8	-3.4	-4.7	-1.0
336	Transportation equipment manufacturing	-0.8	-0.4	-10.6	-3.6
339	Miscellaneous manufacturing	-1.0	-1.1	2.4	-4.3
42	Wholesale trade	4.8	-3.0	-4.0	-4.5
44,45	Retail trade	0.8	-2.8	0.4	-2.8
484	Truck transportation	3.7	-4.1	-0.0	-5.7
486-492	Other transportation and support activities	3.8	-6.7	-6.0	-4.3
511	Publishing, except internet (includes software)	-1.2	-1.9	-2.4	-0.2
515,517	Broadcasting and telecommunications	-0.2	-4.0	-3.5	-2.0
516-519	Data processing and other information services	-3.2	-5.7	2.5	-1.9
524	Insurance carriers and related activities	2.5	-6.8	1.7	-15.8
531	Real estate	-1.3	-1.2	-0.3	-1.2
532,533	Leasing services and lessors of intangible assets	-5.1	-3.5	-6.4	-0.4
541	Professional, scientific, and technical Services	-2.7	-5.9	-2.9	-5.3
561	Administrative and support services	-2.6	-5.5	0.1	-1.7
611	Educational services	4.6	-1.5	5.0	6.4
621	Ambulatory health care services	-1.7	-2.5	-0.4	0.7
622,623	Hospitals and nursing/residential care facilities	-0.5	-0.0	1.3	1.3
721	Accommodation	2.0	-0.8	-4.0	-2.8
722	Food services and drinking places	-2.0	-1.5	-1.6	-3.4
811-813	Other services, except government	-1.4	-4.6	-1.5	-5.2
	Aggregate	-1.1	-0.5	-4.2	-5.9

Notes: BLS Raw series taken from the BLS' Multifactor Productivity Series project. Adjusted series accounts for selection in the human capital levels of employed workers according to equation (24). Aggregate TFP constructed as the mean of sector TFP series, weighted by value-added in each sector. The table excludes the 15 sectors which were among the 20 smallest sectors in both 1990 and 2008, measured by value added. "Data processing and other information services" includes NAICS codes 516, 218, and 519.

TABLE A4: Estimated Γ, m_j and ξ_k, 1984-1989 CPS

	Worker type j								
Occupation k	1	2	3	4	5	6	7	8	ξ_k
1 - Routine	0.855	0.684	0.807	0.090	1.341	1.926	2.552	0.480	-2.17
2 - Low-Skill Service	0.121	0.749	0.921	0.188	1.485	0.628	2.826	6.610	-2.46
3 - Manual	1.037	0.483	0.858	0.089	1.419	2.222	2.384	5.697	-2.67
4 - Sales	0.125	0.706	0.853	1.426	1.036	0.222	2.395	5.671	-2.64
5 - Production	1.213	0.298	0.888	0.093	1.576	2.508	2.567	2.647	-2.94
6 - Clerical	0.107	0.689	0.850	1.384	1.119	0.348	2.369	5.546	-2.37
7 - Construction	1.057	0.426	0.774	0.079	1.420	2.117	2.650	5.977	-3.69
8 - Tradespeople	1.148	0.402	0.103	0.780	1.446	2.264	2.502	5.680	-3.18
9 - Supervisors	0.411	0.642	0.665	1.303	1.119	0.926	2.303	5.091	-3.15
10 - Technicians	0.119	0.490	0.747	1.316	1.295	1.944	2.367	4.858	-3.40
11 - Social Skilled	0.066	0.798	0.806	1.518	1.001	0.325	2.385	1.511	-3.24
12 - Medical	0.071	0.716	1.002	1.561	0.766	0.610	2.321	5.467	-3.54
13 - Computing	0.077	0.661	0.123	1.502	1.228	1.904	2.490	4.886	-3.88
14 - Engineers	0.065	0.575	0.514	1.507	1.067	1.914	2.517	4.860	-4.41
15 - Business Services	0.060	0.680	0.763	1.390	1.092	0.988	2.375	4.953	-3.21
m_j	0.118	0.325	0.124	0.128	0.143	0.041	0.114	0.006	–
$\mathbb{E}_k[\gamma_{jk}]$	0.435	0.600	0.712	0.948	1.227	1.390	2.467	4.662	–
$Var_k(\gamma_{jk})$	0.223	0.021	0.072	0.411	0.050	0.680	0.020	2.996	–

Notes: Table reports the estimated matrix of skills Γ, mass of worker types m_j for the period 1984-1989. A cell (k, j) in the matrix reports the estimated units of human capital that a worker of type j supplies to occupation k on average. The final column reports the net non-pecuniary benefits of each occupation ξ_k. The final three rows report the mass of each worker type, the mean of each type's skill vector (column of the Γ matrix), and variance of each type's skill vector. Estimation procedure laid out in Section 4.1, and carried out using data from 1984-1989 in the CPS.

TABLE A5: In-Sample Model Fit, 2002-2006

	Emp. Shares		Mean Log Wage		SD Log Wage	
	Model	Data	Model	Data	Model	Data
	(1)	(2)	(3)	(4)	(5)	(6)
Non-Employed	20.45	23.1	–	–	–	–
1 Routine	10.04	10.20	9.62	9.60	0.81	0.80
2 Low-Skill Service	5.19	4.78	9.65	9.62	0.85	0.85
3 Manual	3.84	3.54	9.85	9.85	0.71	0.70
4 Salespeople	5.41	4.87	9.84	9.82	0.82	0.81
5 Production	4.16	3.86	10.03	10.03	0.70	0.69
6 Clerical	9.12	8.49	10.04	10.02	0.76	0.74
7 Construction	1.86	1.64	10.10	10.06	0.76	0.77
8 Tradespeople	3.23	3.19	10.20	10.19	0.66	0.65
9 Supervisors	7.35	6.82	10.22	10.20	0.85	0.85
10 Technicians	2.73	2.47	10.38	10.39	0.64	0.63
11 Social Skilled	7.26	7.81	10.17	10.22	0.92	0.90
12 Medical	4.84	5.29	10.45	10.51	0.82	0.80
13 Computing	3.07	2.99	10.59	10.62	0.76	0.73
14 Engineers	1.73	1.67	10.82	10.85	0.65	0.61
15 Business Services	9.72	9.28	10.67	10.72	0.84	0.83
Correlation: Model to Data	0.99		1.00		0.98	

Notes: Table reports the in-sample fit of the estimated model for the period 2002-2006. Columns 1 and 2 report employment shares in each of the 15 occupations and the non-employment rate implied by the model and in the data, respectively. Columns 3 and 4 similarly report the mean log wage, while columns 5 and 6 report the standard deviation of log wages. The final row reports the correlation of model quantities to data quantities at the occupation level.

