Skill Heterogeneity and Aggregate Labor Market Dynamics*

John Grigsby

This paper is evolving. For the latest version, please click here.

November 12, 2019

Abstract

What determines the joint dynamics of aggregate employment and wages over the medium run? This classic question in macroeconomics has received renewed attention since the Great Recession, when real wages did not fall despite a crash in employment. This paper proposes a microfoundation for the medium-run dynamics of aggregate labor markets which relies on worker heterogeneity. I develop a model in which workers differ in their skills for various occupations, sectors employ occupations with different weights in production, and skills are imperfectly transferable. When shocks are concentrated in particular industries, the extent to which workers can reallocate across the economy determines aggregate labor market dynamics. I apply the model to study the recessions of 2008-09 and 1990-91. I estimate the distribution of worker skills using two-period panel data prior to each of these recessions and find that skills became less transferable between the 1980s and 2000s. Shocking the estimated model with industry-level TFP series replicates the increase in aggregate wages in 2008-09, and decline in 1990-91. The model implies that if either the composition of industry shocks or the distribution of skills in the economy had been the same in the 2008-09 recession as in the 1990-91 recession, real wages would have fallen, while employment would have declined less. The declining industries during the 2008-09 all employed a similar mix of skills, which induced many low-skill workers to leave the labor force and limited downward wage pressure on the rest of the economy. Finally, the model inspires a novel reduced form method to correct aggregate wages for selection in the human capital of workers, which accounts for cyclical job downgrading by focusing on the wage movements of occupation-stayers. This correction recovers pro-cyclical wages, suggesting the changing composition of the workforce was crucial for aggregate wage dynamics during the Great Recession.

*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 Stéphane Bonhomme, Ali Hortaçsu, Thibaut Lamadon, Simon Mongey, Kevin Murphy, Jeremy Pearce, Gustavo Souza, Tom Winberry, Liangjie Wu, and Yulia Zhestkova, as well as seminar participants in University of Chicago's Capital Theory, Applications of Economics, and Applied Macro Theory workshops. Author's contact information: jgrigsby@uchicago.edu

1 Introduction

What determines the medium-run dynamics of aggregate employment and wages? Providing a compelling answer to this classic question is complicated by the fact that these dynamics are not constant over time. For example, the US recessions of the 1970s, 80s and early 90s saw real wage declines between 1.5 and 5.4 percent. In contrast, during the Great Recession of 2008-09, real average hourly earnings rose despite a crash in employment.[1] Existing macroeconomic theories generally attribute changes in the relationship between aggregate employment and wages to exogenous fluctuations in preferences or frictions, such as wage rigidity (Christiano et al., 2005; Smets and Wouters, 2007). While useful, these representative agent theories cannot speak to a sizable empirical literature suggesting that muted aggregate wage fluctuations largely result from shifts in the composition of the workforce that arise from low-skill workers leaving the employed pool in a downturn (Solon et al., 1994; Daly et al., 2011; Devereux, 2001). Indeed, quantitative general equilibrium theories generating medium-run fluctuations in the composition of the workforce are scarce. Without such a theory, it remains difficult to predict the conditions under which worker heterogeneity has a large influence on the behavior of the aggregate labor market.

In this paper, I propose a microfoundation for the medium-run behavior of aggregate employment and wages which relies on worker and firm heterogeneity. Sectors experience uneven shocks. If a particular sector collapses, the extent to which its workers can apply their skills to alternative pursuits will govern the aggregate impact of that shock. If workers are unable to reallocate to unshocked sectors, a sectoral decline will induce large flows into non-employment, with little impact on the rest of the economy. On the other hand, if workers may easily reallocate themselves, the workers expelled from a declining sector may seek employment in other industries, reducing the aggregate employment impact of the shock, but exerting downward wage pressure on the rest of the economy. The aggregate response to a shock, over a horizon in which skills are fixed, will therefore depend on both the sectoral composition of that shock and the distribution of skills in the labor force.

