Skill Heterogeneity and Aggregate Labor Market Dynamics*

John Grigsby
Department of Economics, Princeton University

August 17, 2022

Abstract

This paper studies aggregate labor market dynamics when workers have heterogeneous skills for tasks which are subject to non-uniform labor demand shocks. Movements in aggregate wages partly reflect a reallocation of different workers across tasks and into employment. This ensures that there nearly always exists some combination of task-specific demand shocks that induce aggregate employment and wages to negatively comove even in a frictionless economy. Furthermore, such reallocation would be interpreted either as a labor wedge or as a shift in an aggregate labor supply curve in representative agent economies. Developing a method to estimate the multidimensional skill distribution, I show that a frictionless model with realistic heterogeneity can replicate the mean wage increase and employment collapse of the Great Recession. Reduced-form composition-adjustment methods recover positive comovements between employment and wages in recent periods suggesting an increasing role for composition effects through time, which the model rationalizes through changes in the skill distribution and composition of sectoral shocks.

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, Adrien Auclert, Mark Bils, Stéphane Bonhomme, Ariel Burstein, Ricardo Caballero, Jonathon Hazell, Ali Hortaçsu, Thibaut Lamadon, Simon Mongey, Kevin Murphy, Jeremy Pearce, Richard Rogerson, Gustavo Souza, Tom Winberry, Liangjie Wu, and Yulia Zhestkova, as well as numerous seminar participants for their many helpful comments.

Grigsby: Julis Romo Rabinowitz Building, Washington Rd, Princeton, NJ 08544, jg6005@princeton.edu

1 Introduction

Many economic phenomena feature labor demand shocks which do not uniformly affect all jobs. For example, the Great Recession featured a collapse in demand for construction workers, while the Pandemic Recession saw demand for in-person service employees fall. Likewise, workers do not have identical skills for all tasks – the skills required of a roofer are vastly different to those of surgeons. Assessing the impact of diffuse labor demand shocks on aggregate employment and wages is therefore challenging.

This paper shows that diffuse labor demand shocks can lead to all manner of comovements between aggregate employment and wages in a frictionless economy when workers have different skills. This observation helps rationalize a large body of research that has found small or negative covariances between employment and wages in response to various labor demand phenomena such as trade shocks, government contract awards, or business cycles, which have been rationalized either through exogenously specified labor supply shocks or frictions.

I develop a model of labor supply featuring multidimensional skill heterogeneity. 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 according to a Roy model and have access to a home option. Workers are paid in proportion to their skill according to the prevailing price per unit of labor in the market for each occupation. Shocks to labor demand shift these prices, thereby reallocating workers into/out of employment and across occupations according to worker type-by-occupation-specific labor supply curves.

In this model, there will always exist some combination of movements in occupation-specific labor prices that lead aggregate employment and wages to move in opposite directions so long as different workers earn different wages in different jobs. Therefore, the usual intuition that net demand shocks cause prices and quantities to move in the same direction in frictionless economies is broken by aggregation.

Indeed, movements in labor prices induce different relationships between employment and wages in the aggregate depending on the distribution of skills in the economy and the composition of price movements. 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 et al. 2019), or from workers being off their labor supply curve due to some sort of friction or "labor wedge" (Chari et al. 2007; Brinca et al. 2016). Such "aggregate labor supply shocks" may be microfounded through aggregation.

Furthermore, the degree to which skills are specific 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. Skill specificity thus influences the general equilibrium effects of granular labor demand shocks.

I apply the model to study the aggregate labor market dynamics in the Great Recession, when real average wages increased by around 2% despite a crash in employment, which has often been attributed to frictions.[1] I close the model by assuming the existence of a finite set of competitive sectors which produce by employing labor into different occupations with different intensities, and which are subject to idiosyncratic profitability shocks. I then non-parametrically estimate the multidimensional skill distribution by adapting the distributional framework of Bonhomme et al. (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.

Shocking the estimated model with a sequence of industry-specific profitability shocks taken from the data, replicates the observed negative comovement between employment and wages during the Great Recession. In the data, real wages 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 exogenous frictions or labor supply shocks. Both skill specificity and heterogeneity in average worker skills are important for this result.

To provide evidence for the model's mechanism, I perform a reduced form composition adjustment in aggregate wages following Solon et al. (1994). While the cyclicality of aggregate real wages has declined through time – even turning countercyclical since 2000 – selection-corrected wages have had stable procyclicality over the last 60 years. This result suggests that the importance of composition effects has

1. For instance, Janet Yellen in her 2014 Jackson Hole symposium postulated that wages did not fall during the Great Recession because downward nominal rigidities constrained wage cuts.

grown, lending credence to the model's mechanism. Empirically, composition effects occurred both because occupations which paid low wages in 2008 suffered large employment declines in 2009 and because low-wage workers were more likely to separate to non-employment than high-wage workers in the same occupation.

Finally, I use the model to assess why aggregate labor market dynamics were so unusual during the Great Recession. The model suggests that aggregate wages would have fallen by approximately 6% were all sectors subject to the same aggregate profitability 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 addition, I estimate that the distribution of skills has changed. Feeding through the realized shocks from the Great Recession yields aggregate wage declines of around 3% under a skill distribution estimated from 1984–89. This is because the degree of both skill specificity and the variance of workers' average skills has grown over time.

Literature Review – The measured acyclicality of aggregate real wages has received great attention in the literature (see Abraham and Haltiwanger (1995) for a survey). This acyclicality implies that large employment declines in recessions must be rationalized either through a labor supply elasticity that is far larger than that implied by micro estimates, 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.[2]

There are many interpretations of this wedge, including exogenous shifts in labor supply (Hall 1997), 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 et al. 2005; Christiano et al. 2014; Hall 2005), or search frictions.[3] However, both of these approaches have drawbacks. Shimer (2005) points out that reasonable calibrations of canonical search models struggle to match cyclical movements of employment and wages. Meanwhile, Beraja et al. (2019) show that a model which calibrates the degree of wage rigidity using regional wage fluctuations requires large exogenous shocks to aggregate labor supply

2. Bils et al. (2018) argue that the wedge between producers' marginal rate of transformation and wages is roughly as large as the wedge between workers' marginal rate of substitution and wages.

3. See Rogerson et al. (2005) for a survey.

to rationalize the observed fluctuations in aggregate employment and wages.

This paper provides a tractable microfoundation for these aggregate labor supply shocks. In this paper's model, the aggregate employment and wage response to sectoral shocks will differ based on the identities of the shocked sectors. If workers leaving the sector may not easily employ their skills elsewhere, the aggregate response of employment may be large relative to the response of labor prices. In addition, if workers expelled from employment are low-skill, the changing composition of the workforce limits fluctuations in measured mean wages. Many models would attribute such a change in the measured relationship between aggregate employment and wages as an inward shift (or flattening) of an aggregate labor supply curve.

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 et al. 2018). Chang and Kim (2007) show that a model with imperfect capital markets and idiosyncratic labor income risk can generate large cyclical movements in the labor wedge, and a low correlation between aggregate hours and productivity.

The role of composition effects in aggregate wage fluctuations was recognized by Solon et al. (1994). These authors studied the cyclical property of wages in the Panel Survey of Income Dynamics (PSID) and found that wages were far more cyclical in a balanced panel of workers which removes compositional shifts. 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 et al. (2020), and Mueller (2017)).[4] This paper builds on this literature in three ways. First, I offer a model which shows how composition effects arise endogenously as a result of heterogeneous labor demand shocks, how skill heterogeneity affects general equilibrium labor supply spillovers to unshocked sectors, and how composition effects shift the relationship between measured wages and employment. Second, I show how the importance of composition effects has changed through time. Finally, I show how to estimate the distribution of skills from the data, and therefore provide a predictive framework for the effect of particular combinations of sectoral shocks.

Barlevy (2002) and Hagedorn and Manovskii (2013) consider variations of job-

4. 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.

search models incorporating 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 in my paper. My paper contributes to this literature by elucidating the role of ex ante skill heterogeneity and diffuse labor demand shocks for equilibrium aggregate labor market dynamics.

The remainder of the paper is organized as follows. Section 2 manipulates accounting identities to show the general existence of direct, reallocation and composition effects on the aggregate wage. Section 3 develops a tractable model of labor supply featuring multidimensional skill heterogeneity, and explores its implications for aggregate employment-wage comovements. Section 4 applies the model to study the labor market dynamics of the Great Recession. It closes and estimates the model, and assesses its ability to match the dynamics of employment and wages in the data. Finally it offers reduced form evidence that composition effects drive the change in wage cyclicality, and studies why composition effects were so large in 2008-09 in counterfactual exercises. Section 5 concludes.

2 Accounting Identity Manipulation

Suppose there is a unit mass of workers in the economy, which may be partitioned into J types of worker with a mass m_j of type j. Workers may either be non-employed or employed in one of K different types of job, which I conceptualize as occupations. Let E_{jk} be the share of type j workers employed in job k. The employment rate of type j workers, denoted $\tilde{E}_j$, is the sum of employment shares across the K occupations:

$$\tilde{E}_j \equiv \sum_k E_{jk} \tag{1}$$

Aggregate employment $\bar{E}$ is the weighted average of type-specific employment:

$$\bar{E} \equiv \sum_j m_j \tilde{E}_j \tag{2}$$

Let ω_{jk} be the average wages of type j workers in job k. One can again calculate average wages of worker type j, denoted $\tilde{\omega}_j$, as the weighted average of that type's job-specific wages:

$$\tilde{\omega}_j = \sum_k \left(\frac{E_{jk}}{\tilde{E}_j} \right) \omega_{jk} \tag{3}$$

Likewise, aggregate wages $\bar{\omega}$ may be calculated as:

$$\bar{\omega} = \sum_j \left(\frac{m_j \tilde{E}_j}{\bar{E}} \right) \tilde{\omega}_j \tag{4}$$

Suppose that there is some arbitrary vector of shocks to labor demand given by $d \ln \mathbf{z}$. Assuming differentiability of type-occupation-specific wages and employment, the first order response of employment to these shocks is

$$\frac{d \ln \bar{E}}{d \ln \mathbf{z}} = \sum_j \underbrace{\left(\frac{m_j \tilde{E}_j}{\bar{E}} \right)}_{\substack{\text{Employment} \\ \text{Share of } j}} \sum_k \underbrace{\left(\frac{E_{jk}}{\tilde{E}_j} \right)}_{\substack{\text{Employment} \\ \text{Share of } k \\ \text{amongst type } j}} \underbrace{\frac{d \ln E_{jk}}{d \ln \mathbf{z}}}_{\substack{\text{Elasticity of} \\ E_{jk} \text{ to shock}}} . \tag{5}$$

This shows that the first order response of aggregate employment to an arbitrary set of shocks is larger if the shock affects workers and jobs (i) with more elastic employment, and (ii) which constitute a larger share of employment.

Meanwhile, the response of aggregate wages to the shock is given by

$$\frac{d \ln \bar{\omega}}{d \ln \mathbf{z}} = \sum_j \left(\frac{m_j \tilde{E}_j}{\bar{E}} \right) \left(\frac{\tilde{\omega}_j}{\bar{\omega}} \right) \left(\underbrace{\frac{d \ln \tilde{E}_j}{d \ln \mathbf{z}} - \frac{d \ln \bar{E}}{d \ln \mathbf{z}}}_{\text{Composition Effect}} + \sum_k \underbrace{\left(\frac{E_{jk} \omega_{jk}}{\tilde{E}_j \tilde{\omega}_j} \right)}_{\substack{\text{Earnings} \\ \text{Share of } k \\ \text{amongst type } j}} \left(\underbrace{\frac{d \ln E_{jk}}{d \ln \mathbf{z}} - \frac{d \ln \tilde{E}_j}{d \ln \mathbf{z}}}_{\text{Reallocation effect}} + \underbrace{\frac{d \ln \omega_{jk}}{d \ln \mathbf{z}}}_{\text{Direct Effect}} \right) \right) \tag{6}$$

The shock has three first order effects on aggregate wages. First is the direct effect: if the shock increases the earnings of workers within their existing job, 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. This effect is stronger (a) the larger the gap between affected workers' wages and aggregate wages, (b) if the affected workers are a larger share of pre-shock employment, and (c) if the affected workers are more elastic to the shock than the average worker. These latter two effects are only present in economies with heterogeneity in both workers and jobs; only the direct effect operates in a representative agent economy with $K = J = 1$.

Importantly, the reallocation and composition effects can operate in the opposite direction to the direct effect. If these effects are strong enough, aggregate wages could fall even if every worker type realizes wage gains within every job type. Therefore, demand shocks might induce aggregate employment increases even while aggregate wages are stable or falling.