A.5 In-Sample Model Fit

This section presents measures of model fit in-sample. Table A5 presents the employment shares (columns (1) and (2)), average log wage (columns (3) and (4)) and standard deviations of log wages (columns (5) and (6)) for each of the 15 occupation clusers over the estimation period 2002-06. Odd columns report model-implied numbers, while even columns report data numbers. The model matches the occupation-specific employment shares and wage distribution of the data almost exactly: the correlations between the model and data columns are very close to 1. Table A4 shows a similar ability to match the employment and wage data in the 1984-89 estimation window.

The model also has predictions over the probability of occupation switching. Figure A4 plots the data's occupation switching probabilities against those implied by the model for the period 2002-06. Each dot is a (k, k') pair. The figure plots the share of workers who switch to k' given that they switch out of occupation k. The horizontal

TABLE A6: In-Sample Model Fit, 1984-1989

	Emp. Shares		Mean Log Wage		SD Log Wage	
	Model	Data	Model	Data	Model	Data
	(1)	(2)	(3)	(4)	(5)	(6)
Non-Employed	26.04	27.72	–	–	–	–
1 Routine	9.12	9.11	9.58	9.58	0.81	0.81
2 Low-Skill Service	4.53	4.34	9.55	9.55	0.84	0.84
3 Manual	5.12	4.91	9.82	9.82	0.74	0.74
4 Salespeople	5.04	4.81	9.73	9.72	0.78	0.78
5 Production	4.52	4.35	10.03	10.05	0.71	0.71
6 Clerical	9.89	9.61	9.86	9.86	0.74	0.73
7 Construction	1.35	1.22	10.04	10.02	0.79	0.80
8 Tradespeople	3.64	3.61	10.13	10.13	0.71	0.70
9 Supervisors	4.62	4.36	10.13	10.12	0.83	0.83
10 Technicians	3.69	3.49	10.34	10.36	0.66	0.63
11 Social Skilled	5.92	6.16	10.14	10.16	0.84	0.83
12 Medical	3.46	3.60	10.25	10.28	0.77	0.75
13 Computing	2.25	2.13	10.42	10.43	0.70	0.68
14 Engineers	1.78	1.74	10.72	10.74	0.61	0.57
15 Business Services	9.03	8.84	10.47	10.51	0.78	0.76
Correlation: Model to Data	1.00		1.00		0.99	

Notes: Table reports the in-sample fit of the estimated model for the period 1984-1989. Columns 1 and 2 report employment shares in each of the 15 occupations and the non-employment rate implied by the model and in the data, respectively. Columns 3 and 4 similarly report the mean log wage, while columns 5 and 6 report the standard deviation of log wages. The final row reports the correlation of model quantities to data quantities at the occupation level.

axis is the model-implied swicthing probability while the vertical axis is the switching probability in the data. The correlation between model and data is very high at 0.84. Indeed, a regression line of data switching probability on model switching probability has slope 0.99 with R^2 of 0.71. The model is a good predictor of source-destination switching pairs. However, the model implies an overall switching probability which is roughly twice that in the data so that gross flows are too large, even if net flows are well captured. This is a limitation of the largely standard assumption that ζ_{ikt} are i.i.d.

A.6 Reduced Form Evidence: Changes in Composition Effects

A strong composition effect is central to the model's ability to match the labor market dynamics of the Great Recession. This section provides reduced form evidence that shifting composition effects account for the change in the cyclicality of real wages and were unusually important during the Great Recession. I first estimate the relationship

FIGURE A4: In-Sample Fit: Occupation Switching Patterns, 2002-2006

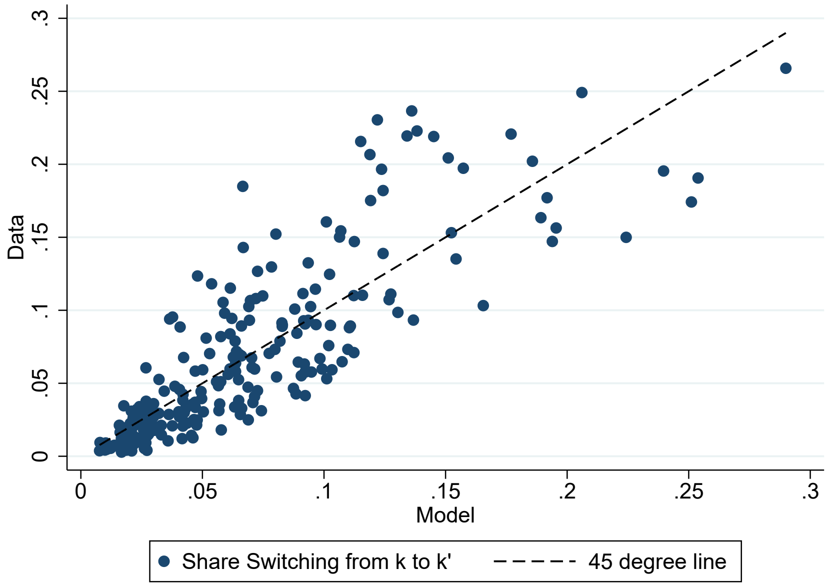

Notes: Figure shows the model fit of occupation switching patterns. Each dot represents a pair of occupations (k, k') and the axes plot the share of occupation switchers who are in that (k, k') pair in the estimated model (horizontal axis) and the data (vertical axis). The dashed line shows a 45 degree line.

between year-over-year GDP growth rates and year-over-year changes in aggregate real wages, allowing their cyclicality to differ every 20 years. Specifically, I estimate time series regressions of the form

$$\text{(A2)} \qquad \Delta \ln W_t = \sum_p \left[\alpha_p + \beta_p \Delta \ln GDP_t\right] \mathbf{1}\{t \in p\} + \epsilon_t$$

with OLS, where t is a quarter, Δ is the year-over-year change operator, W_t is real average hourly earnings of production workers, p indexes 20-year periods (e.g. 1960-1979, 1980-1999 etc.), GDP_t is quarterly real GDP per capita, and ϵ_t is an error term. The coefficient β_p represents the cyclicality of the wages in period p. Average hourly earnings of production and non-supervisory employees come from the Current Employment Statistics and are deflated by core CPI. GDP per capita is from the BEA and deflated by the GDP price deflator.