I build and estimate a macroeconomic model in which multiple industries employ workers in a variety of occupations to produce output. The key innovation is that labor is supplied by workers who belong to one of a discrete set of skill types, characterized by a vector describing the effective human capital that the worker can supply to each occupation. The model nests multiple common representations of the skill distribution, such as representative agent economies, or a model in which workers have specific skills that may only be applicable

[1] According to the Current Employment Statistics (CES) provided by the Bureau of Labor Statistics. See Appendix Table A1 for wage, employment, hours, and price index changes for the last six US recessions.

in one occupation.[2] Workers choose whether to supply their labor to the market and, if so, their occupation according to a standard Roy Model. Industries combine occupations with different weights in their production function and are subject to occupation-neutral total factor productivity (TFP) shocks.

A decline in a particular sector's TFP in this setup has three effects. The first effect is common to many models - a decline in a sector's TFP lowers the employment and price of occupations heavily employed by that sector. Here, however, there is an additional effect arising from labor supply spillovers: workers displaced from the declining occupation exert downward wage pressure on other occupations in the economy. The strength of this spillover is dictated by the extent to which skills are transferable from declining occupations to growing occupations. Finally, there is a selection effect. As the price of labor declines in a set of occupations, workers employed in those occupations may choose to leave employment. If these expelled workers are generally low-skill, the decline in sectoral TFP will induce positive selection in the set of workers employed, pushing up the measured average wage. Indeed, if the skill gap between low- and high-skill workers is sufficiently large, and the workers employed in the declining sector are generally low-skill, this selection force could generate increases in measured aggregate wages from sectoral declines in TFP.

The model remains tractable enough to be estimated by building off the distributional framework of Bonhomme et al. (2019). By observing the inter-occupation mobility patterns of workers, as well as the wages before and after the occupation switch, the econometrician can recover the distribution of types, as well as the mean and variance of the wages in every occupation for each type of worker. Intuitively, the principal determinant of wage changes for workers who switch occupations is their occupation-specific skill vector. The approach consistently estimates these parameters of interest in two-period panel data, under some standard rank and exogeneity conditions.

I apply the model to study the US recessions of 2008-09, which experienced increases in real wages and a crash in employment, and 1990-91, which saw both wage and employment declines. I estimate the distribution of latent skill types and their returns to different occupations using the March supplement of the Current Population Survey before the each of these recessions. The estimation reveals that skills have become less transferable, with the variance of skills growing both within workers across occupations, and across workers. The elasticity of non-employment to changes in the price of occupational services varies greatly across occupations, suggesting that the identity of shocked occupations will matter for employment fluctuations.

[2]See, for example Alvarez and Shimer (2012), Kambourov and Manovskii (2009a), Cosar (2013), and Adão (2019) for examples of models with occupation-specific human capital.

The Roy model of occupation choice permits explicit estimation of the degree to which employed workers are selected along unobservable human capital levels at business cycle frequencies. My estimates show countercyclical positive selection in the quality of employed workers: on average, low quality workers leave employment during a recession. However, these patterns differ greatly for each occupation. While production workers become more positively selected in recessions, the selection of medical workers exhibits no cyclical pattern. The estimates imply that the selection patterns have changed so that the average human capital level of employed production workers increased by 50% between 2008 and 2009, compared with just 14% between 1990 and 1991. In aggregate, the mean human capital of employed workers rose by 10% from 2008-09, compared with just 4% in 1990-91, suggesting that positive selection propped up aggregate wages by a further 6% in 2008-09.

Explicitly estimating the selection in the unobserved human capital level of employed workers allows model-consistent estimation of industry-level total factor productivity series. Standard measures of TFP, such as those provided by the Bureau of Labor Statistics' (BLS) KLEMS project, do not account for the selection of employed workers according to their unobservable skill. Purging estimated industry-level TFP series of cyclical selection patterns of employed workers reveals a substantially more procyclical TFP series than is offered by the BLS' numbers. While the raw BLS TFP series reveals no cyclical change in the productivity of the construction sector in the 2008-09 recession, the series adjusted for human capital selection reveals a 6% decline in construction TFP, and an aggregate TFP shock which is 2 percentage points larger than that implied by the raw series.