This result is very general – the only assumption necessary for its derivation is differentiability of employment and wages with respect to the shock. While this is useful to illustrate the sources of aggregate employment and wage movements, this accounting identity offers little predictive power or economic content. Next, I write down a model with heterogeneous workers and jobs to assess the conditions under which these forces are large and small.

3 A Labor Supply Model with Skill Heterogeneity

This section studies a labor supply model with skill heterogeneity. This section considers a partial equilibrium model so that labor price movements are reduced form labor demand shocks. I extend the model to general equilibrium in Section 4.

3.1 The Model Economy

The economy is static and, as above, consists of K distinct job types indexed by k and a unit mass of workers who belong to one of J skill types indexed by j. The mass of workers who are type j is given by mass m_j.[5]

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

5. For the rest of the paper, I think of the job types as reflecting occupations, and so use the words "job" and "occupation" interchangeably.

simplicity, let Γ denote the matrix whose (j,k) element is γ_{jk}. A law of one price holds for occupational skill so that a worker of type j earns $\omega_{jk} \equiv \gamma_{jk} w_k$ if they work in occupation k. One may think of these γ_{jk} as a metaphor for the skill level of a type j worker in the various tasks employed by occupation k. Each worker values their income with the utility function $u(\gamma_{jk} w_k)$. One may interpret this as assuming workers are hand-to-mouth or as the indirect utility from earning $\gamma_{jk} w_k$.

Workers' only decision is their occupation choice. Each occupation provides some average non-pecuniary benefits ξ_k to workers.[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 is given by

$$k(i) = \underset{k \in \{0,1,\ldots,K\}}{\operatorname{argmax}} \{u(\gamma_{j(i)k} w_k) + \xi_k + \zeta_{ik}\} \tag{7}$$

where $k = 0$ represents the non-employed state.

As above, let $E_{jk}(\mathbf{w})$ denote the probability that a worker of type j chooses to supply her labor to occupation k given the occupation price vector $\mathbf{w} = \{w_1, \ldots, w_K\}$. These are the primitive labor supply curves in the model. Movements in $\mathbf{w}$ will induce workers of different types to reallocate across occupations and to non-employment. This produces selection in the types of workers employed in each occupation.

The preference shocks ζ_{ik} are assumed to be distributed according to an i.i.d. mean zero type 1 extreme value 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 distributional assumptions on ζ are standard in the discrete choice literature and generate a tractable form for the occupational choice probabilities of workers:

$$E_{jk}(\mathbf{w}) = \frac{\exp\left(\frac{u(\gamma_{jk} w_k) + \xi_k}{\nu_j}\right)}{\sum_{k'=0}^{K} \exp\left(\frac{u(\gamma_{jk'} w_{k'}) + \xi_{k'}}{\nu_j}\right)} \tag{8}$$

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

Aggregating workers' individual decision problems yields occupation-level labor supply curves. The mass of workers employed in each occupation E_k is

$$E_k(\mathbf{w}) = \sum_{j=1}^{J} m_j E_{jk}(\mathbf{w}) \tag{9}$$

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

$$L_k(\mathbf{w}) = \sum_{j=1}^{J} m_j E_{jk}(\mathbf{w})\gamma_{jk} \tag{10}$$

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

$$\bar{E}(\mathbf{w}) = \sum_{k=1}^{K} E_k(\mathbf{w}), \qquad L(\mathbf{w}) = \sum_{k=1}^{K} L_k(\mathbf{w}) \tag{11}$$

When $\mathbf{w}$ moves, it may induce separation between $E_k(\mathbf{w})$ and $L_k(\mathbf{w})$ 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:

$$\bar{\gamma}_k = \frac{L_k}{E_k} \tag{12}$$

Since workers are remunerated according to their human capital levels, movements in $\bar{\gamma}_k$ 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, and is key to the model's ability to generate both positive and negative responses of wages to exogenous labor demand shocks. One may express the measured aggregate wage as the employment-share-weighted average wage of each of the occupations:

$$\bar{\omega} = \sum_{k=1}^{K} w_k \bar{\gamma}_k \left(\frac{E_k}{E}\right) \tag{13}$$

where the symbol ω represents take-home pay. 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

$$\tilde{\omega}_j(\mathbf{w}) = \sum_{k=1}^{K} \frac{E_{jk}(\mathbf{w})}{\tilde{E}_j} \gamma_{jk} w_k \tag{14}$$

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 $E_{jk}(\mathbf{w})$. Workers with higher γ_{jk} *relative* to $\gamma_{jk'}$ are more likely to work in occupation k. Next, I discuss the model's structure and study its implications for aggregate employment-wage comovements and labor supply spillovers.

3.2 Discussion of Assumptions

Before considering the properties of the model, it is worth remarking on its structure. 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 of 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 model of, for example, Abowd et al. (1999). 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. Estimating the Γ matrix, as well as the mass of each type and the parameters determining the non-pecuniary benefits of job choice therefore permits a detailed estimation of flexible labor substitution patterns. I present an approach to estimating Γ in Section 4.2.1. In this way, the model provides a way to estimate job-specific labor supply elasticities.

The Γ matrix is a reduced form for a much larger array of traits that individuals may possess. For instance, construction workers may require high levels of strength and dexterity, while managers require organizational skills. Under the assumption that occupation skill may be linearly decomposed into these traits, Welch (1969) shows that the unidimensional occupation-skills captured by Γ entirely describes the relevant skill distribution of the economy. It is worth noting, however, that this reduced form may not hold if bundles of traits are non-linearly combined in the production of each occupation's tasks (Rosen 1983).[7] While Γ represents a useful reduced form representation of skills that grants great analytical tractability, these caveats confound attempts to decompose Γ into more fundamental components.

3.3 Aggregate Employment and Wage Responses to Shocks

Under this model, one can specialize the response of aggregate employment and wages to a set of shocks $\ln \mathbf{z}$ derived in Section 2 above (see Appendix B). Specifically, the response of aggregate employment to these shocks may be written

$$\frac{d\ln \bar{E}}{d\ln \mathbf{z}} = \sum_j \underbrace{\left(\frac{m_j \tilde{E}_j}{\bar{E}}\right)}_{\substack{\text{Employment}\\ \text{Share of } j}} \underbrace{(1-\tilde{E}_j)}_{\substack{\text{Pool of}\\ \text{Workers who}\\ \text{Respond}}} \underbrace{\tilde{\omega}_j}_{\substack{\text{Income}\\ \text{Effect}}} \sum_k \underbrace{\left(\frac{u'(\omega_{jk})}{\nu_j}\right)}_{\substack{\text{LS}\\ \text{Elasticity}}} \underbrace{\left(\frac{\omega_{jk} E_{jk}}{\tilde{\omega}_j \tilde{E}_j}\right)}_{\substack{\text{Share of } k\\ \text{in } j\text{'s earnings}}} \frac{d\ln w_k}{d\ln \mathbf{z}} \tag{15}$$

In this labor supply model, 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 elastic is labor supply, which is governed by the standard deviation of the idiosyncratic preference shocks ν_j and their marginal valuation of income $u'(\omega_{jk})$. There is an income effect on employment in this model: if workers earn high wages on average, then they are further from the outside option of non-employment and thus more likely to work in response to shocks. This income effect can be killed through appropriate choice of $u'(\omega_{jk})$. Finally, employment responds more to the shock if there are more non-employed workers that can be drawn into employment. If all workers are already employed ($\tilde{E}_j = 1 \forall j$), small changes in

7. Edmond and Mongey (2019) explore this idea in the context of technology adoption and show that the law of one price for skills may fail if workers cannot unbundle their fundamental talents.

the price of labor have no effect on employment.

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

$$\frac{d\ln\bar{\omega}}{d\ln\mathbf{z}} = \sum_j \underbrace{\frac{m_j\tilde{\omega}_j\tilde{E}_j}{\bar{\omega}\bar{E}}}_{\substack{\text{Share of } j\\ \text{in Agg. Earnings}}} \sum_k \underbrace{\left(\frac{\omega_{jk}E_{jk}}{\tilde{\omega}_j\tilde{E}_j}\right)}_{\substack{\text{Share of } k\\ \text{in } j\text{'s earnings}}} \left[\overbrace{1}^{\text{Direct}} + \overbrace{\underbrace{\left(\frac{u'(\omega_{jk})}{\nu_j}\right)}_{\substack{\text{LS}\\ \text{Elasticity}}} \underbrace{(\omega_{jk}-\tilde{\omega}_j)}_{\substack{\text{Cross-Occ.}\\ \text{Wage Change}}}}^{\text{Reallocation}} \right.$$
$$\left. + \overbrace{\underbrace{\left(\frac{u'(\omega_{jk})}{\nu_j}\right)}_{\substack{\text{LS}\\ \text{Elasticity}}} \underbrace{(1-\tilde{E}_j)}_{\substack{\text{Pool of}\\ \text{Workers who}\\ \text{Respond}}} \underbrace{(\tilde{\omega}_j-\bar{\omega})}_{\substack{\text{Relative}\\ \text{Wages}}}}^{\text{Composition}} \right] \frac{d\ln w_k}{d\ln\mathbf{z}} \tag{16}$$

As before, this clarifies that changes in the price of occupation-specific labor prices have direct, reallocation and composition effects on aggregate wages. 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 important if the occupations being shocked tend to constitute a large share of income for low-wage workers. As a real world example, suppose those with a high school degree earn less than college educated workers and construction constitutes a larger share of income for high school workers than college educated workers; then negative shocks to construction induce a positive composition effect on the aggregate wage. This is why the interaction of skill (and thus earnings) heterogeneity and the composition of shocks is critical for understanding aggregate wage movements.

To illustrate this interaction, consider the following simple example. Suppose there are only two types of workers and two types of job: $K = J = 2$. For simplicity, suppose further that each worker type constitutes half of the population ($m_j = 0.5$ for $j = 1, 2$), workers share a common fundamental labor supply elasticity $\nu_j^{-1} = 2$ and have log utility. As equations (15) and (16) show, the first order aggregate responses of employment and wages depend on the labor price movements in occupation 1 and 2 induced by the shocks. It is instructive, therefore, to consider how employment and wages respond for different movements in occupation labor prices dw_1 and dw_2 under different assumed skilled distributions Γ.

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.

Figure 1 partitions the shock space (dw_1, dw_2) for a variety of skill matrices. In each panel, the dashed black line plots the combination of labor price movements that induce no employment change, while the solid black line plots the combination of labor price movements that induce no wage change. Panel A assumes a representative agent skill matrix and that the pre-shock labor prices are 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 (15) and (16) 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

Figure 1: Splitting Regions of Shock Space by Direction of Aggregate Wage and Employment Movements in Economy with Two Job Types

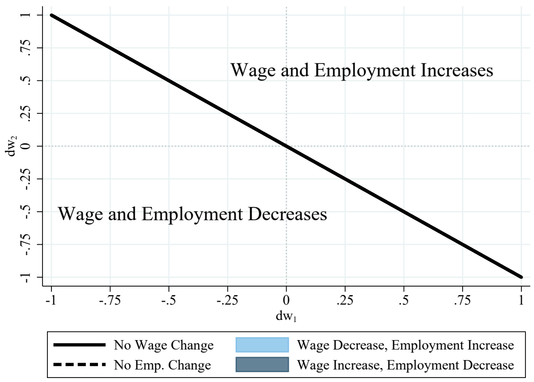

Panel A: $\Gamma^{(RA)}$: Representative Agent, $w_1 = w_2 = 1$

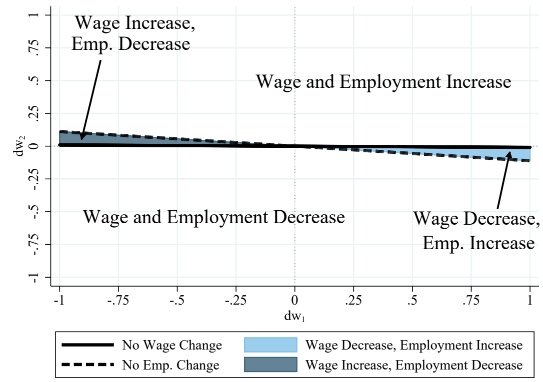

Panel B: $\Gamma^{(RA)}$: Representative Agent, $w_1 = 0.75,\ w_2 = 1.5$

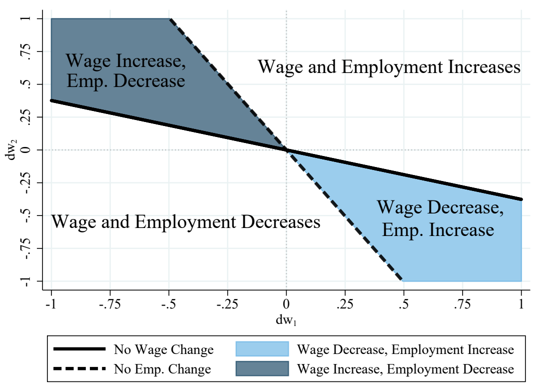

Panel C: $\Gamma^{(CA)}$: Comparative Advantage, $w_1 = 0.75,\ w_2 = 1.5$

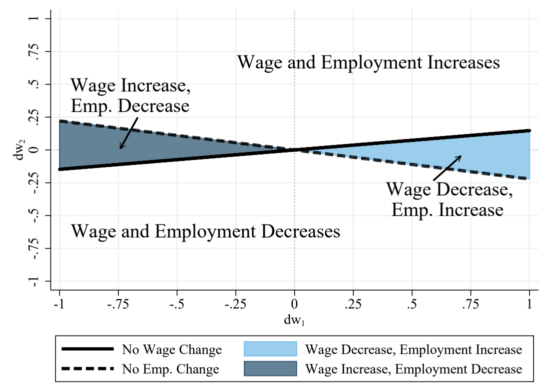

Panel D: $\Gamma^{(AA)}$: Absolute Advantage, $w_1 = 0.75,\ w_2 = 1.5$

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.