The estimated cyclicalities of aggregate wages β_p are presented in Panel A of Figure A5, along with 90% confidence intervals calculated using Newey-West standard errors allowing for 10-year autocorrelation. The figure shows a stark decline in real aggregate wage cyclicality. While the elasticity of real aggregate wage growth to GDP growth was 0.50 (standard error: 0.07) in the 1960s and 70s, this fell to 0.31 (SE: 0.15) from 1980-1999. Strikingly, aggregate wages have become countercyclical since 2000, with

FIGURE A5: The Cyclicality of Real Wages in Aggregate and Controlling for Composition

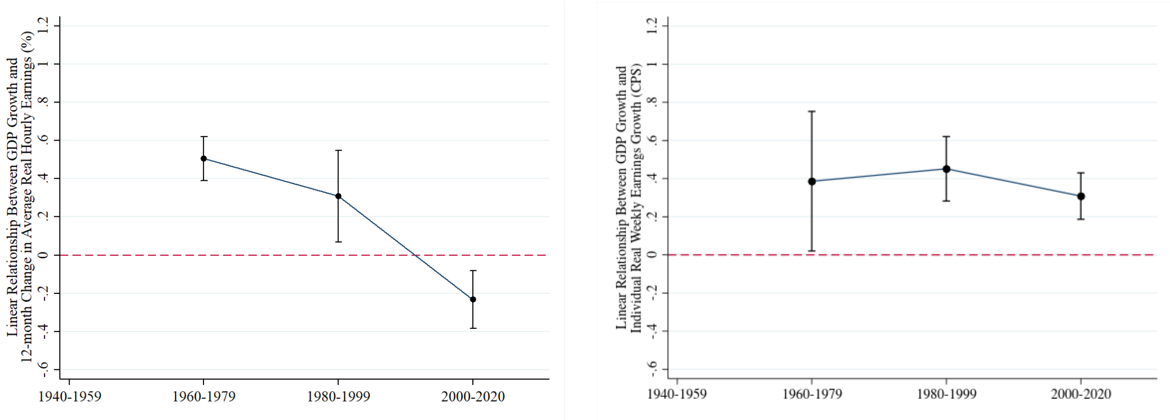

PANEL A: REAL AVERAGE HOURLY EARNINGS

PANEL B: COMPOSITION-ADJUSTED AVERAGE WAGES

Notes: Figure presents estimates of wage cyclicality in the U.S. for twenty-year time periods going back to the 1940s. It shows the estimated linear relationship between growth rates in real wages and year-over-year real GDP per capita growth rates. Panel A plots the correlation with real average hourly earnings from the BLS by estimating β_p from Equation (A2) using OLS. Panel B presents estimates of wage cyclicality after controlling for composition bias. It shows the estimated linear relationship between year-over-year real GDP per capita growth rates and individual growth rates in real weekly earnings in the Current Population Survey by estimating $\tilde{\beta}_1$ from Equation (A4) following Solon et al. (1994) (SBP). Bars represent 90% confidence intervals. Panel A uses Newey-West standard errors allowing for up to 10 years of autocorrelation, while Panel B clusters standard errors at the year level. Wages are deflated by core CPI, while GDP per capita is deflated by the GDP deflator.

an elasticity of wage growth to per capita GDP growth of -0.23 (SE: 0.09). This trend continues past 2020, given the large increase in aggregate wages observed during the COVID-19 pandemic (Cajner et al., 2020).

To assess the role of composition effects for the decline in wage cyclicality, I follow the selection-correction method of Solon et al. (1994). Specifically, assume the following statistical model for individual wages

$$\ln \omega_{it} = \alpha_i + \tilde{\beta}_p \ln GDP_t + \eta_0 \cdot t + \eta_1 \cdot t^2 + \eta_2 X_{it} + \epsilon_{it} \tag{A3}$$

where t is an aggregate time trend, and X_{it} is a control for worker experience. The worker fixed effect α_i is the source of the composition bias in the aggregate statistics. If the selection of workers employed during a recession have higher α_i on average than those employed during a boom, then the estimate of $\tilde{\beta}_p$ will be biased upward in aggregate data. By estimating equation (A3) in first differences, one controls for characteristics of a worker which are fixed over time, such as the worker's permanent ability. Therefore estimating

$$\Delta \ln \omega_{it} = \tilde{\eta}_0 + \tilde{\beta}_p \Delta \ln GDP_t + \tilde{\eta}_1 \cdot t + \tilde{\eta}_2 X_{it} + \Delta\epsilon_{it}. \tag{A4}$$

where ΔZ_t represents the change in a variable Z between $t-1$ and t, yields a consistent estimate of the cyclicality of wages $\tilde{\beta}_p$. I thus estimate this specification using microdata from the March Supplement of the Current Population Survey (CPS), measuring individual wages as real weekly earnings. As above, I allow the relationship between the aggregate cycle and selection-corrected wages to vary across 20-year periods p. Note that linking workers across time in CPS March Supplement is far easier after 1977, so the 1960s-70s period should not be heavily interpreted.

Panel B of Figure A5 shows that the cyclical behavior of wages is much more stable after controlling for composition bias. It presents estimates of $\tilde{\beta}_p$ and 90% confidence bands constructed using standard errors clustered at the year level. In the 1960s and 70s, the elasticity of individual real wage growth to GDP per capita growth was 0.39 (SE: 0.22) after controlling for selection, and 0.41 in the 1980s and 90s.[43] This elasticity fell only slightly to 0.31 (SE: 0.07) by the 2000-2020 period. The relative stability of this coefficient through time, coupled with the sharp decline in aggregate wage cyclicality, provides reduced form evidence that cyclical compositional shifts in the workforce have become especially important over the last twenty years.