The model replicates the unusual patterns of aggregate wages and employment during the Great Recession. I calibrate the macro model to match the estimated skill distributions before the 1990-91 and 2008-09 recessions and hit the economy with the realized sequence of selection-corrected TFP shocks at the 3-digit NAICS level.[3] Even though it relies solely on TFP shocks, the model reproduces the procyclicality of wages in the early 90s, and the countercyclical wages around 2009.

The change in labor market dynamics may arise in the model due to changes in either the skill distribution or industry shock composition. The model implies that if the shocks of 2009 had hit the distribution of skills of the early 1990s, real wages would have fallen 3 percent with employment falling 2 percent. Meanwhile, if the recession of 2009 had arisen from the industry shocks of 1991, then real wages would have declined approximately 8 percent, with employment declines of 2 percent. This implies that the 2009 recession was unique in that multiple sectors, all of which employ the same low-skill workers, declined at once, limiting

[3]These TFP shocks are isomorphic to preference shocks inducing fluctuations in the demand for sectoral output.

the ability of these low-skill workers to supply their labor elsewhere in the economy.

Finally, the model suggests a novel reduced form approach to correcting aggregate wage series for the selection of workers employed during the cycle. Existing approaches generally assume workers' skills are determined by a worker fixed effect: while some workers are persistently high-earners, others are low-earners. In this paper's framework, workers differ in skills for a variety of occupations. As a result, they may choose to apply their skills to tasks to which they are worse suited in response to movements in occupational labor prices – manufacturing workers may become cashiers in a downturn, or a shale gas boom may attract workers with little mining ability. Considering the wage changes of occupation-stayers isolates wage movements arising from the price of labor by holding fixed workers' on-the-job human capital, after controlling for standard tenure effects. Holding fixed the composition and allocation of workers using this method restores the pro-cyclicality of aggregate wages in the Great Recession, suggesting an important role for composition bias. However, this new composition adjustment generates similar wage pro-cyclicality as the classic fixed-effect approach of Solon et al. (1994), suggesting that the changed allocation of workers to tasks had little effect on the cyclicality of wages in recent periods.

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 manifest themselves as a wedge between a representative agent's marginal rate of substitution (MRS) and the economy's marginal rate of transformation (MRT, 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. Bils et al. (2018) argue that the wedge between producers' MRT and wages is of roughly the same size as the wedge between wages workers' MRS, urging deviations from the baseline representative agent model on both the production and worker sides.

To rationalize these wedges, economists have principally considered the many frictions present in the labor market. An enormous literature considers the role of search frictions for the behavior of employment and wages.[4] Shimer (2005) points out, however, that standard calibration of such models struggles to match the joint movements of employment and wages in most recessions, and urges the consideration of models incorporating wage rigidity.[5] Many papers incorporating wage rigidity therefore followed (Hall, 2005; Schmitt-Grohé and Uribe, 2012). However, the size of labor wedge fluctuations have varied greatly across recessions. As a result, models calibrated to aggregate data estimate vastly different degrees of wage

[4] This literature is far too extensive to be adequately surveyed here. See Rogerson et al. (2005) for a classic survey.

[5] Hagedorn and Manovskii (2008) argue that a different calibration of classic search models based on the cost of vacancy creation and cyclicality of wages is able to jointly match aggregate employment and wages.

rigidity depending on the time period of the calibration. For instance, Christiano et al. (2005) estimate a New Keynesian dynamic stochastic general equilibrium (DSGE) model for the period 1965-1995 and find that 83.2% of workers can change their wage in a given year, while Christiano et al. (2014) estimate a monetary DSGE model augmented with a financial accelerator on the period 1985-2010, finding that just 57% of workers see a wage change in a given year. My model provides an alternative unifying framework to predict the behavior of the labor wedge across different time periods through variations in the degree of skill transferability out of declining sectors. The shifting dynamics of aggregate employment and wages that arise from the variable sectoral composition of shocks will manifest as fluctuations in the labor wedge in a representative agent economy.