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 swage. 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.[8]

This section shows that the interaction of skill heterogeneity and the composition of labor price shocks has a first order impact on aggregate employment and wages. In this simple frictionless model of labor supply, skill and shock heterogeneity can induce aggregate employment and wages to move in opposite directions, breaking the standard intuition that demand shocks induce positive comovements between measured prices and quantities. The next section explores how this result influences the interpretation of aggregate employment and wage fluctuations in the data.

3.4 Aggregating Labor Supply Curves

The canonical aggregate labor supply curve traces out the measure of workers willing to be employed as a function of the prevailing wage. In this paper's model, the slope and location of this curve will depend on the set of occupational prices used to construct it. This results from differences in the characteristics of workers employed in the occupations subject to labor price movements.

To build intuition for the relationship between aggregate employment and wages, I return to the simple two-type, two-occupation calibration of the previous subsection. Using the model, one can trace out an aggregate employment-wage schedule as the price of occupational labor services changes. Specifically, one can vary the vector of labor prices $\mathbf{w}$ to generate occupation choices according to equation (7). One can

8. A.1 considers additional specifications of the skill matrix. The employment-wage divergence region grows large when the skill matrix combines absolute and comparative advantage.

then plot the relationship between $\bar{E}(\mathbf{w})$ and $\bar{\omega}(\mathbf{w})$ as implied by equations (11) and (13), respectively. I do this for the same three specifications of the Γ matrix as above.[9]

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 for the three Γ matrices are presented in Panels A and B of 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.

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

The red line shows the response under the comparative advantage skill matrix. In this case, 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 resembles an upward shift of the representative agent schedule as workers sort into their occupation of skill, thereby realizing higher wages for any given employment rate. The strength of this force is mediated by the ratio of the marginal value of income $u'(\omega_{jk})$ to the variance of non-pecuniary benefits ν_j.

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

9. 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.

Figure 2: Aggregate Employment-Wage Schedule As Vary Γ and Occupation Prices

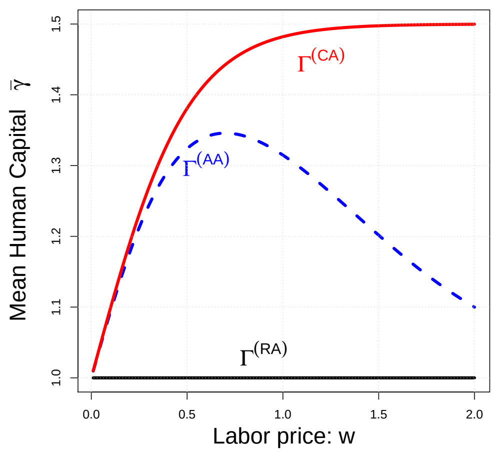

Panel A: Mean Human Capital
$w_1 = w_2 = w$

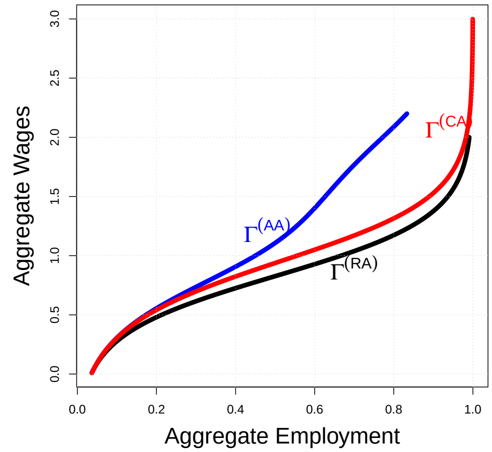

Panel B: Agg. Employment-Wage
Schedule; $w_1 = w_2 = w$

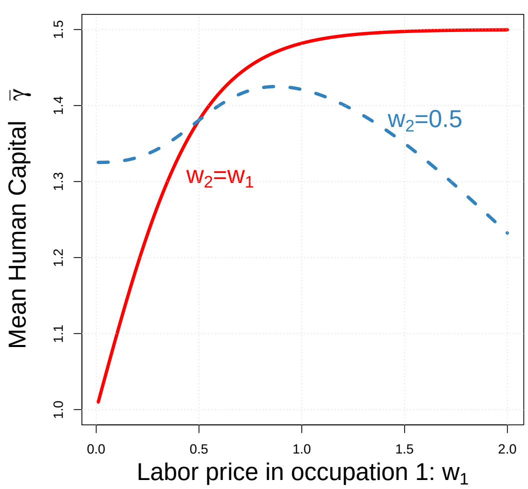

Panel C: Mean Human Capital
w_2 fixed at 0.5, w_1 varies

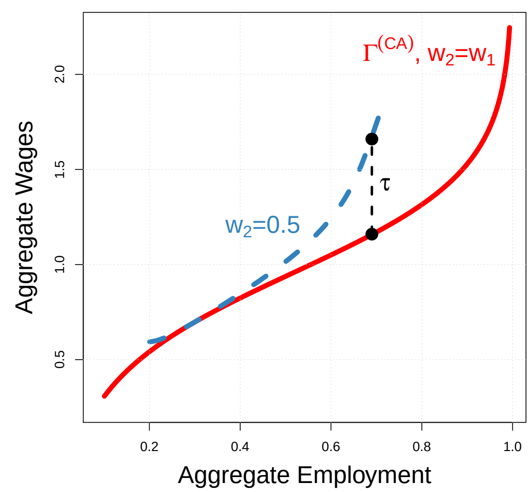

Panel D: Agg. Employment Wage
Schedule; w_2 fixed at 0.5, w_1 varies

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. Panels A and B plots the implied movements when the price of labor in occupation 1 is constrained to equal the price in occupation 2, while Panels C and D plot the implied curves when occupation 2's labor price is fixed at 0.5 and occupation 1's price is allowed to vary between 0 and 2. Panels A and C plot the mean human capital of employed workers $\bar{\gamma}$ against the prevailing price of labor, while Panels B and D plot 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 $w_1 = w_2$, and Γ exhibits comparative advantage. The blue dashed line in Panels C and D is the curve when w_2 is fixed to 0.5, and Γ exhibits comparative advantage. Γ matrices defined as in equation (17).

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. Indeed, if the composition effect is strong enough (e.g. if the variance of ζ_{ikt} were very small), the model with absolute advantage could generate a non-monotone aggregate relationship between employment and wages if increases in the price of labor induced large enough inflows of low-type workers.

Idiosyncratic Shock – Now consider the opposite extreme in which w_2 were fixed at 0.5, while only w_1 varies. This case is depicted in Panels C and D of Figure 2. I restrict attention to the case with comparative advantage which most naturally permits price dispersion across k. The red lines recreate the red curves from panels A and B, while the blue dashed line shows the curves after fixing w_2 at 0.5. 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.

This has important implications for macroeconomic accounting frameworks. 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 D. 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 (Chari et al. 2007), or as an exogenous shock to labor supply (Beraja et al. 2019). Therefore, compositional shifts can microfound the labor wedge or labor supply shocks which have been shown to be important to account for recent business cycle fluctuations.

This exercise highlights two primary reasons why wage-employment comovements may differ over time. First, the distribution of skills may change. Second, labor price movements may change if different labor demand shocks hit the economy at different times, inducing variable aggregate relationships between employment and wages.

3.5 Cross-Occupation Labor Supply Spillovers

When the price of labor falls in one occupation, workers begin to 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.

Consider the simple two-occupation, two-type labor supply model. Note that the labor supply curve in each occupation $L_k(\mathbf{w})$ as defined in equation (10) depends on the entire vector of labor prices $\mathbf{w} = (w_1, w_2)$. Nevertheless, one can plot the labor supply curve in occupation 1 given a fixed price of labor in occupation 2.

Panel A of Figure 3 plots this curve for the three skill matrix specifications 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 strongly resembles that of the absolute advantage skill distribution. In the absolute advantage case, there are two offsetting effects. Low-skill workers do not respond to increases in the price of labor much, which makes the labor supply curve more inelastic *ceteris paribus*. However, for a given increase in the price of labor, it is principally high skill workers who enter occupation 1. Since these workers carry more human capital with them, this makes labor supply curve more elastic in labor-units space. The lack of response of low skill workers is almost exactly offset by the fact that the responding high skill workers are more productive.

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. Thus, the distribution of skills affects the behavior of skill prices in response to labor demand shocks, which will be important in the general equilibrium application of the model in Section 4.

Panel A is plotted assuming that w_2 were equal to 1. Suppose now that the price of labor in occupation 2 exogenously moved to $w_2 = 0.5$. This simulates a large

Figure 3: General Equilibrium Labor Supply Spillovers

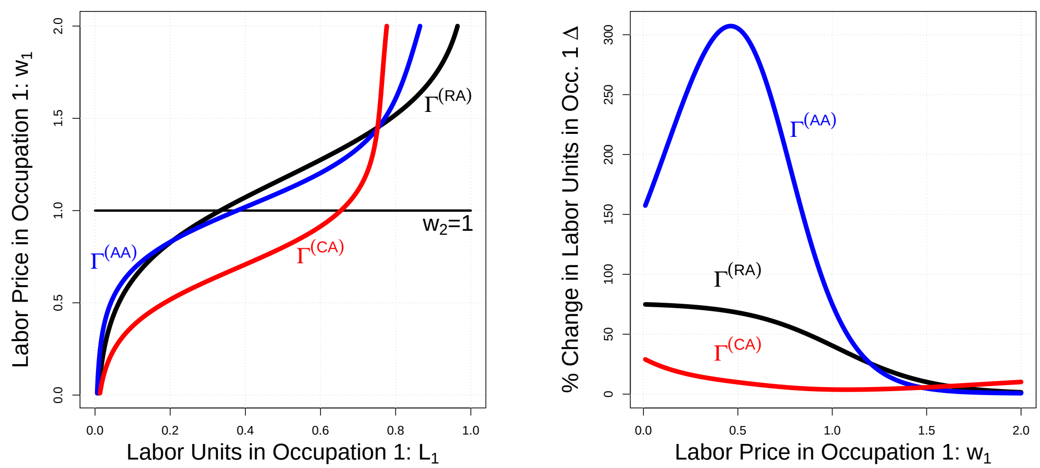

Panel A: Labor Supply Curves in Occ. 1

Panel B: Outward Shift in Occ. 1 Labor Supply Curve from Decline in w_2 (Δ)

Notes: Figure shows the behavior of 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).

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 horizontal movement in the labor supply curve for every given level of w_1. One can calculate the percentage change in labor units supplied to occupation 1 induced by the change of price in occupation 2 for every given price of labor 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 3 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. 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. Limiting the strength of labor supply spillovers across jobs will turn out to be crucial in the quantitative application of the model in Section 4, suggesting that models must incorporate both absolute and comparative advantage of workers to generate realistic employment and wage fluctuations.

To summarize, skill heterogeneity exerts three novel effects on the relationship between aggregate employment and wages. First, shifts in the distribution of skills change the sorting patterns in response to a shock to the price of labor, thereby affecting the relationship between aggregate employment and measured wages. Second, changes in the relative price of labor between two occupations similarly affect sorting patterns, which implies that idiosyncratic shocks to occupations influence the relationship between aggregate employment and wages. There always exists some combination of labor price shocks that induce negative comovements between aggregate employment and wages so long as different workers earn different amounts in different jobs. Movements in this relationship manifest as labor wedges or labor supply shocks in representative agent economies. Finally, the multidimensional nature of workers' skills implies that shocks to the price of labor in one occupation induce shifts in the labor supply curves of other occupations, thereby exerting pressure on the price of labor elsewhere in the economy. This force is especially strong if skills are easily transferred between tasks.