Figure A6 provides reduced form evidence for strong composition effects specifically during the Great Recession. Panel A plots a bin-scatter plot in which the x-axis is an occupation's average weekly earnings in 2008 according to the OES, while the vertical axis is the percentage employment change for the occupation between 2008 and 2009. Each dot is a different percentile of the occupational average weekly earnings distribution. It shows that low-wage occupations were significantly more likely to experience large employment declines, reflecting a net reallocation of workers towards high-wage occupations. A one-standard deviation decrease in an occupation's average earnings is associated with a 1 percentage point larger decline in employment in that occupation. This is one of the core mechanisms of the model: a net reallocation away from occupations k with low wages towards those with high wages.

The second core mechanism in the model is a selection across different types of workers j within an occupation. Panel B of Figure A6 shows this selection. It is constructed by regressing individual weekly earnings and 12-month separation rates on occupation fixed effects in the 2008 outgoing rotations group of the CPS, and plotting a binned scatter of the residuals from this regression. The figure shows that those who are paid

[43] Some estimates are noisy due to relatively small sample sizes in the early iterations of the CPS.

FIGURE A6: Composition Effects Within and Between Occupations, 2008-2009

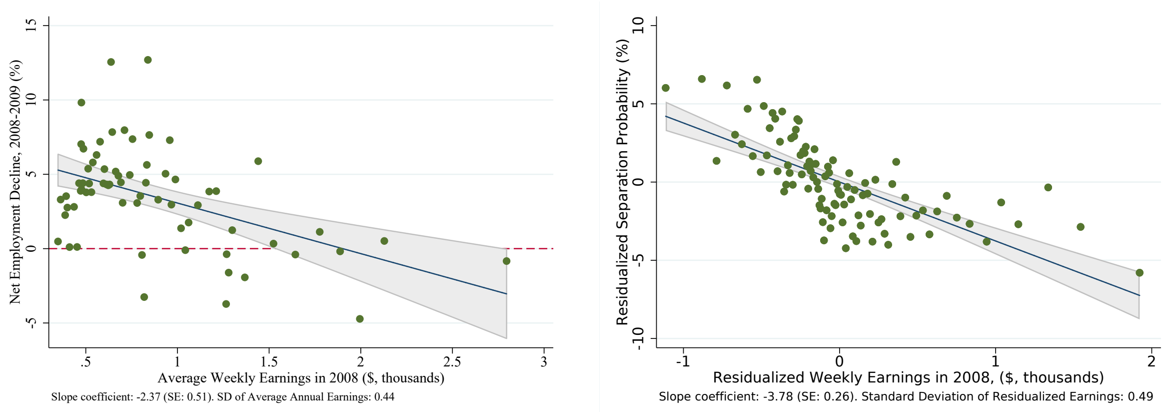

PANEL A: BETWEEN EFFECTS PANEL B: WITHIN EFFECTS

Notes: Figure plots relationship between 2008 weekly earnings and 2008-09 occupational employment changes (Panel A) or 12-month separation probabilities (Panel B). Panel A uses data from the OES, while Panel B uses data from the Outgoing Rotation Group of the CPS, and residualizes both earnings and separation rates against 6-digit SOC code fixed effects.

less relative to others in their occupation in 2008 were more likely to be non-employed 12 months later. A one standard deviation increase in residualized weekly earnings was associated with a 1.8 percentage point (pp) decline in the residualized probability of separation between 2008 and 2009. This is unusually large: the equivalent number for the 1990-91 recession – the last recession to feature declining real wages – is 1.6pp. The evidence presented here shows that, consistent with the model's mechanism, composition effects were important both within and across occupations during the Great Recession: that is, low-wage occupations saw declining employment and lower wage workers within occupations were more likely to separate.

Furthermore, the BLS' Employment Cost Index (ECI), which builds a wage index holding fixed the share of industry-occupation cells in employment, shows declines in real wages during the Great Recession. Although the ECI only begins in 2000 and is thus not well-suited to studying how composition effects have changed over time, the fact that its movements during the Great Recession are so different to those of average hourly earnings further suggests the importance of composition effects at this time.

A.7 Changes in Skill Distribution Over Time

This section further examines the estimated changes in the skill distribution between 1984-89 and 2002-06, studied in section 4.4.1. Consider two key variances in the skill

FIGURE A7: Absolute and Comparative Advantage: 1984-1989 and 2002-2006

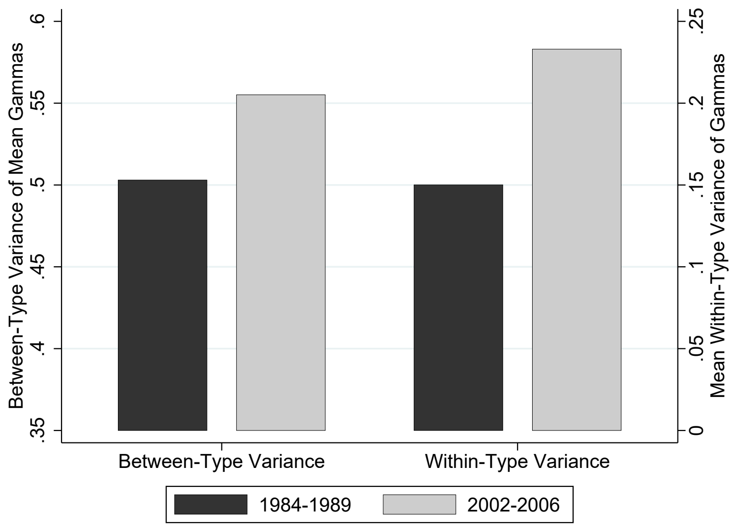

Notes: Figure plots the estimated within and between type variance of skills in the economy, captured by the Γ matrix of Table A2 and Appendix Table A4. Estimation follows the procedure outlined in Section 4.1, and carried out separately in the CPS March Supplement for the periods 1984-89 (gray bars) and 2002-2006 (black bars). Within and between variance defined as in equation (A5).

distribution:

$$\text{(A5)} \qquad Var^{BTWN} := \sum_{j=1}^{J} m_j(\mathbb{E}_k[\gamma_{jk}] - 1)^2; \qquad Var^{WTHN} := \sum_{j=1}^{J} m_j Var_k(\gamma_{jk}),$$

Figure A7 plots the within and between variance of skills in the economy prior to the 1991 and 2008 recessions. The cross-type "between" variance is informative about the difference in mean skill for various workers, while the within-type variance is a measure of skill specificity. Between-type variance is plotted against the left axis while within-type variance is plotted against the right axis. The black bars represent the estimation period 1984-1989, while the gray bars represent the period 2002-2006.