Although the base wages of job-stayers display evidence of downward nominal rigidity (Grigsby et al., 2019), the microdata suggest that average hourly earnings cuts are relatively common (Kurmann and McEntarfer, 2019; Jardim et al., 2019). Using regional data, Beraja et al. (2019) argue that reasonable calibrations of nominal rigidity are insufficient to explain aggregate wage fluctuations during the Great Recession, arguing that labor supply shocks must have been a key feature of the period.

My paper provides a microfoundation for these aggregate labor supply shocks. In my 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, then the aggregate response of employment will be large relative to the response of labor prices. In addition, if workers expelled from employment as a result of a sectoral productivity shock are low-skill, the changing composition of the workforce will limit fluctuations in measured mean wages. In either case, standard 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. The volatility of these implied aggregate supply responses will therefore be larger the more heterogeneous are skills.

The role of selection in determining aggregate wage fluctuations was recognized by, among others, Solon et al. (1994). These authors studied the cyclical property of wages for a panel of workers in the Panel Survey of Income Dynamics (PSID) and found that wages were far more cyclical when one removes the influence of selection by considering a balanced panel of workers. This influential paper spawned a number of papers seeking to understand the cyclical selection patterns in the labor market (e.g. Gertler and Trigari (2009); Gertler et al. (2016)). My paper builds on this literature in two ways. First, my model shows how the selection arises endogenously as a result of heterogeneous sectoral shocks, and how that selection generates general equilibrium spillovers to unshocked sectors.[6] Second, the

[6] Hagedorn and Manovskii (2013) provides an alternative mechanism for procyclical selection in the labor

model suggests a novel reduced form method to correct for the selection of workers in an environment in which workers are both vertically and horizontally differentiated. Finally, I show how the distribution of skills may be estimated from the data, and therefore provide a predictive framework for the effect of particular combinations of sectoral shocks.

The paper proceeds as follows. Section 2 introduces the quantitative model with multiple skill types, and explores its implications in simple two-occupation, two-type frameworks. Section 3 describes the approach to estimating the model, including the details of the data used to do so. Section 4 presents the estimated skill distribution and the changing cyclical pattern of selection. Section 5 presents the results of the calibrated model, and estimates the importance of the changing skill distribution for the changing cyclical wage dynamics. Inspired by the model, section 6 proposes a simple reduced form approach to correcting aggregate wage series for the selection of workers employed. Section 7 discusses the model's implications in the context of other active debates in macroeconomics. Section 8 concludes.

2 Quantitative Model

This section builds a quantitative model with a non-trivial skill distribution which may be estimated using two-period panel data. The model features multiple sectors, each employing multiple occupations. Workers belong to one of a finite number of types and are each endowed with one unit of indivisible time. Types differ in the units of effective capital that they can supply to each occupation. Sectors hire workers in each occupation to produce output, which is sold to a competitive final goods producer. The final goods producer sells numeraire to a risk-neutral household sector.

2.1 Setup

Time is discrete. The economy consists of S sectors, indexed by s, each of which employs workers in K distinct occupations, indexed by k. Workers belong to one of J skill types, indexed by j. Neither workers, firms, nor households make dynamic decisions; therefore, the model may considered period-by-period.[7]

2.1.1 Households

There is a large representative household containing a measure 1 of infinitely-lived workers. The household is risk-neutral and consumes a final numeraire consumption good C. The

market in a search theoretic model in which the match quality of existing workers is predicted by the number of outside offers she has received during her tenure.

[7] In Appendix D, I discuss extensions to the model which capture the dynamic nature of worker decisions.

household takes as given income from labor $\bar{\omega}$, which is determined below, and from firms' profits Π, which it uses to finance consumption. The household additionally gains non-pecuniary benefits Ξ from the workers' activities, to be described in depth below. The household consumes its total income each period: $C = \bar{\omega} + \Pi$.