4 Application: the Great Recession Labor Market

I apply the model to study employment and wage dynamics during the Great Recession by embedding the labor supply model into a general equilibrium model of the labor market, which can then be calibrated to industry-level data. Next I describe how to estimate the multidimensional skill matrix using two-period panel data on earnings and occupation choices. While the model is stylized in nature, this section shows that realistic calibrations of the skill distribution and diffuse labor demand shocks are sufficient to generate negative comovements between employment and wages. Finally, I provide reduced form evidence in support of the model's mechanism by demonstrating that composition-adjusted wages have had stable cyclicality through time.

4.1 Closing the Model

Households – There is a large representative household containing a measure 1 of infinitely-lived workers. These workers supply their labor according to the model of Section 3. The household and its workers are risk-neutral and hand-to-mouth consume a numeraire consumption good C. The household takes income from labor I and from firms' profits Π as given, which it uses to finance consumption. The household additionally gains non-pecuniary benefits Ξ from the workers' activities. The household consumes its total income each period: $C = I + \Pi$.

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

$$Y = \left(\sum_{s=1}^{S} \tau_s \hat{y}_s^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \tag{19}$$

where $\hat{y}_s$ the demand for sector s's output from the final goods producer and τ_s is a shifter of demand for sector s. The demand curve for sector s's output is

$$p_s = \tau_s \left(\frac{Y}{\hat{y}_s} \right)^{\frac{1}{\eta}} \tag{20}$$

Intermediate Goods Firms – Each sector s is populated by a representative competitive firm. The firm hires labor in the K occupations to produce output $y_s = z_s F^{(s)}(l_{s1}, l_{s2}, \ldots, l_{sK})$, where z_s denotes the productivity (TFP) of sector s, l_{sk} is the quantity of occupation k services hired by sector s, and $F^{(s)}(\cdot)$ is a sector-specific production function which is increasing and concave in each of its arguments.

The price of sector s's output is given by p_s, which firms take as given. Each occupation k has one price w_k. Therefore, intermediate firms solve

$$\pi_s = \max_{\{l_{s1}, l_{s2}, \ldots, l_{sK}\}} \quad p_s z_s F^{(s)}(l_{s1}, l_{s2}, \ldots, l_{sK}) - \sum_{k=1}^{K} w_k l_{sk} \tag{21}$$

Total profits in the economy is the sum of all sectors' profits: $\Pi := \sum_{s=1}^{S} \pi_s$. Conditional on the choice of occupation, workers are indifferent between sectors.

In the quantitative exercise below, I explore the economy's response to changes in the distribution of sector TFP z_s. Note that τ_s and z_s both affect the marginal revenue product of labor in sector s isomorphically: τ_s affects it through the price of output from sector s while z_s affects the marginal quantity product of labor. Therefore, I normalize $\tau_s = 1$ for all sectors and interpret industry shocks as a combination of productivity and demand shocks.

Equilibrium Definition − A competitive equilibrium is a set of output prices $\mathbf{p} = \{p_s\}_{s=1}^{S}$, occupation prices $\mathbf{w} = \{w_k\}_{k=1}^{K}$, and decision rules $\{E_{jk}(\mathbf{w})\}_{j,k}, \mathbf{l} = \{l_{sk}(p_s, \mathbf{w}|z_s)\}_{s,k}, \hat{\mathbf{y}} = \{\hat{y}_s(\mathbf{p})\}_s$ such that, given sectoral productivities $\mathbf{z} = \{z_1, \ldots, z_S\}$,

1. The occupation demand functions $\{l_{sk}(p_s, \mathbf{w}|z_s)\}_{s,k}$ solve the intermediate sectors' firm's problem (21),
2. The workers' occupation choice decisions solve (7),
3. The demand for each sector's output from the final goods producer $\hat{y}_s(\mathbf{p})$ is equal to the supply of that sector's output $z_s F^{(s)}(\mathbf{l}_s(p_s, \mathbf{w}|z_s))$,
4. The final goods market clears; that is, aggregate output equals total income: $Y = C = I + \Pi$
5. Occupation-specific labor markets clear

$$L_k(\mathbf{w}) = \sum_{s=1}^{S} l_{sk}(\mathbf{p}, \mathbf{w}|z_s) \qquad \text{for all } k$$

The approach to characterizing equilibrium is detailed in Appendix B.

4.2 Model Calibration

4.2.1 Estimation of Labor Supply Parameters

The identification and estimation of the skill distribution follows closely the distributional framework developed by Bonhomme et al. (2019). Estimation follows 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.[10] Log observed wages $\ln \omega_{it}$ are then

$$\ln \omega_{it} = \ln \gamma_{j(i)k_t(i)} + \ln w_{k_t(i)t} + \ln \epsilon_{it}. \tag{22}$$

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, occupation 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.

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 $Z^t = \{Z_{i1}, \dots, Z_{it}\}$. 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. *(Connecting Cycles) - For any two occupations k and $k' \in \{0, \dots, K\}$, there exists a connecting cycle $(k_1, \dots, k_R), (\tilde{k}_1, \dots, \tilde{k}_R)$ such that $k_1 = k$ and $k_r = k'$ for some r, and such that the scalars $a(1), \dots, a(J)$ are all distinct where*

$$a(j) = \frac{\mathbb{P}_{k_1\tilde{k}_1}(j)\mathbb{P}_{k_2\tilde{k}_2}(j)\dots\mathbb{P}_{k_R\tilde{k}_R}(j)}{\mathbb{P}_{k_2\tilde{k}_1}(j)\mathbb{P}_{k_3\tilde{k}_2}(j)\dots\mathbb{P}_{k_1\tilde{k}_R}(j)}.$$

In addition, for all k, k' possibly equal, there exists a connecting cycle $(k'_1, \dots, k'_R), (\tilde{k}'_1, \dots, \tilde{k}'_R)$ such that $k'_1 = k$ and $\tilde{k}'_r = k'$ for some r

10. 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.

4. *(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, which is derived in Appendix C. The formal argument closely follows that of Bonhomme et al. (2019).[11]

Assumption 1 has four pieces. The first is that workers' idiosyncratic wage draws are uncorrelated with their occupation choice, conditional on their type and choice of occupation. 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 wage draws are serially independent, conditional on a worker's type and occupation choice. In some settings, this is a reasonable assumption: for instance, tip workers or those in the gig economy have nearly i.i.d. fluctuations around a mean wage. Similarly, upper executives have roughly i.i.d. fluctuations in their earnings as a result of random stock performance.[12]

The third item of Assumption 1 requires that any two occupations belong to a connecting cycle for every type of worker. This does *not* require that every worker type must flow between every pair of occupations (k, k') bilaterally. Rather, it imposes graph connectedness in the sense of Abowd et al. (1999). This will hold under the model given the distributional assumptions on the idiosyncratic preference shocks. In addition, workers of different types must flow in different ways, which occurs so long as the γ_{jk} differ by worker type.

11. Assumption 1 may be relaxed at the expense of greater data requirements. Unfortunately, the set of large, representative, long-run panel datasets containing information on occupation and wages is small, requiring the use of panel data with just two periods.

12. 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 higher or low wage workers switch jobs.

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

Identification is achieved through occupation switchers. Consider the wage change of an individual of type j who moves from occupation k to k' between $t-1$ and t

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

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, the model assumes that there is one price in each occupation that applies to all workers. Therefore, the endogenous price of labor in occupation k in period t 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.[13]

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 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. Engineering, for instance, may have a low ξ_k not because engineering is an unpleasant occupation, but rather because

13. 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.

the annualized cost of acquiring and maintaining engineering knowledge is high.

To maximize the likelihood function, I assume that measurement error in wages follows a log-normal distribution which is type-occupation specific, following Bonhomme et al. (2019). Second, that taste shocks are assumed 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 closely aligns with those found in the literature estimating micro labor supply elasticities.

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 employ two data sources. First, I rank SOC occupations according to the share of workers with at least some college education using data from the Current Population Survey (CPS).[14] 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. 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.[15]

A full picture of the $K = 15$ clustered occupations is provided in Appendix D. Clusters are ordered according to their mean annual income in the period 2002-2006,

14. 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 D.

15. Gathmann and Schönberg (2010) build indices of skill relatedness using these surveys.

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, cluster 12 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.[16] 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 D. Worker earnings ω_{it} are measured by the total labor income of workers over the prior year, deflated by the CPI-U.[17] 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) in order to minimize sampling noise.

Estimated Skill Distributions – Appendix Table A1 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 two 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

16. 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. In contrast, Bonhomme et al. (2019), on which the estimation is based, allowed for the equivalent of $K = 10$ and $J = 6$.

17. 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.

management consultants). In contrast, type 6 workers supply 2.57 units of human capital to business services occupations, but only 0.38 units of human capital to routine occupations. Recall that the γ_{jk} are normalized to have unit mean (weighted by worker type shares) within each occupation. As a result, these γ_{jk} may be interpreted as the 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. For brevity, 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. However, the model overpredicts the share of people who switch occupation. This is due to the i.i.d. assumption on the idiosyncratic preference shocks ζ_{ikt}. As explored in Appendix A.4, the estimation also produces intuitive patterns of skill relatedness. For instance, Figure A7 shows that workers who are skilled in routine occupations, such as stock clerks and cashiers, are also skilled as manual laborers.

4.2.2 Calibrating Labor Demand

Table 1 summarizes the model's calibration. The parameters governing labor supply are estimated using the maximum likelihood approach outlined above. 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).[18] 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

18. Broda and Weinstein (2006) estimate the mean elasticity of substitution across 3-digit SITC products, rather than sectors. The true elasticity of substitution across 3-digit sectors may therefore be somewhat lower than 4. Reducing 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.

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

s's wage bill that is accounted for by occupation k.[19] 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 be fixed to the average share in each sector over the period 2002-2006.

Estimating Sector-Level TFP Series – I estimate TFP using a Solow residual approach. A challenge arises when there is selection on unobservable quality in labor inputs. Through the lens of my model, the problem arises because $\bar{\gamma}_{kt}$ fluctuates over the cycle. 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

19. 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.

implies that the TFP of sector s in period t may be estimated using the equation

$$\ln z_{st} = \ln \text{Value Added}_{st} - x_s \sum_{k=1}^{K} \alpha_{sk} \ln(\bar{\gamma}_{kt} n_{kt}^s) - (1 - x_s) \ln(\text{Non-Labor Input}). \quad (23)$$

The employment in each sector in each occupation, n_{kt}^s, 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. Equation (7) yields an estimate of the mean human capital of workers employed in every occupation in every year. Recall that, under the model, the labor supply estimation remains valid even with labor demand shocks because the occupation-by-year fixed effect absorbs all variation in labor prices.

I employ sectoral value added and non-labor input data from the BLS' KLEMS Multifactor Productivity Series.[20] These data go back to 1987. 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. Finally, I use equation (23) to estimate sectoral TFP series.

The adjustment for $\bar{\gamma}$ is meaningful. Appendix Table A2 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%,

20. These value added data include sectoral prices. Thus the z_{st} arising from this calculation may be interpreted as reflecting shocks both to productivity and to τ_s – the demand for sector s – and may be broadly interpreted as profitability shocks.

Figure 4: Aggregate Wage and Employment Responses, 2008-2012

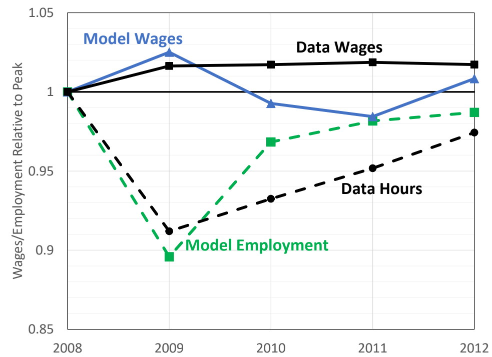

Notes: Figure plots the model-implied aggregate behavior of wages and employment in response to the calibrated sectoral TFP series around the Great Recession. Wages and employment normalized to be 1 in 2008. Black lines plot aggregate hours and average weekly earnings of all private employees, deflated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

compared to a 4.2% decline in the unadjusted BLS series.

4.3 Great Recession Aggregate Labor Market Dynamics

Figure 4 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 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. Reassuringly, the model successfully predicts occupational wage dynamics and the strength of the relationship between employment changes and wage levels: the model's mechanism is observable in the data.

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.[21] 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 remarkable: 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 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 et al. (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

21. The data numbers consider year over year changes from March 2008 to March 2009.

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%
Model: Aggregate TFP Shock	-6.4%	-16%

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. The final row considers an estimated model in which all sectors experienced the same 5.9% decline in TFP.