I estimate that the cross-type variance of γ_{jk} has increased from 0.50 to 0.56, an increase of 10.4% in the 20 years leading up to the Great Recession. Since the mean of γ_{jk} is one within each occupation, these variances may be interpreted as the coefficient of variation of skills, squared. There is an even larger increase in within-type variance, while the mean variance of the γ_{jk} vectors was 0.15 in the late 1980s, it was 0.23 prior to the 2008 recession, an increase of 55.2%. This suggests both that skills have become more specific and that the gap between the highest- and lowest-skill workers has grown.

The majority of the variance of skills is across types, rather than within types. In the

FIGURE A8: Correlation of Skill Relatedness in Γ with Euclidean Skill Distance in O*NET

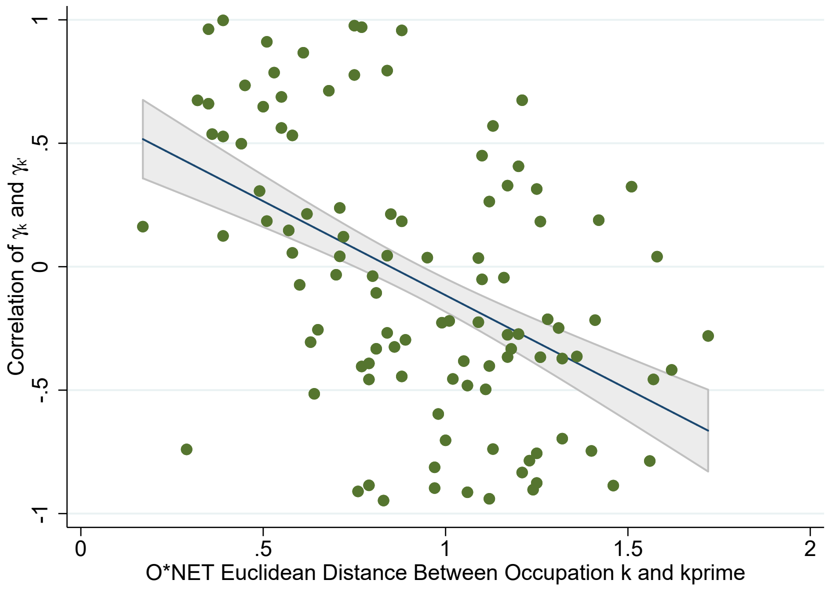

Notes: Figure compares the structurally-estimated skill transferability from the model to a common measure of skill relatedness from O*NET. Each dot corresponds to a pair of occupation clusters (k, k'). Occupations clustered by O*NET skill and knowledge vectors within terciles of the share with at least some college education. The horizontal axis reports the Euclidean distance between skill vectors in O*NET. The vertical axis reports the correlation of row vectors in Γ in 2002-2006, as in Panel B of Figure A9. Line of best fit reported, with shaded area representing 95% confidence interval using White heteroskedasticity robust standard errors.

1980s, cross-type variance accounted for 85% of total skill variance, while within-type variance accounts for 25%. In the 2000s, cross-type variance accounted for 76% of total variance, with within-type variance accounting for 31%. In both periods, this indicates a negative covariance between within-type variance and mean skill, suggesting that low-skill workers have more specific skills. This negative covariance arises because low-wage workers rarely enter high-skill occupations, such as engineering or skilled business services.

Next, I consider changes in the cross-occupation skill transferability. In the model, a natural proxy for skill transferability between any two occupations is the correlation of the row vectors of the Γ matrix. For instance, if the correlation between the Manual and Production occupations' γ vectors is high, it suggests that workers who have high-skills in Manual occupations tend to also have skills in Production occupations. Put differently, workers who are good at manual labor may easily transfer their skills to production occupations to be serviceable welders or machinists. Indeed, this covariance appears in the expression for cross-price labor supply elasticities (15).

FIGURE A9: Correlation of Estimated Occupation Skills: 2002-2006

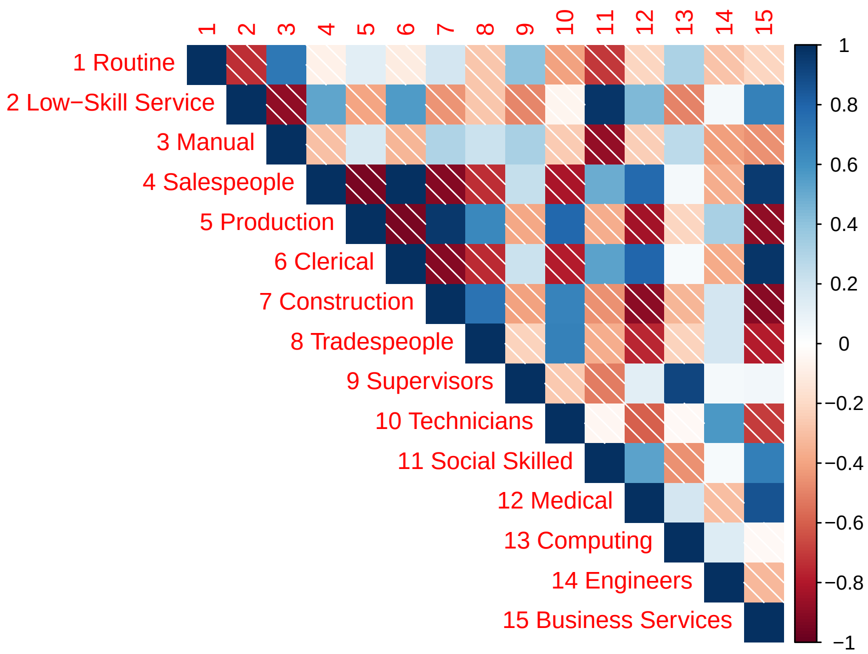

Notes: Figure plots the correlation of the row vectors of the estimated Γ, normalized by workers' mean skill in each occupation. Estimation follows procedure outlined in Section 4.1, and carried out separately in the CPS March Supplement for the periods 1984-89 (Panel A) and 2002-2006 (Panel B). Blue squares indicate that the correlation of skills between occupations is positive, while red checked squares indicate a negative correlation. Darker colors indicate that the magnitude of the correlation is closer to 1.