2.1.2 Intermediate Goods Firms

Each industry s is populated by a representative competitive firm. The firm hires workers into each of the K occupations in order to produce output y_s according to

$$y_s = z_s F^{(s)}(l_{s1}, l_{s2}, \ldots, l_{sK})$$

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. In the quantitative exercise below, I explore the economy's response to changes in the distribution of industry TFP z_s.

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

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

Total profits in the economy is the sum of all sectors' profits: $\Pi := \sum_{s=1}^{S} \pi_s$.

2.1.3 Workers

Workers, indexed by i, may be one of J types. Let the type of worker i be given by $j(i)$, and suppose that the mass of workers of type j is given by m_j. Each worker may supply her labor to only one of the K occupations in each period.

The J types of worker differ according to their skill in each occupation k. A worker of type j can supply γ_{jk} efficiency units of labor to occupation k. For notational simplicity, let Γ denote the matrix whose (j, k) element is γ_{jk}. Units of human capital are perfectly substitutable; therefore, the law of one price holds for occupational skill, and a worker of type j will earn $\gamma_{jk} w_{kt}$ if she were to work in occupation k.

One may think of these γ_{jk} as being a metaphor for the skill level of a type j worker in the various tasks employed by occupation k. For instance, if tax accountants require acumen in mathematics, economics, and tax law, those workers who are strong in these more fundamental skills will have a high γ for the accounting profession. Similarly, those who are manually dextrous will see higher γ's in carpentry or other manual occupations.

Workers' only decision is their occupation choice. In addition, each occupation provides some fixed non-pecuniary benefits ξ_k to workers.[8] 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 without loss of generality. 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 determined by workers solving

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

where $k = 0$ represents the non-employed state.[9]

Let $\mathbb{P}_k(j|\mathbf{w})$ 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 themselves across occupations and to non-employment. In turn, this produces selection in the types of workers employed in each occupation.

Conditional on the choice of occupation, workers are indifferent between sectors. The idiosyncratic preference shocks ζ_{ik} are assumed to be i.i.d. across workers and occupations. In particular, they as assumed to have marginal (cross-sectional) distribution which is type 1 extreme value with standard deviation ν. This assumption is standard in the discrete choice literature following McFadden (1974), and generates the tractable form for the cross-sectional choice probabilities of workers below:

$$\mathbb{P}_k(j|\mathbf{w}) = \frac{\exp\left(\frac{\gamma_{jk} w_k + \xi_k}{\nu}\right)}{\sum_{k=0}^{K} \exp\left(\frac{\gamma_{jk'} w_{k'} + \xi_{k'}}{\nu}\right)} \tag{3}$$

Aggregating workers' individual decision problems yields occupation-level labor supply

[8] Sorkin (2018) shows that approximately 40% of workers receive a wage cut when switching employers, and, as a result, estimates that non-pecuniary benefits account for over half of the firm component of the variance of earnings.

[9] Note that, since the household to which the worker belongs is risk-neutral, the dollar wage is the same as the utility wage for each worker. With strictly concave utility, there would be an additional income effect on labor supply, which makes workers less responsive to the dollar wage as total income increases. This would have the effect of making the aggregate labor supply curve less elastic as the economy grows.

curves. The mass of workers employed in each occupation E_k is

$$E_k(\mathbf{w}) = \sum_{j=1}^{J} m_j \mathbb{P}_k(j|\mathbf{w}) \tag{4}$$

This $E_k(\mathbf{w})$ schedule returns, for any set of labor prices, the measure of workers in each occupation. This quantity does *not* correspond to true labor supply curve in the model, but does match the employment concept generally measured in the data. Because workers differ in their effective labor units based on their type, the true labor supply curve in each occupation is instead given by the human-capital-weighted employment in each occupation:

$$L_k(\mathbf{w}) = \sum_{j=1}^{J} m_j \mathbb{P}_k(j|\mathbf{w}) \gamma_{jk} \tag{5}$$