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 A1 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, 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. Thus type 1 workers have 0.55 units of human capital while type 8 workers have 7.54 units of human capital, and they may supply those units equally well across all occupations. In this model, employment falls by 7.9% while real wages fall by 1.3%. Here, there remains a strong 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 have labor supply which is relevant to all possible pursuits. Therefore, as foreshadowed in Section 3.5, 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. The downward price pressure more than overcomes the selection force generated in the pure absolute advantage model, thereby preserving a positive covariance 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%.

The final row of Table 2 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. Under this counterfactual set of labor demand shocks, the model implies that real wages would have declined by approximately 6%. The negative demand shocks during the Great Recession were unusually concentrated amongst sectors which primarily employ manual laborers (see Appendix Figure A3). This ensured that workers who have skills in manual occupations received a large negative demand shock for their skills and furthermore 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 paid (i.e. have low γ_{jk}s), this generated a large composition effect with limited spillovers to the rest of the economy.

This section shows that worker and demand shock heterogeneity is 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 employment and wages over this period. Both skill specificity and absolute advantage are necessary to this result surviving in equilibrium.

The role of the skill distribution – The model suggests that changes in the skill distribution change the aggregate dynamics of employment and wages. This section estimates these changes in the years leading up to the Great Recession.

I re-estimate the labor supply side of the model using CPS data from 1984-89, immediately prior to the 1990-91 recession. Feeding through the realized sequence of Great Recession TFP shocks generates real wage declines of 3% and employment declines of just 2% (see Appendix Figure A2), standing in stark contrast to the wage increases the baseline model predicts. This suggests meaningful changes in the skill distribution have occurred over the past 30 years.

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 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 movements along different aggregate wage-employment schedules. 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.[22]

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.

22. 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

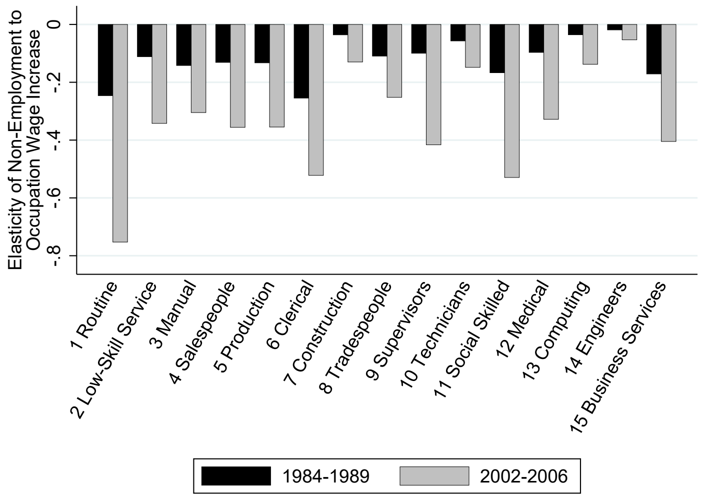

Notes: Figure reports the estimated model-implied elasticity of non-employment to a change in the price of each occupation's price of labor w_k. The estimation procedure is outlined in Section 4.2.1, and carried out separately in the CPS March Supplement for the periods 1984-1989 (black bars) and 2002-2006 (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 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.4 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.

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. If this variance is high, 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. Finally, the transferability of skills has declined amongst high-skill occupations, occupations employing manual labor, and occupations employing social skills. There may be many reasons for these changes, such as changes in education policy or a 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. These changes have conspired to increase the primitive occupation-specific labor supply elasticities 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 in the future.

4.4 Reduced Form Evidence for Mechanism

A strong composition effect is central to the model's ability to match the labor market dynamics of the Great Recession. If the mechanism is correct, it must be that composition effects have grown in importance over time, because wages were generally not countercyclical during the latter half of the 20th century. This section provides reduced form evidence that shifting composition effects do indeed account for the change in the cyclicality of real wages, while Appendix A.5 examines the role of composition effects specifically during the Great Recession.

I first estimate the relationship 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

$$\Delta \ln W_t = \sum_p \left[\alpha_p + \beta_p \Delta \ln GDP_t\right] \mathbf{1}\{t \in p\} + \epsilon_t \tag{24}$$

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.

The estimated cyclicalities of aggregate wages β_p are presented in Panel A of Figure 6, 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 an elasticity of wage growth to per capita GDP growth of -0.23 (SE: 0.09). This trend looks set to continue past 2020, given the large increase in aggregate wages observed during the COVID-19 pandemic (Cajner et al. 2020).[23]

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

$$\ln \omega_{it} = \alpha_i + \tilde{\beta}_p \ln GDP_t + \eta_0 \cdot t + \eta_1 \cdot t^2 + \eta_2 X_{it} + \epsilon_{it} \tag{25}$$

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

23. Variables downloaded from FRED, maintained by the St. Louis Fed. Average hourly earnings of production and non-supervisory employees come from the Current Employment Statistics. GDP per capita is from the BEA and deflated by the GDP price deflator. Wages deflated by core CPI.

Figure 6: The Cyclicality of Real Wages in Aggregate and Controlling for Composition

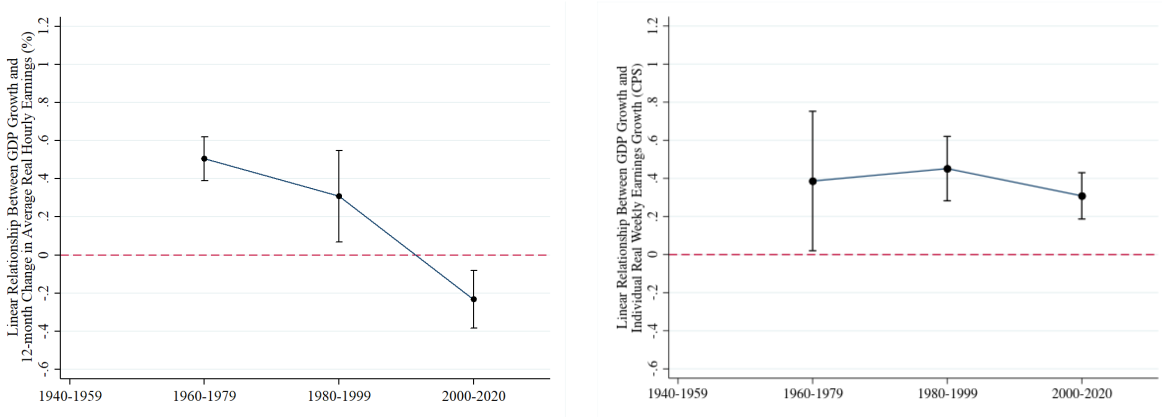

Panel A: Real Average Hourly Earnings

Panel B: Composition-Adjusted Average Wages (following SBP)

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 (24) 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 (26) 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. Data sources: BLS, BEA, CPS, and FRED.

in aggregate data. By estimating equation (25) in first differences, one controls for characteristics of a worker which are fixed over time, such as the worker's permanent ability. Therefore estimating

$$\Delta \ln \omega_{it} = \tilde{\eta}_0 + \tilde{\beta}_p \Delta \ln GDP_t + \tilde{\eta}_1 \cdot t + \tilde{\eta}_2 X_{it} + \Delta \epsilon_{it}. \tag{26}$$

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 micro-data 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.

Panel B of Figure 6 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.[24] 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, lending credence to the core mechanism of the structural model in this paper.[25]

5 Conclusion

This paper argues that skill and shock heterogeneity have a first order effect on the behavior of aggregate employment and wages. When workers differ in their skills for a variety of tasks, diffuse shocks to labor demand 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. This is accentuated by skill specificity, which constrains the extent to which localized demand shocks exert downward pressure on labor prices elsewhere in the economy.

I make this case in a frictionless Roy model featuring workers with heterogeneous skills for a variety of tasks. Whenever workers earn different wages in different tasks, there will always exist some combination of labor price shocks – a short-hand for labor demand shocks in a partial equilibrium labor supply model – that lead employment and wages to move in opposite directions, even in this frictionless economy.

I assess whether this mechanism is quantitatively important in equilibrium by studying the aggregate labor market dynamics during the Great Recession. 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. 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

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

25. The BLS' Employment Cost Index (ECI), which builds a wage index holding fixed the share of industry-occupation cells in employment, also shows mild declines in real wages during the Great Recession, further suggesting the importance of composition effects during this period.

directions. Skill specificity, which constrains cross-occupation labor supply spillovers, and differences in average skill across workers, which generate composition effects, are both necessary for this result. Reduced-form selection-correction methods show that composition forces have indeed been behind the stark decline in aggregate wage cyclicality over the last 60 years. The model suggests that these strong composition forces 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. Understanding the source of these changes is fertile ground for future research.

The core insight of this paper is that aggregating employment and wage movements is non-trivial whenever labor demand shocks are diffuse. The complexity of this aggregation offers a frictionless microfoundation for a number of phenomena, such as the strong procyclicality of the labor wedge, exogenous representative agent labor supply shocks, the large gap between estimated micro and macro labor supply elasticities, or the instability of parameters governing frictions estimated from aggregate time series. Furthermore, the possibility of reallocation and composition effects necessitates careful thought when interpreting estimates from research designs which exploit differential exposure to diffuse labor demand shocks. Embedding the tractable and estimable framework of this paper into more complex models may provide a relatively straightforward method to account for cross-sector labor supply spillovers and to control for equilibrium compositional effects in future research.

References

Abowd, John M, David N Margolis, and Francis Kramarz. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2): 251–333.

Abraham, Katharine G., and John C. Haltiwanger. 1995. "Real Wages and the Business Cycle." *Journal of Economic Literature* 33 (3): 1215–1258.

Barlevy, Gadi. 2002. "The Sullying Effect of Recessions." *Review of Economic Studies* 69 (1): 65–96.

Benhabib, Jess, and Richard Rogerson. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." 99 (6): 1166–1187.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2019. "The Aggregate Implications of Regional Business Cycles." *Econometrica* 87 (6).

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2018. "Resurrecting the Role of the Product Market Wedge in Recessions." *American Economic Review* 108 (4-5).

Bils, Mark J. 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93 (41): 666–689.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A Distributional Framework for Matched Employer Employee Data." *Econometrica* 87.

Brinca, Pedro, V V Chari, Patrick J Kehoe, and Ellen R McGrattan. 2016. "Accounting for Business Cycles." *NBER Working Paper Series* 22663.

Broda, Christian, and David E. Weinstein. 2006. "Globalization and the Gains from Variety." *The Quarterly Journal of Economics* 121 (2): 541–585.

Cajner, Tomaz, Leland D. Crane, Ryan A. Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. 2020. "The U.S. Labor Market during the Beginning of the Pandemic Recession." *Brookings Papers on Economic Activity* Spring.

Chang, Yongsung, and Sun-Bin Kim. 2007. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations." *American Economic Review* 97.

Chang, Yongsung, Sun-Bin Kim, Kyooho Kwon, and Richard Rogerson. 2018. "Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins."

Chari, V V, Patrick J Kehoe, and Ellen R McGrattan. 2007. "Business Cycle Accounting." *Econometrica* 75 (3): 781–836.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea 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 (2): 471–75.

Christiano, Lawrence J, and Martin Eichenbaum. 1992. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations." *American Economic Review* 82 (3): 430–450.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review* 104 (1): 27–65.

Edmond, Chris, and Simon Mongey. 2019. "Unbundling Labor." *Working Paper.*

Gathmann, Christina, and Uta Schönberg. 2010. "How General Is Human Capital? A Task-Based Approach." *Journal of Labor Economics* 28 (1): 1–49.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari. 2020. "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires." *The Review of Economic Studies* 87 (4): 1876–1914.

Gertler, Mark, and Antonella Trigari. 2009. "Unemployment Fluctuations with Staggered Nash Wage Bargaining." *Journal of Political Economy* 117 (1): 38–86.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data." *American Economic Review* 111 (2): 428–71.

Hagedorn, Marcus, and Iourii Manovskii. 2013. "Job Selection and Wages over the Business Cycle." *American Economic Review* 103 (2): 771–803.

Hall, Robert E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *The American Economic Review* 95 (1): 50–65.

———. 1997. "Macroeconomic Fluctuations and the Allocation of Time." *Journal of Labor Economics* 15 (1).

Mueller, Andreas I. 2017. "Separations, sorting, and cyclical unemployment." *American Economic Review* 107 (7): 2081–2107.

Rogerson, Richard, Robert Shimer, and Randall Wright. 2005. "Search-Theoretic Models of the Labor Market: A Survey." *Journal of Economic Literature* 43.

Rogerson, Richard, and Johanna Wallenius. 2009. "Micro and Macro Elasticities in a Life Cycle Model with Taxes." *Journal of Economic Theory* 144.