I calculate these cross-occupation skill correlations in both the 1984-89 and 2002-06 estimated skill distributions. As validation, Figure A8 compares these correlations from 2002-2006 with the Euclidean distance between the clusters' O*NET skill vectors, a measure of skill distance employed by Poletaev and Robinson (2008) among others. The distance between clusters in O*NET negatively predicts the correlation between occupational human capital in the Γ matrix, with a correlation coefficient of -0.48.

Figure A9 plots a correlogram of the Γ matrix's row vectors. Before calculating the correlations, I divide each element of Γ by the mean γ for type j workers, so that absolute advantage does not dominate the correlations. Panel A reports the correlation of skills in the 1984-1989 period, while Panel B plots the same correlation for the estimation sample 2002-2006. Each row and column of the correlogram correspond to one of the 15 occupations used for estimation. Blue squares in the figure indicate that the correlation of skills between occupations is positive, while red checked squares indicate a negative correlation. Deeper colors indicate that the magnitude of the correlation is closer to 1.

The majority of the correlations are highly intuitive. For instance, routine occupations employ similar skills to manual, production, and construction occupations, but have low correlations with business service occupations. Similarly, engineers are strong technicians or computer workers in both periods, while salespeople are adept in low-skill service, clerical, social skilled, and business service occupations.

One noteworthy outlier is the medical field, which appears to have correlated skills with clerical, social, sales, and business services occupations. Intuitively, medical occupations should be highly specialized, with relatively low correlations throughout the matrix. The fact that it is not is instructive to the variation used to identify the Γ matrix. Since the matrix is principally identified using information on occupation switchers, the skill correlations will tilt towards those who switch occupations. The medical workers who switch occupations are principally nurses and medical technicians, for whom soft skills may be more valuable than they are for surgeons. Framed in this way, it is unsurprising that job-switchers out of medical professions tend to have similar skills to teachers and salespeople.

APPENDIX B. MODEL APPENDIX

This subsection derives equations (15), (19) and (20). Suppressing t and differentiating the sum of the occupation choice probabilities from equation (6) with respect to price of labor w_k in k yields the response of type j employment to w_k:

$$\frac{d\tilde{E}_j}{dw_k} = \frac{\exp\left(\frac{u(c_{jk})+\xi_k}{\nu_j}\right)\frac{u'(c_{jk})\gamma_{jk}}{\nu_j}}{1+\sum_{k'\neq 0}\exp\left(\frac{u(c_{jk'})+\xi_{k'}}{\nu_j}\right)} - \frac{\left[1+\sum_{k'\neq 0}\exp\left(\frac{u(c_{jk'})+\xi_{k'}}{\nu_j}\right)-1\right]\exp\left(\frac{u(c_{jk})+\xi_k}{\nu_j}\right)\frac{u'(c_{jk})\gamma_{jk}}{\nu_j}}{\left[1+\sum_{k'\neq 0}\exp\left(\frac{u(c_{jk'})+\xi_{k'}}{\nu_j}\right)\right]^2}$$

$$= \left(\frac{u'(c_{jk})\gamma_{jk}}{\nu_j}\right)\left(\frac{\exp\left(\frac{u(c_{jk})+\xi_k}{\nu_j}\right)}{1+\sum_{k'\neq 0}\exp\left(\frac{u(c_{jk'})+\xi_{k'}}{\nu_j}\right)}\right)\left(\frac{1}{1+\sum_{k'\neq 0}\exp\left(\frac{u(c_{jk'})+\xi_{k'}}{\nu_j}\right)}\right)$$

(A6)

$$= \left(\frac{u'(c_{jk})\gamma_{jk}}{\nu_j}\right)E_{jk}(1-\tilde{E}_j)$$

Noting that employment probabilities only depend on labor demand shocks $\mathbf{z}$ through occupation prices w_k, we can write the response of aggregate employment as:

$$\begin{aligned}
\frac{d\ln\bar{E}}{d\ln z} &= \sum_j\left(\frac{m_j\tilde{E}_j}{\bar{E}}\right)\sum_{k\neq 0}\frac{d\ln\tilde{E}_j}{d\ln w_k}\cdot\frac{d\ln w_k}{d\ln\mathbf{z}}\\
&= \sum_j\left(\frac{m_j\tilde{E}_j}{\bar{E}}\right)(1-\tilde{E}_j)\tilde{\omega}_j\sum_{k\neq 0}\left(\frac{u'(c_{jk})}{\nu_j}\right)\left(\frac{\omega_{jk}E_{jk}}{\tilde{\omega}_j\tilde{E}}\right)\frac{d\ln w_k}{d\ln\mathbf{z}}
\end{aligned}\tag{A7}$$

which is equation (19) in the text. This uses the fact that $d\ln x = dx/x$ for any variable x, multiplies and divides by $\tilde{\omega}_j$, and collects terms.

We now build up the response of wages. First note that the response of aggregate wages can be written as

(A8)

$$\frac{d\ln\bar{\omega}}{d\ln\mathbf{z}} = \sum_j\left(\frac{m_j\tilde{E}_j}{\bar{E}}\right)\left(\frac{\tilde{\omega}_j}{\bar{\omega}}\right)\left(\underbrace{\frac{d\ln\tilde{E}_j}{d\ln\mathbf{z}}-\frac{d\ln\bar{E}}{d\ln\mathbf{z}}}_{\text{Composition Effect}}+\sum_k\underbrace{\left(\frac{E_{jk}\omega_{jk}}{\tilde{E}_j\tilde{\omega}_j}\right)}_{\substack{\text{Earnings}\\ \text{Share of } k\\ \text{amongst type } j}}\left(\underbrace{\frac{d\ln E_{jk}}{d\ln\mathbf{z}}-\frac{d\ln\tilde{E}_j}{d\ln\mathbf{z}}}_{\text{Reallocation effect}}+\underbrace{\frac{d\ln\omega_{jk}}{d\ln\mathbf{z}}}_{\text{Direct Effect}}\right)\right)$$

By assumption, $\omega_{jk} = \gamma_{jk}w_k$, so that the direct effect $d\ln\omega_{jk}/d\ln\mathbf{z} = d\ln w_k/d\ln\mathbf{z}$ as in the text. The composition effect is:

$$\begin{aligned}
\text{Composition effect} &= \sum_j\frac{m_j\tilde{\omega}_j\tilde{E}_j}{\bar{\omega}\bar{E}}\frac{d\ln\tilde{E}_j}{d\ln z}-\frac{d\ln\bar{E}}{d\ln z}\\
&= \sum_j\frac{m_j\tilde{\omega}_j\tilde{E}_j}{\bar{\omega}\bar{E}}(1-\tilde{E}_j)(\tilde{\omega}_j-\bar{\omega})\sum_k\left(\frac{\omega_{jk}E_{jk}}{\tilde{\omega}_j\tilde{E}_j}\right)\left(\frac{u'(c_{jk})}{\nu_j}\right)\frac{d\ln w_k}{d\ln z}
\end{aligned}$$

which is the expression in equation (20). This derivation involves plugging in equations (A6) and (A7) into equation (19) and rearranging.