When $\mathbf{w}$ moves, it may induce separation between $E_k(\mathbf{w})$ and $L_k(\mathbf{w})$ depending on the sets of workers who respond labor price changes. This changes the mean human capital of employed workers. Since workers are remunerated according to their human capital levels, this selection force can induce all manner of relationships between employment changes and mean earnings. Summing over each occupation yields the aggregate employment and labor supply curves, which depend on the vector of occupation prices $\mathbf{w}$:

$$E(\mathbf{w}) = \sum_{k=1}^{K} E_k(\mathbf{w}), \qquad L(\mathbf{w}) = \sum_{k=1}^{K} L_k(\mathbf{w}) \tag{6}$$

At the worker level, note that the earnings of type j workers is given by

$$\omega_j(\mathbf{w}) = \sum_{k=1}^{K} \mathbb{P}_k(j|\mathbf{w}) \gamma_{jk} w_k \tag{7}$$

where the symbol ω represents take home pay, which increases with worker skill, unlike the price of labor w. This equation shows that skill influences workers' earnings in two ways. The first is the direct effect: workers with high γ_{jk} earn higher wages from working in occupation k simply by virtue of being more productive in that occupation. This is an absolute advantage effect. In addition, there is a comparative advantage effect, that operates through $\mathbb{P}_k(j|\mathbf{w})$. Workers with higher γ_{jk} *relative* to $\gamma_{jk'}$ are more likely to work in occupation k. Mean wages are given by summing over each worker type's mean earnings

$$\bar{\omega}(\mathbf{w}) = \sum_{j=1}^{J} m_j \omega_j(\mathbf{w}) \tag{8}$$

2.1.4 Final Goods Producers

There is a representative competitive firm which produces numeraire using the output from each sector as inputs to a constant elasticity (CES) production function. That is, the output of the final good is given by

$$Y = \left(\sum_{s=1}^{S} \hat{y}_s^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \tag{9}$$

for $\hat{y}_s$ the demand for sector s's output from the final goods producer. As is standard with this specification, the demand curve for sector s's output is given by:

$$p_s = \left(\frac{Y}{\hat{y}_s} \right)^{\frac{1}{\eta}} \tag{10}$$

2.2 Equilibrium Definition

A static 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 $\{\mathbb{P}_k(j|\mathbf{w})\}_k, \{l_{sk}(p_s, \mathbf{w}|z_s)\}_{s,k}, \{\hat{y}_s(\mathbf{p})\}_s$ 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 (1),
2. The workers' occupation choice decisions are consistent with maximizing expected utility, solving (2),
3. The demand for each sectors output from the final goods producer $\hat{y}_s(\mathbf{p})$ is equal to the supply of that sectors 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: $Y = \bar{\omega} + \Pi$
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 E.

2.3 Discussion

Before considering the identification and estimation of the model, it is worth remarking on its structure. I first elaborate on its relation to the existing paradigms of skill specificity.

Next I provide intuition for the nature of labor supply by considering partial equilibrium responses to labor price changes in a two-occupation, two-type version of the model. Finally, I highlight the role of the structure of labor demand shocks for determining the cyclical response of employment and wages using a simple 4-sector calibration of the model.

2.3.1 Skill Heterogeneity

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. Those with high average skill will be strongly attached to the labor force, as the benefit of working will generally outstrip the value of the non-employment outside option. Meanwhile, the ratio of γ_{jk} to $\gamma_{jk'}$ measures the comparative advantage of type j workers in k relative to k'. In this way, the Γ matrix determines the transferability of skills across occupations. Workers with less variance in their skill vector will generally have transferable skills, as the return of working similar across all occupations.

This structure nests the three principal paradigms for skill heterogeneity. If $\gamma_{jk} = \gamma_k$ for all j, then every worker type is equally good at each occupation. This is a standard 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 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 other parameters determining the non-pecuniary benefits of job choice therefore permits a detailed structural estimation of labor substitution patterns.