Rosen, Sherwin. 1983. "A Note on Aggregation of Skills and Labor Quality." *The Journal of Human Resources* 18 (3): 425–431.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 169–235.

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): 1–25.

Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *The Quarterly Journal of Economics* 133 (3): 1331–1393.

Welch, Finis. 1969. "Linear Synthesis of Skill Distribution." *The Journal of Human Resources* 4 (3): 311–327.

Online Appendix

Appendix A   Additional Analyses

This appendix presents additional results that are supplementary to those contained in the main text. First, I show additional versions of Figure 1 from Section 3.3 for alternative specifications of the Γ matrix. Next I present reduced form evidence for labor supply spillovers of the sort hypothesized in Section 3.5. After that, 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. Finally, I present the changes in the estimated skill distribution between 1984-89 and 2002-06.

A.1   Further Divisions of Partial Equilibrium Shock Space under Different Γ Matrices

Figure A1 produces additional versions of Figure 1 under alternative specifications of the Γ matrix of skills. As in Section 3.3 of the main text, this figure considers a 2-type, 2-occupation ($K = J = 2$) version of the labor supply model, calibrating the average non-pecuniary benefit of each occupation to be $\xi_1 = \xi_2 = -1$, the variance of the idiosyncratic non-pecuniary benefit to be $\nu_1 = \nu_2 = 0.5$, and each type of worker occupies one half of the population $m_1 = m_2 = 0.5$. Workers have log utility.

Panels A and B considers shocks from a base in which the two occupations have the same skill price, i.e. $w_1 = w_2 = 1$. Panel A considers a case where the skill distribution features only comparative advantage, while Panel B considers a case where there is only absolute advantage/a worker fixed effect. In both of these specifications, there is no region in which employment and wages can diverge from one another. In the case of pure comparative advantage, when initial labor prices are identical to one another, all worker types earn the same wage on average. Thus there is no composition effect (because $\tilde{\omega}_j = \bar{\omega}$). Likewise, since workers and jobs are ex ante symmetric and there are only two job types, a positive reallocation effect for one worker type is exactly offset by a negative reallocation effect for the other worker type. Thus the economy

Figure A1: Additional Splits of Shock Space by Direction of Aggregate Wage and Employment Movements in Economy with Two Job Types

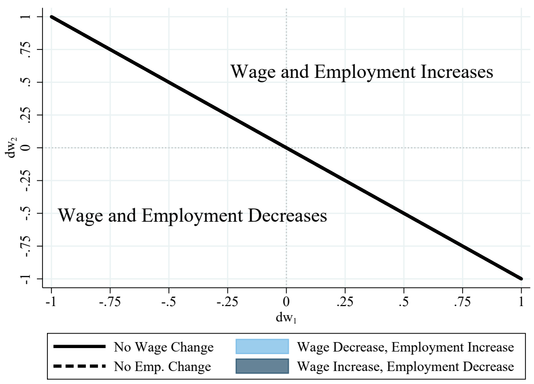

Panel A: $\Gamma^{CA} = \begin{pmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{pmatrix}$ and equal prices $w_1 = w_2 = 1$

Panel B: $\Gamma^{AA} = \begin{pmatrix} 1.5 & 1.5 \\ 0.5 & 0.5 \end{pmatrix}$ and equal prices $w_1 = w_2 = 1$

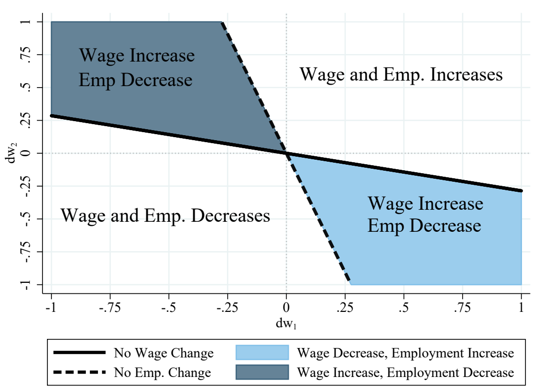

Panel C: $\Gamma^{Both} = \begin{pmatrix} 0.75 & 0.25 \\ 1.25 & 1.75 \end{pmatrix}$ and equal prices $w_1 = w_2 = 1$

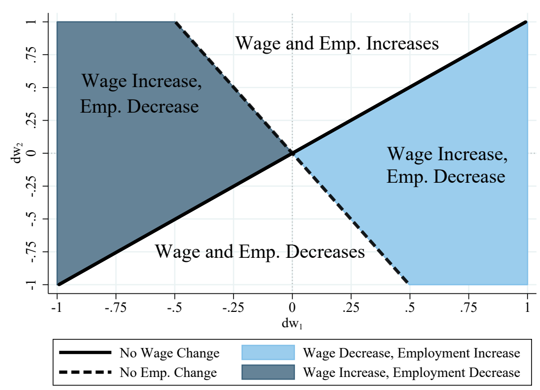

Panel D: $\Gamma^{Both} = \begin{pmatrix} 0.75 & 0.25 \\ 1.25 & 1.75 \end{pmatrix}$ and $w_1 = 0.75,\ w_2 = 1.5$

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) and in the text below. Regions constructed from equations (15) and (16). Text describes additional calibration.

behaves as if there were a representative agent when all prices are ex ante equal and workers have symmetric comparative advantages.

Meanwhile, if workers have symmetric absolute advantage and prices are equal, then shocks induce no reallocation effects because workers earn the same amount in each occupation. While there may be a composition effect if there is a correlation between employment probabilities and average wages, this does not affect the slope of the zero-wage line as shocks to both w_1 and w_2 will induce the same composition effect. Thus, in this knife edge case where all occupation prices are identical and workers have perfectly transferable skills between those occupations, there also exists no combination of first order labor price shocks that induce negative comovements between employment and wages.

For the slope of the line to change, it must be that worker types are not symmetrically affected by labor price shocks. There are two ways to accomplish this. The first is to start from a point where there is ex ante price dispersion. This possibility is explored in the main text – Panels B through D of Figure 1 consider shocks from a base of $w_1 = 0.75$ and $w_2 = 1.5$. An alternative is to have a non-symmetric skill matrix. Panel C considers the shock space when the skill distribution features both absolute and comparative advantage. That is, suppose the Γ matrix is given by

$$\Gamma^{Both} = \begin{pmatrix} 0.75 & 0.25 \\ 1.25 & 1.75 \end{pmatrix}$$

Under this skill matrix, type 1 workers are both low skill and have relative expertise in occupation 1. As a result, they will both earn lower wages on average (opening the possibility for composition effects) and be relatively more exposed to shocks in occupation 1. As a result, positive shocks to occupation 1 induce a negative composition effect as more low type workers respond to that shock. If a positive shock to occupation 1 if offset by a negative shock to occupation 2, it is possible that aggregate wages would increase as employment declines, even if labor prices were ex ante identical.

Finally, Panel D combines all of these ingredients. It uses the skill matrix combining both absolute and comparative advantage as well as allowing for ex ante heterogeneity in labor prices, assuming shocks hit the economy from a base of $w_1 = 0.75$ and $w_2 = 1.5$. In this case, there is a huge region of shocks which generate negative employment and wage comovements. Now, low type workers have skill in the low-price occupation 1. They thus sort to this occupation, realizing very low wages on average. Thus a positive shock to occupation 1 exerts a very large composition effect on the aggregate wage, so large in fact that a positive shock to occupation 1 induces wage declines and employment increases on its own.

Could a positive shock to one type of job plausibly induce negative aggregate wage effects in the real world? Consider the following simple back-of-the-envelope exercise. As above, suppose that there are two types of job and two types of worker – high type H and low type L. For simplicity, let us assume away the reallocation effect, which one can do by supposing that workers have perfectly specific skills so high type workers only earn wages in high type jobs, and vice versa for low type workers. Then a shock to the low type job induces an aggregate wage increase if the composition effect more than compensates for the direct effect, that is (from equation (16)):

$$1 + \frac{u'(c_L)}{\nu_L}(1 - E_L)(\tilde{\omega}_L - \bar{\omega}) < 0 \tag{A1}$$

Now suppose that $u'(c_L) = \tilde{\omega}_L^{-1}$ and the employment rate of low type workers is 50%. As a point of comparison, the employment-to-population ratio of those 25 years or older with less than a high school diploma has hovered around 40%, while that of those with a high school diploma ranged from 60% to 55% between 2001 and 2019. Substituting these suppositions into equation (A1) and rearranging yields the following condition

$$\frac{\bar{\omega}}{\tilde{\omega}_L} - 1 > 2\nu_L.$$

Therefore, if $\nu_L = 0.5$ (implying a labor supply elasticity of around 2), a shock affecting only low skill workers will reduce aggregate wages if low skill workers earn no more than one half of the average wage. According to data from the BLS,[26], the median weekly earnings of those with less than a high school education was \$520, compared with \$1,305 for those with a Bachelor's Degree, and \$907 for the whole economy in 2018. Considering mean earnings yields a starker picture, with those having less than a high school diploma earning an average of 29 thousand dollars, compared with over 58 thousand for the whole economy. Thus the lowest skill workers earn approximately half that of average workers. This back-of-the-envelope

26. See https://www.bls.gov/careeroutlook/2018/data-on-display/education-pays.htm.

calculation suggests that isolated shocks affecting only one segment of the economy could plausibly induce negative comovements between employment and wages. Of course, the possibility that multiple shocks might offset one another makes it much more likely that such behavior could be realized in reality, as shown in the main text of the paper.

A.2 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. The skill matrices are recovered using the procedure of section 4.2.1 and the TFP series are calculated using the procedure of 4.2.2. Details included in main text. Table A1 shows the estimated Γ matrix from the period 2002-2006. Table A2 shows the estimated selection-adjusted TFP series and the raw BLS KLEMS project TFP series. Table A3 plots the estimated Γ matrix from the period 1984-89.

Figure A2 plots counterfactual evolutions of 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. The black dash-dot line shows the evolution of employment and wages in a model in which a skill distribution estimated using CPS data from 1984-89 were subjected to the TFP shocks of the Great Recession: it studies the impact of the Great Recession's labor demand shocks were they to occur 20 years earlier. The green dashed line shows the evolution of employment and wages were all sectors subject to a 5.9% decline in TFP.

Figure A3 is an industry-level scatter plot in which the horizontal axis is the share of that industry's wage bill accounted for by "manual occupations," and the vertical axis is the selection-adjusted TFP series between 2008 and 2009 (Panel A) or 1990 and 1991 (Panel B). 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.

Table A1: 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 4.2.1, and carried out using data from 2002-2006 in the CPS.

Table A2: 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 (23). 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.

Figure A2: Predicted Wage and Employment Dynamics in Great Recession under Counterfactual Skill Distributions and Labor Demand Shocks

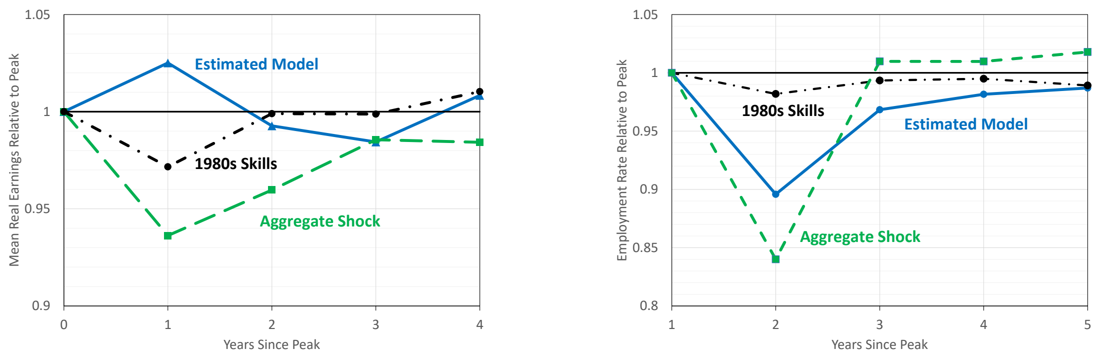

Panel A: Wages

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.