Finally to calculate the reallocation effect, differentiate the occupation choice probabilities from equation (6) with respect to $w_{k'}$:

$$\begin{aligned}
\frac{d\ln E_{jk}}{dw_{k'}} &= \left[\gamma_{jk'}\left(\frac{u'(c_{jk'})}{\nu_j}\right)\right]\mathbf{1}\{k'=k\} - \frac{\exp\left(\frac{u(c_{jk'})+\xi_{k'}}{\nu_j}\right)\gamma_{jk'}u'(c_{jk'})/\nu_j}{1+\sum_{\tilde{k}\neq 0}\exp\left(\frac{u(c_{j\tilde{k}})+\xi_{\tilde{k}}}{\nu_j}\right)}\\
&= \gamma_{jk'}\left(\frac{u'(c_{jk'})}{\nu_j}\right)(\mathbf{1}\{k'=k\}-E_{jk'})
\end{aligned}$$

Plugging this and equation (A6) into equation (A8) and rearranging obtains the result.

Now we derive the elasticities of labor supply to each occupation k. We have

$$\begin{aligned}\frac{d\ln L_k}{d\ln w_{k'}} &= \frac{d\left(\sum_j m_j\gamma_{jk}E_{jk}\right)}{d\ln w_{k'}}\frac{1}{\left(\sum_j m_j\gamma_{jk}E_{jk}\right)} \\ &= \frac{1}{\left(\sum_j m_j\gamma_{jk}E_{jk}\right)}\sum_j m_j E_{jk}\gamma_{jk}\frac{d\ln E_{jk}}{d\ln w_{k'}} \\ &= \sum_j \frac{1}{L_k}\gamma_{jk}\gamma_{jk'}w_{k'}\left(\frac{u'(c_{jk'})}{\nu_j}\right)(\mathbf{1}\{k'=k\}-E_{jk'})E_{jk}\end{aligned}$$

Appendix C. Data Appendix

This section contains additional details of the paper's data. I use the March Supplement of the IPUMS Current Population Survey (CPS) for the labor supply estimation. The CPS is a rotating panel. Respondents are surveyed for four consecutive months, followed by an eight-month hiatus, before being surveyed again for the subsequent four months. For example, if an individual is first surveyed in January 2005, they will be surveyed between January and April in both 2005 and 2006.

The CPS contains information on individuals' employment status, demographics, and educational attainment at a monthly frequency. A supplemental survey - the Annual Social and Economic Supplement - which solicits additional information on respondents income sources and hours is administered every March. I restrict attention to the sample of individuals who are between the age of 21 and 60 years old in both years in which they are surveyed. I include both men and women in the analysis. I drop workers who earn positive labor income that is less than \$1,000 in a given year, fearing that these records may suffer from undue measurement error. I drop individuals living in group quarters, retired workers, those serving in the armed forces and employed workers with missing wage information.

I harmonize all sector codes to the 2010 NAICS coding using the crosswalks of provided by the Census bureau, and available at `https://www.census.gov/topics/employment/sector-occupation/guidance/code-lists.html`. I harmonize occupation codings to the 2010 Standardized Occupation Classification (SOC) using Census crosswalks, available from the same location. Much of the work to generate this crosswalk was performed by IPUMS, and is contained in the IPUMS CPS variable OCC2010.

Crucial to the estimation routine outlined in section 4.1 is the availability of panel data on earnings and occupations. Therefore, it is crucial that one is able to construct a consistent individual identifier over time using the CPS. This is not a trivial task, as highlighted by Flood and Pacas (2008). IPUMS has constructed a unique identifier for individuals for the period from 1990 onward. I follow their approach and state that two workers are the same individual in period t and $t+1$ if they: 1) share the same household identifier (IPUMS variable HRHHID), 2) share the same person number within the household (LINENO), 3) have the same race (RACE) and sex (SEX), and 4) have aged by one year between t and $t+1$ (i.e. the variable AGE in t is one less than its value in $t+1$). Using this routine, I find only 0.01% of records before 1989 have non-unique worker matches. These rare non-unique matches are dropped from the analysis. Note this imputation procedure is only relevant for the reduced form evidence in Section A.6 and for re-estimation during 1984-89 in section 4.4.1. Finally, I include only individuals for whom two years of data are available.

In addition to providing the microdata for estimation, the CPS is used to calculate employment levels in occupation-by-sector cells, which is an input into the estimation of sector-level total factor productivity series. Using the CPS, I calculate the share of employees in each 3-digit NAICS code who belong to each of the K occupation clusters. I then interact this share with the sector-level employment provided by the Bureau of Economic Analysis (BEA) to construct an estimate for the total employment in each occupation-sector cell for every year.

I use the Occupation Employment Statistics (OES) to calculate the share of sector wage bills that accrue to each occupation group α_{sk}. The OES is an employer survey conducted by the BLS which asks for total employment and wages of workers in each standardized occupation code. The survey has been run annually at the 3-digit level since 1997, and every 3 years prior. I consider the period 2003-2007 - the period immediately prior to the Great Recession - to construct the wage bill shares.

Finally, Tables A7-A9 report the results of the occupation clustering algorithm detailed in the main text. The tables list the 8 largest SOC occupations for each occupation cluster. Occupation size is measured by the total employment in the occupation as of 2013 in the OES. The mean annual income in each SOC code is also listed.