In effect, 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 manual dexterity, while managers require organizational and negotiation 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 individuals represent bundles of traits which are non-linearly combined in the production of each occupation's tasks (Rosen, 1983).[10] While Γ represents a useful reduced

[10] Edmond and Mongey (2019) explore this idea further in the context of technology adoption, and show that the law of one price for particular skills may fail if workers are unable to unbundle the fundamental talents they have into one task-specific skill level.

form representation of skills that grants great analytical tractability, these caveats confound attempts to decompose Γ into more fundamental components.

2.3.2 Partial Equilibrium: Aggregating Labor Supply Curves

The canonical aggregate labor supply curve traces out the measure of workers willing to be employed as function of the prevailing wage. That is, the aggregate labor supply curve relates movements in aggregate employment $E(\mathbf{w})$ to movements in the aggregate wage $\bar{\omega}(\mathbf{w})$. In this model, the slope and location of this curve will depend on the set of occupational prices used to construct it. A change in the price of routine manual labor may induce a very different aggregate response than a change in the price of engineering for instance. This results from differences in the kinds of workers employed in those two occupations along two dimensions. First is an absolute ability effect: if those who opt to become engineers are high ability (i.e. have especially high $\gamma_{j\text{Engineering}}$), they may be inframarginal to small changes in the price of engineering, and are unlikely to drop out of the employer labor pool when the price of labor falls. The reverse may be true for those employed in low-skill routine occupations such as cashiers. This effect exists in models with vertically differentiated workers, such as the framework of Smith (1995).

Here, there is an additional skill specificity effect: workers are less likely to drop out of the labor force if they may apply their skills to alternative pursuits. For instance, a drop in the price of the services rendered by academic economists may lead to a flow of economists into the private sector to become financial analysts or data scientists. This is possible because the skills of economists are related to those of financial analysts: $\gamma_{j\text{Financier}}$ tends to be high among those employed as academic economists - i.e., those with a high $\gamma_{j\text{Economist}}$. The specificity of the skills of workers employed in the affected occupation will therefore have an influence on the aggregate labor supply curve.

To build intuition for the behavior of the aggregate labor supply curve, consider the following partial equilibrium exercise. Suppose that there are two occupations, and two worker types, each accounting for half of the population. One can trace out an aggregate labor supply curve by relating aggregate employment to aggregate wages as one changes the price of occupational labor services. I do this for three specifications of the Γ matrix:[11]

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

[11] For this exercise, the variance of the idiosyncratic preference shocks ζ_{ik} is 0.25, while the fixed non-pecuniary benefit is set to -1 across both occupations.

The matrix $\Gamma^{(RA)}$ is the representative agent skill matrix: every worker can supply one unit of human capital to each occupation. Meanwhile, $\Gamma^{(AA)}$ is a model with absolute advantage: type 1 workers can supply 1.5 units of human capital to each occupation, while type 2 workers can only supply 0.5 units. Finally, $\Gamma^{(CA)}$ is a model with comparative advantage: both types of workers have the same mean level of labor supply units, but type 1 workers are better at occupation 1, while type 2 workers have a comparative advantage in occupation 2. In all three settings, the aggregate human capital in each occupation is normalized to 1.

Now suppose that price movements were such that $w_1 = w_2 = w$: that is, both occupations had an equal price at all times. This would be the case if occupations were perfect substitutes in firms' production functions – in such a model, a law of one price must hold. By varying the price of labor w, one can trace out an labor supply curve relating aggregate employment (6) to aggregate wages (8). The implied labor supply curves under the three Γ matrices are presented in Panel A of Figure 1. The black line shows the familiar upward-sloping labor supply curve implied by the representative agent skill matrix.

The red line shows the implied labor supply curve under comparative advantage. This curve behaves similarly to the representative agent schedule, only shifted upwards as workers sort into their occupation of skill, thereby realizing higher wages for any given labor price.