Figure A3: Industry TFP shocks by Share of Wage Bill in Manual Occupations

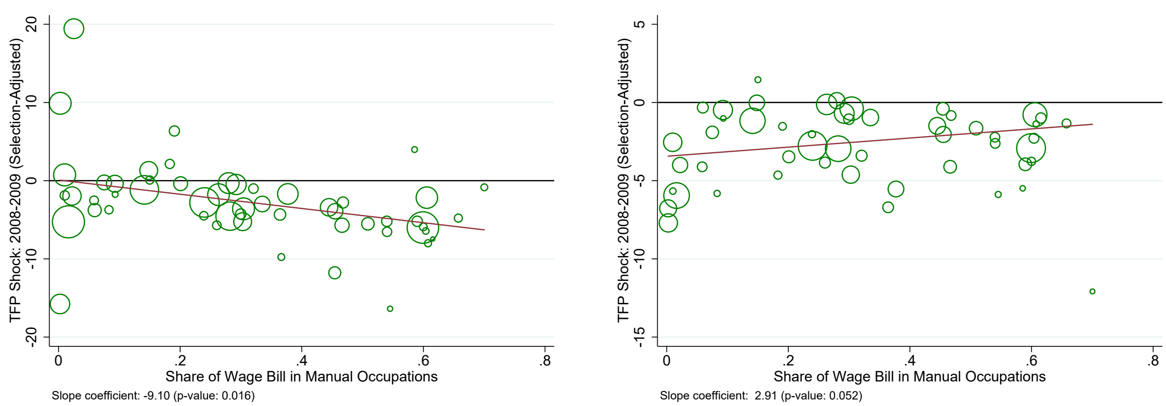

Panel A: 2008-2009

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. Wage bill shares calculated as of 2002-2006 in the Occupational Employment Statistics

Table A3: 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 4.2.1, and carried out using data from 1984-1989 in the CPS.

Table A4: In-Sample Model Fit, 2002-2006

	Emp. Shares		Mean Log Wage		SD Log Wage	
	Model (1)	Data (2)	Model (3)	Data (4)	Model (5)	Data (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.3 In-Sample Model Fit

This section presents measures of model fit in-sample. Table A4 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 correlation between the model and data columns are very close to 1. Table A3 shows a similarly strong ability to match the employment and wage data in the 1984-89 estimation window.

Table A5: In-Sample Model Fit, 1984-1989

	Emp. Shares		Mean Log Wage		SD Log Wage	
	Model (1)	Data (2)	Model (3)	Data (4)	Model (5)	Data (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.

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 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 very good predictor of source-destination switching pairs. However, the model implies an overall switching probability which is roughly twice that in the data. This is a limitation of the largely standard assumption that ζ_{ikt} are i.i.d.

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

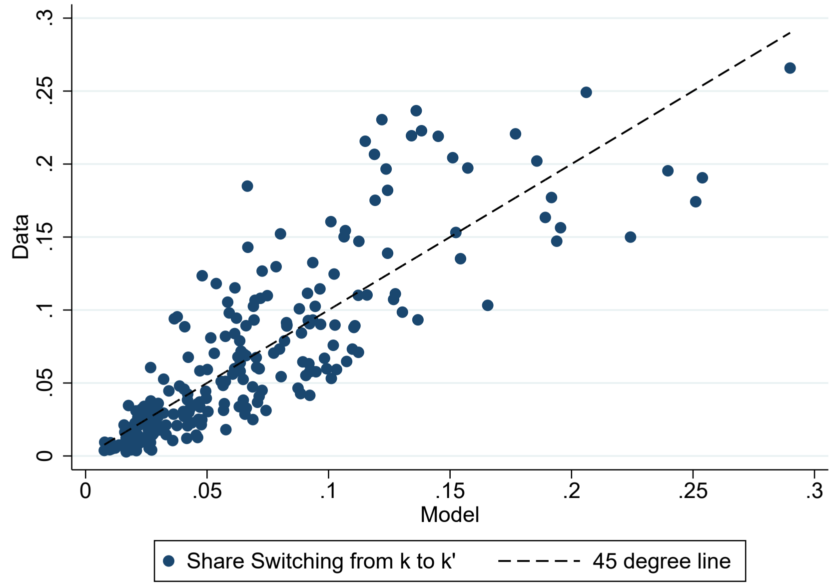

Notes: Figure shows the model fit of occupation switching patterns. Each dot represents a pair of occupations (k, k') and the axes plot the share of occupation switchers who are in that (k, k') pair in the estimated model (horizontal axis) and the data (vertical axis). The dashed line shows a 45 degree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

A.4 Changes in Skill Distribution Over Time

This section further examines the estimated changes in the skill distribution between 1984-89 and 2002-06. As suggested in that section, consider two key variances in the skill distribution:

$$Var^{BTWN} := \sum_{j=1}^{J} m_j(\mathbb{E}_k[\gamma_{jk}] - 1)^2; \qquad Var^{WTHN} := \sum_{j=1}^{J} m_j Var_k(\gamma_{jk}), \tag{A2}$$

Figure A5 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,

Figure A5: Absolute and Comparative Advantage: 1984-1989 and 2002-2006

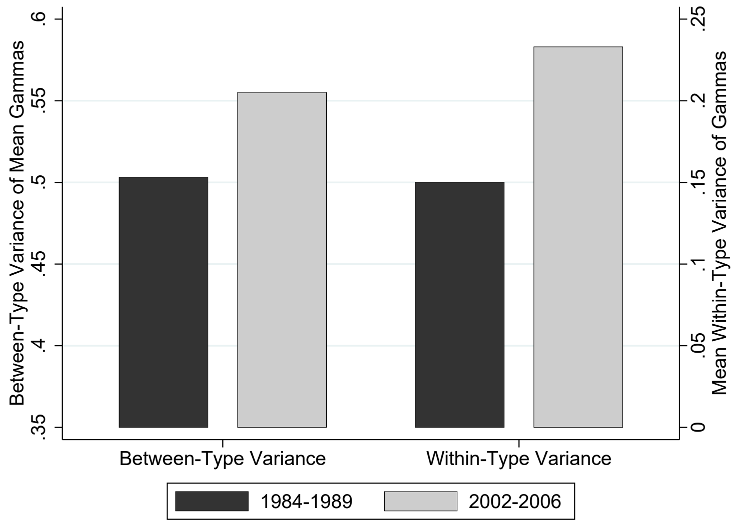

Notes: Figure plots the estimated within and between type variance of skills in the economy, captured by the Γ matrix of Table A1 and Appendix Table A3. Estimation follows the procedure outlined in Section 4.2.1, and carried out separately in the CPS March Supplement for the periods 1984-89 (gray bars) and 2002-2006 (black bars). Within and between variance defined as in equation (A2).

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 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 variance in their skill. This negative covariance is driven by an inability to engage in the 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

Figure A6: Correlation of Skill Relatedness in Γ with Euclidean Skill Distance in O*NET

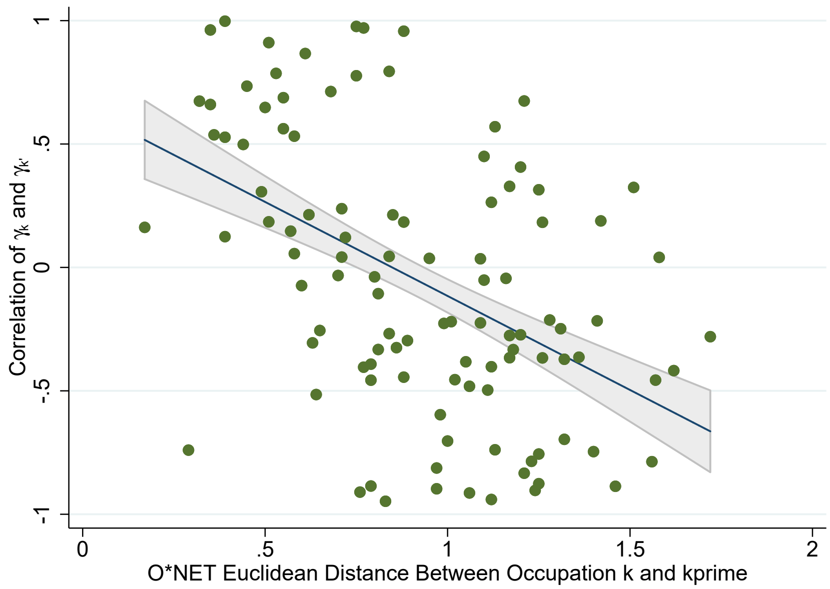

Notes: Figure compares the structurally-estimated skill transferability from the model to a common measure of skill relatedness from O*NET. Each dot corresponds to a pair of occupation clusters (k, k'). Occupations clustered by O*NET skill and knowledge vectors within terciles of the share with at least 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 A7. Line of best fit reported, with shaded area representing 95% confidence interval using White heteroskedasticity robust standard errors.

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.

I calculate these cross-occupation skill correlations in both the 1984-89 and 2002-06 estimated skill distributions. As a validation check, Figure A6 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 A7 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 figure shows numerous interesting patterns. First, 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.

The comparison between 1984-1989 and 2002-2006 is also instructive. In the late 1980s, skills were highly transferable across high-skill occupations, as represented by the large amount of blue squares in the bottom right corner of the correlogram. This reflects that workers would move between these high-skill occupations relatively frequently, realizing high wages in both source and destination occupations. In addition, skills were highly transferable across many of the low-skill tasks - the correlations between manual, routine, production, construction, and tradespeople jobs were all above 0.73, with the correlation between manual, production, and construction occupations reaching 0.93 or higher. When there were declines in construction demand,

Figure A7: Correlation of Occupation Skills, 1984-1989 and 2002-2006

Panel A: 1984-1989

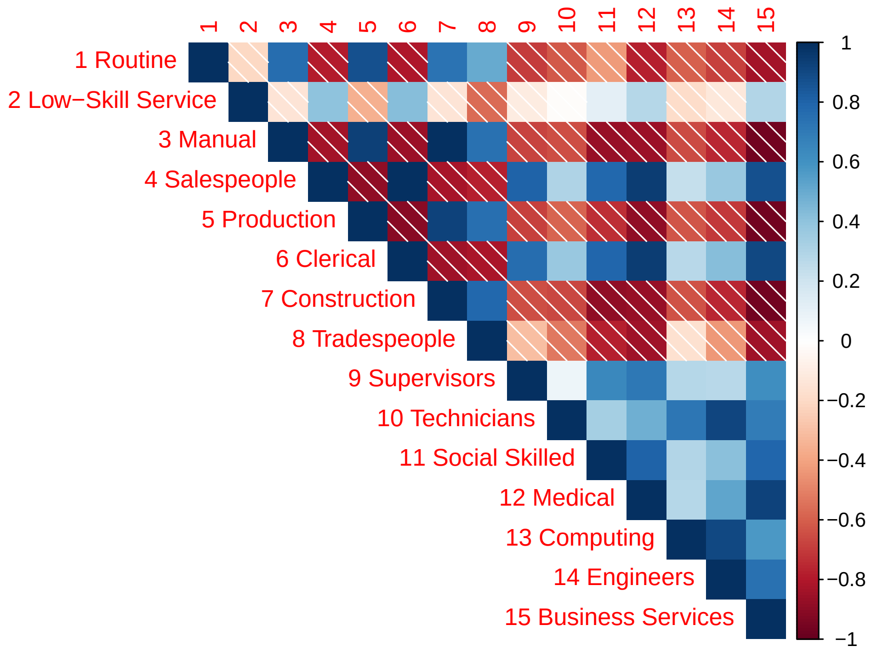

Panel B: 2002-2006

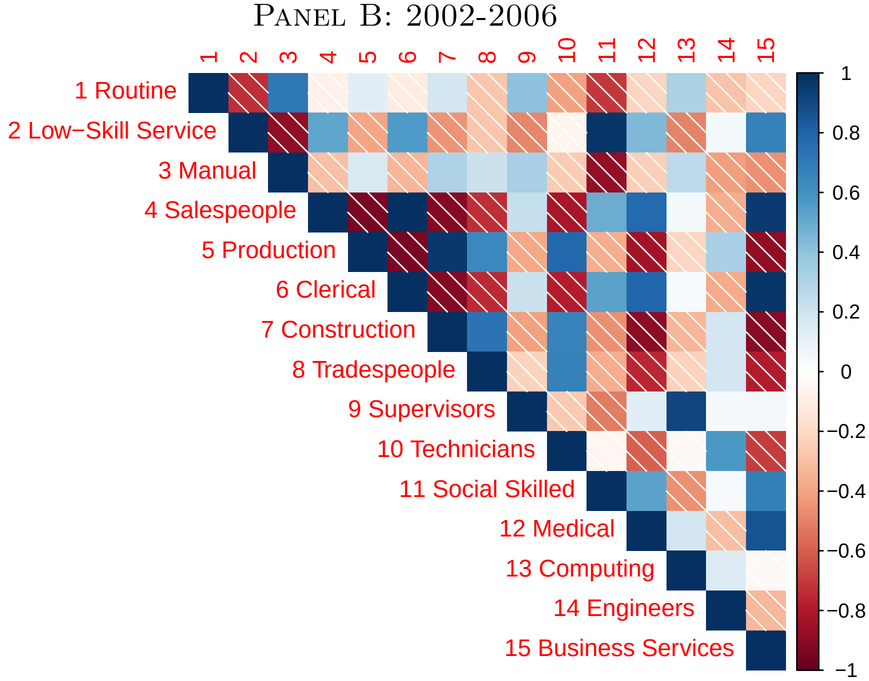

Notes: Figure plots the correlation of the row vectors of the estimated Γ, normalized by workers' mean skill in each occupation. Estimation follows procedure outlined in Section 4.2.1, and carried out separately in the CPS March Supplement for the periods 1984-89 (Panel A) and 2002-2006 (Panel B). Blue 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.

construction workers would exert substantial negative wage pressure on production line workers, as well as the routine manual occupations.