TABLE A7: Largest Employment SOC Codes within Occupation Clusters, Set 1

Cluster #	SOC Title Examples	Income
1	Cashiers	20561
Routine	Driver/Sales Workers and Truck Drivers	37017
	Combined Food Preparation and Serving Workers, Including Fast Food	19099
	Stock Clerks and Order Fillers	25190
	Nursing, Psychiatric, and Home Health Aides	24758
	Janitors and Cleaners, Except Maids and Housekeeping Cleaners	25977
	Maids and Housekeeping Cleaners	22175
	Shipping, Receiving, and Traffic Clerks	31275
2	Waiters and Waitresses	20884
Low-Skill	Receptionists and Information Clerks	27502
Service	Personal Care Aides	21242
	Inspectors, Testers, Sorters, Samplers, and Weighers	37941
	Hairdressers, Hairstylists, and Cosmetologists	27533
	Childcare Workers	21942
	Counter and Rental Clerks	27143
	Hosts and Hostesses, Restaurant, Lounge, and Coffee Shop	19683
3	Laborers and Freight, Stock, and Material Movers, Hand	26744
Manual	Miscellaneous Assemblers and Fabricators	30123
Laborers	Industrial Truck and Tractor Operators	32699
	Helpers–Production Workers	25086
	Miscellaneous Agricultural Workers	21410
	Electrical, Electronics, and Electromechanical Assemblers	31824
	Painting Workers	35751
	Machine Feeders and Offbearers	29516
4	Retail Salespersons	25376
Salespeople	Security Guards and Gaming Surveillance Officers	28015
	Health Practitioner Support Technologists and Technicians	33698
	Bartenders	21777
	Bailiffs, Correctional Officers, and Jailers	44405
	Dental Assistants	35699
	Production, Planning, and Expediting Clerks	46726
	Hotel, Motel, and Resort Desk Clerks	22027
5	Grounds Maintenance Workers	27432
Construction/	Welding, Soldering, and Brazing Workers	38874
Production	Machinists	41251
	Packaging and Filling Machine Operators and Tenders	28753
	Operating Engineers and Other Construction Equipment Operators	46164
	Production Workers, All Other	31055
	Helpers, Construction Trades	28581
	Crushing, Grinding, Polishing, Mixing, and Blending Workers	34240

Notes: Table reports the 8 SOC occupations with the largest employment within each of the 15 occupation clusters. Employment and mean annual income taken from the Occupation Employment Statistics as of 2013. Cluster labels supplied by the author. Occupations grouped using a k-means clustering algorithm based on the skill and knowledge vectors of each SOC occupation in O*NET, within terciles of share of worker with at least some college education in the CPS.

TABLE A8: Largest Employment SOC Codes within Occupation Clusters, Set 2

Cluster #	SOC Title Examples	Income
6	Secretaries and Administrative Assistants	38381
Clerical	Customer Service Representatives	33407
	Office Clerks, General	30196
	Bookkeeping, Accounting, and Auditing Clerks	37374
	Sales Representatives, Wholesale and Manufacturing	68877
	First-Line Supervisors of Office and Administrative Support Workers	53851
	Tellers	26264
	Bill and Account Collectors	34683
7	Construction Laborers	35095
Skilled	First-Line Supervisors of Construction Trades and Extraction Workers	63479
Construction	Painters, Construction and Maintenance	39887
	First-Line Supervisors of Housekeeping and Janitorial Workers	39124
	Highway Maintenance Workers	36977
	Hazardous Materials Removal Workers	42536
	Ship and Boat Captains and Operators	71295
	Locksmiths and Safe Repairers	40715
8	Maintenance and Repair Workers, General	38058
Trades-	Carpenters	45071
people	Automotive Service Technicians and Mechanics	39863
	Pipelayers, Plumbers, Pipefitters, and Steamfitters	51922
	Industrial Machinery Mechanics	49777
	Heating, Air Conditioning, and Refrigeration Mechanics and Installers	46352
	Bus and Truck Mechanics and Diesel Engine Specialists	44493
	Heavy Vehicle and Mobile Equipment Service Technicians and Mechanics	46200
9	First-Line Supervisors of Retail Sales Workers	41465
Supervisors	First-Line Supervisors of Food Preparation and Serving Workers	32078
	Teacher Assistants	25778
	Business Operations Specialists, All Other	71403
	Supervisors of Transportation and Material Moving Workers	52864
	First-Line Supervisors of Mechanics, Installers, and Repairers	63513
	Firefighters	48600
	Purchasing Agents, Except Wholesale, Retail, and Farm Products	64456
10	First-Line Supervisors of Production and Operating Workers	58373
Technicians	Electricians	53707
	Engineering Technicians, Except Drafters	56521
	Radio and Telecommunications Equipment Installers and Repairers	53719
	Telecommunications Line Installers and Repairers	52771
	Miscellaneous Plant and System Operators	58163
	Water and Wastewater Treatment Plant and System Operators	45074
	Aircraft Mechanics and Service Technicians	57481

TABLE A9: Largest Employment SOC Codes within Occupation Clusters, Set 3

Clust #	SOC Title Examples	Mean Income
11	Elementary and Middle School Teachers	56909
Social	Secondary School Teachers	58491
Skilled	Other Teachers and Instructors	36646
	Postsecondary Teachers	74068
	Special Education Teachers	58420
	Designers	46437
	Lawyers	126710
	Human Resources Workers	61057
12	Registered Nurses	68801
Medical	Licensed Practical and Licensed Vocational Nurses	42685
	Physicians and Surgeons	191843
	Counselors	50523
	Diagnostic Related Technologists and Technicians	59563
	Social Workers	49607
	Pharmacists	116015
	Dental Hygienists	71356
13	Computer Support Specialists	53141
Software/	Software Developers, Systems Software	104103
Computing	Computer Programmers	80073
	Network and Computer Systems Administrators	76764
	Computer and Information Systems Managers	130036
	Clinical Laboratory Technologists and Technicians	50111
	Drafters	53670
	Database Administrators	79358
14	Industrial Engineers, Including Health and Safety	83202
Engineers	Electrical and Electronics Engineers	95607
	Mechanical Engineers	86182
	Architectural and Engineering Managers	134778
	Civil Engineers	84849
	Compliance Officers	65586
	Architects, Except Naval	78241
	Chemists and Materials Scientists	78884
15	General and Operations Managers	115124
Managers/	Accountants and Auditors	71718
Skilled	Sales Representatives, Services, All Other	61414
Business	Financial Managers	124469
Services	Management Analysts	87539
	Securities, Commodities, and Financial Services Sales Agents	102509
	Financial Analysts	90968
	Education Administrators	90877