By 2002-2006, these patterns had changed. The skill correlation between manual, routine, construction, and production occupations all fell. The high-skill occupations became more specific, with correlations falling throughout the bottom right of the correlogram. In addition, many of the occupations that employ soft skills such as salespeople, clerical workers, and those occupations employing social skills such as teachers and lawyers, saw declines in skill correlation.

A.5 Reduced Form Composition Effects: 2008-09

Figure A8 provides reduced form evidence for strong composition effects 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 A8 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 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

Figure A8: Composition Effects Within and Between Occupations, 2008-2009

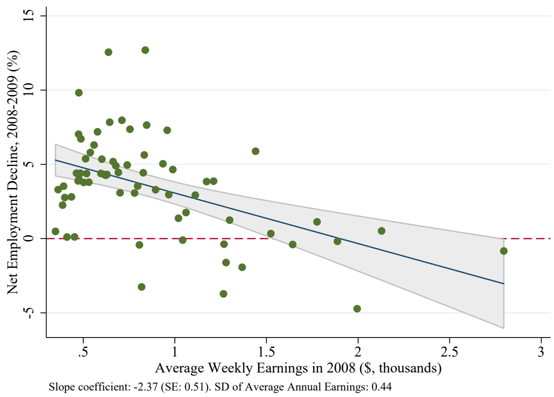

Panel A: Between Effects

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.

across occupations during the Great Recession.

Appendix B Model Appendix

B.1 Deriving Aggregate Responses to Shocks in Model

This subsection derives equations (15) and (16). Differentiating the sum of the occupation choice probabilities from equation (8) with respect to price of labor w_k in k yields the response of type j employment to w_k. From the quotient rule, this is

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

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

$$= \left(\frac{u'(\omega_{jk})\gamma_{jk}}{\nu_j}\right)E_{jk}(1-\tilde{E}_j) \tag{A3}$$

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:

$$\begin{aligned}\frac{d\ln \bar{E}}{d\ln z} &= \sum_j \left(\frac{m_j \tilde{E}_j}{\bar{E}}\right) \sum_{k\neq 0} \frac{d\ln \tilde{E}_j}{d\ln w_k} \cdot \frac{d\ln w_k}{d\ln \mathbf{z}} \\ &= \sum_j \left(\frac{m_j \tilde{E}_j}{\bar{E}}\right) (1-\tilde{E}_j)\tilde{\omega}_j \sum_{k\neq 0} \left(\frac{u'(\omega_{jk})}{\nu_j}\right) \left(\frac{\omega_{jk} E_{jk}}{\tilde{\omega}_j \tilde{E}}\right) \frac{d\ln w_k}{d\ln \mathbf{z}} \end{aligned} \tag{A4}$$

which is equation (15) 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. 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:

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

which is the expression in equation (16). This derivation involves plugging in equations (A3) and (A4) into equation (15) and rearranging.

Finally to calculate the reallocation effect, differentiate the occupation choice probabilities from equation (8) with respect to $w_{k'}$:

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

Plugging this and equation A3 into equation (6) and rearranging obtains the result.

B.2 Characterizing and Computing Equilibrium

Consider the problem of the sector s firm. The first order conditions for optimality for this firm is given by

$$l_{sk} = \frac{p_s x_s z_s \alpha_{sk} \left(\prod_{k'=1}^{K} l_{sk'}^{\alpha_{sk'}} \right)^{x_s}}{w_k} \tag{A5}$$

Divide the equivalent expression for $l_{sk'}$ by the above expression to arrive at

$$l_{sk'} = l_{sk} \left(\frac{\alpha_{sk'}}{\alpha_{sk}} \frac{w_k}{w_{k'}} \right) \tag{A6}$$

Substitute this into equation (A5) to arrive at

$$l_{sk}^{1-x_s} = p_s x_s z_s \left(\frac{\alpha_{sk}}{w_k} \right)^{1-x_s} \left(\prod_{k'=1}^{K} \left(\frac{\alpha_{sk'}}{w_{k'}} \right)^{\alpha_{sk'}} \right)^{x_s} \tag{A7}$$

To save on notation, let $M_s := \prod_{k'=1}^{K} \left(\frac{w_{k'}}{\alpha_{sk'}} \right)^{\alpha_{sk'}}$. M_s is the marginal cost of a cost-minimizing firm with a constant returns to scale Cobb-Douglas production function.

Using the demand curve for sector s's production, substitute in for p_s to arrive at

$$l_{sk}^{1-x_s} = \frac{(Y)^{\frac{1}{\eta}} x_s z_s \left(\frac{\alpha_{sk}}{w_k} \right)^{1-x_s}}{M_s^{x_s} y_s^{\frac{1}{\eta}}} \tag{A8}$$

Plugging equation (A6) into the production function for sector s reveals that

$$y_s = z_s \left(\frac{M_s \alpha_{sk}}{w_k} \right)^{-x_s} l_{sk}^{x_s} \tag{A9}$$

which we may then substitute into the amended first order condition (A8)

$$l_{sk}^{\eta - x_s(\eta-1)} = Y x_s^{\eta} z_s^{\eta-1} M_s^{-x_s(\eta-1)} \left(\frac{\alpha_{sk}}{w_k} \right)^{\eta - x_s(\eta-1)} \tag{A10}$$

We may do this same process for sector s' to arrive at an analogous expression for that sector. Divide this analogous sector's expression by the one for sector s to eliminate Y and see that, letting $\nu_s = \eta - x_s(\eta - 1)$

$$l_{s'k} = l_{sk}^{\frac{\nu_s}{\nu_{s'}}} \underbrace{\left(\frac{\alpha_{s'k}}{w_k}\right)\left(\frac{w_k}{\alpha_{sk}}\right)^{\frac{\nu_s}{\nu_{s'}}} \left[\left(\frac{x_{s'}}{x_s}\right)^{\eta}\left(\frac{z_{s'}}{z_s}\right)^{\eta-1}\left(\frac{M_s^{x_s}}{M_{s'}^{x'_s}}\right)^{\eta-1}\right]^{\frac{1}{\nu_{s'}}}}_{:=\psi_{s',s}} \tag{A11}$$

As a result, A9 implies that the equilibrium output in sector s' is given by

$$y_{s'} = z_{s'} \left(\frac{M_{s'}\alpha_{s'k}}{w_k}\right)^{-x_{s'}} \psi_{s',s}^{x_{s'}} l_{sk}^{\frac{x_{s'}\nu_s}{\nu_{s'}}} \tag{A12}$$

so that the output of final goods may be expressed as a function of l_{sk}:

$$Y(l_{sk}) = \left(\sum_{s'=1}^{S}\left[z_{s'}\left(\frac{M_{s'}\alpha_{s'k}}{w_k}\right)^{-x_{s'}} \psi_{s',s}^{x_{s'}} l_{sk}^{\frac{x_{s'}\nu_s}{\nu_{s'}}}\right]^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}} \tag{A13}$$

Finally, plug this into equation (A10) to have one equation in l_{sk} which may be solved numerically. Once this is done for some arbitrarily selected sector s and occupation k, we may use equations (A6) and (A11) to solve for the full system of occupation demands, given an exogenous productivity vector $\mathbf{z}$ and endogenous vector of wages $\mathbf{w}$. As a result, the aggregate demand for occupation k is given by summing over the demands from each of the sectors:

$$L_k^D(\mathbf{w}|\mathbf{z}) = \sum_{s=1}^{S} l_{sk}(\mathbf{w}|\mathbf{z}) \tag{A14}$$

Labor supply of occupation k is given by the total labor units supplied to k by the J worker types. That is, supply of services for occupation k is given by

$$L_k(\mathbf{w}) = \sum_{j=1}^{J} m_j \gamma_{jk} E_{jk} \tag{A15}$$

One may solve for equilibrium by equating labor demand for occupation k, given by equation (A14), with the labor supply for this occupation, given by equation (A15). Note that since workers do not have preferences over which sector to work for, and because workers are perfect substitutes within an occupation conditional on their units of effective labor, the law of one price will hold within each occupation. These occupation prices will determine the quantities of effective labor in each occupation employed by each sector. Furthermore, Walras' Law implies that equating the labor demand and labor supply in each occupation will imply that final goods clearing is also satisfied. That is, total income, given by

$$C = \underbrace{\sum_{j=1}^{J} m_j \sum_{k=1}^{K} \gamma_{jk} w_k E_{kj}}_{I} + \underbrace{\sum_{n=1}^{N} (1 - x_s) p_s y_s}_{\Pi} \tag{A16}$$

will equal aggregate output given by equation (A13).

Note that the structure of the model implies that one need only solve for the K occupation prices in order to characterize the equilibrium. For this reason, one can consider sectors at a fine level of aggregation without adding substantial computational burden. Indeed, the model features both sectors and occupations separately primrarily to allow for sufficient granularity in the construction of labor demand shocks from the data.

To compute equilibrium, I employ the R package nlopt's implementation of the Improved Stochastic Ranking Evolution Strategy (ISRES) optimizer to minimizer the largest squared difference between labor supply (A15) and labor demand (A14) subject to a choice of wage vector $\mathbf{w}$. I additionally include the squared difference between aggregate output and consumption as an equilibrium condition, as doing so improves performance of the optimizer. The ISRES routine is a semi-global optimization method put forward by Runarsson and Yao (2005). Arnoud et al. (2019) finds that ISRES performs well in many economic applications.

Appendix C Estimation Details

Consider the likelihood of observing a single worker i for two periods, labeled 1 and 2. This worker chooses occupation k in period 1 and k' in period 2, realizing wages ω_{i1}

and ω_{i2} in periods 1, and 2, respectively. Let the parameters of the model be given by θ, which includes γ_{jk}, ξ_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 her occupation choices and wages is given by the probability that her type made her 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:

$$\mathcal{L}(\theta) = \sum_i \sum_{k=0}^{K} \sum_{k'=0}^{K} \mathbf{1}\{k_1(i) = k\}\mathbf{1}\{k_2(i) = k'\} \ln l_i(k, k', \omega_{i1}, \omega_{i2}|\theta) \tag{A17}$$

A formal proof of identification is identical to that provided in Appendix A of Bonhomme et al. (2019) and the interested reader may find a version of the argument tailored to my particular setting in the working paper version of this manuscript.

Appendix D 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.[27] 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 4.2.1 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. 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 important input into the

27. Solon et al. (1994) highlights important differences in the cyclicality of real wages for men and women between 1967 and 1987.

Table A6: 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 A7: 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 A8: Largest Employment SOC Codes within Occupation Clusters, Set 3

Clust #	SOC Title Examples	Mean Income
11 Social Skilled	Elementary and Middle School Teachers	56909
	Secondary School Teachers	58491
	Other Teachers and Instructors	36646
	Postsecondary Teachers	74068
	Special Education Teachers	58420
	Designers	46437
	Lawyers	126710
	Human Resources Workers	61057
12 Medical	Registered Nurses	68801
	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 Software/ Computing	Computer Support Specialists	53141
	Software Developers, Systems Software	104103
	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 Engineers	Industrial Engineers, Including Health and Safety	83202
	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 Managers/ Skilled Business Services	General and Operations Managers	115124
	Accountants and Auditors	71718
	Sales Representatives, Services, All Other	61414
	Financial Managers	124469
	Management Analysts	87539
	Securities, Commodities, and Financial Services Sales Agents	102509
	Financial Analysts	90968
	Education Administrators	90877

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 A6-A8 report additional 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 according to the BLS is also listed.

References

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg. 2019. "Benchmarking Global Optimizers." *NBER Working Paper Series,* no. 26340.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A Distributional Framework for Matched Employer Employee Data." *Econometrica* 87.

Flood, Sarah M., and José Pacas. 2008. "Using the Annual Social and Economic Supplement as Part of a Current Population Survey Panel." *Census Bureau Working Paper,* 1–36.

Poletaev, Maxim, and Chris Robinson. 2008. "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000." *Journal of Labor Economics* 26 (3).

Runarsson, Thomas Philip, and Xin Yao. 2005. "Search biases in constrained evolutionary optimization." *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 35 (2): 233–243.

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): 1–25.