In the case where there is absolute advantage (the blue line), the aggregate labor supply curve becomes relatively inelastic at low levels of employment. This is because of a selection 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 higher wages than observed in the representative agent economy. Eventually, nearly all of the high type workers are employed. When this occurs, additional increases in the price of labor w only has an impact on the employment of low type workers. All high type workers are inframarginal to the increases in the wage, but still receive sizable wage increase. As a result, any given increase in the price of labor will generate little increase in employment for a given wage movement, yield a highly inelastic aggregate labor supply curve. Indeed, if the selection is strong enough (e.g. if the variance of the idiosyncratic preference shocks were zero), the model with absolute advantage could generate a backward-bending aggregate labor supply curve if an increase in the price of labor induced a large enough inflow of low-type workers.

The analysis thus far has assumed that the price of both occupations' services in tandem. Now consider the similarly extreme case in which the price of occupation 2, w_2, were fixed at 0.5, while the price of occupation 1 varies to trace out the labor supply curve. This case is depicted in Panel B of Figure 1. I restrict attention to the case with comparative advantage which most easily permits deviations from the law of one price for labor.

Figure 1: Aggregate Labor Supply Curves for Various Labor Price Relations

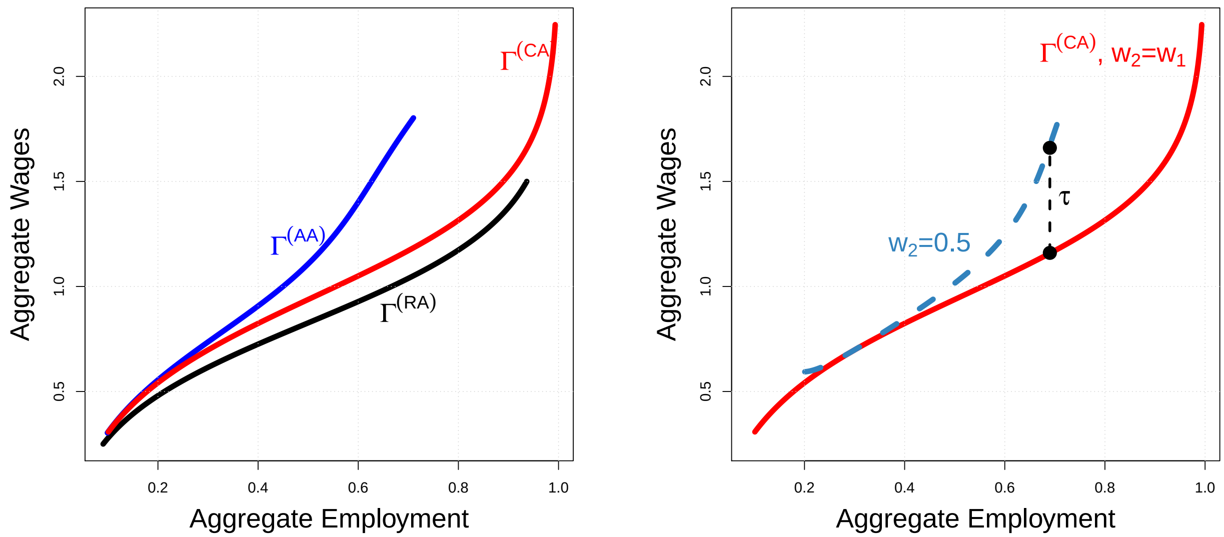

Panel A: $w_1 = w_2 = w$ Panel B: w_2 fixed at 0.5, w_1 varies

Notes: Figure presents aggregate labor supply curves implied by movements in occupational labor prices w_k in a two occupation, two-type model. Panel A plots the implied labor supply curves when the price of labor in occupation 1 is constrained to equal the price in occupation 2, while Panel B plots the implied curves when occupation 2's labor price is fixed at 0.5, while occupation 1's price is allowed to vary between 0.2 and 1.5. The solid black line is the representative agent labor supply curve with $\gamma_{jk} = 1$ for all j and k, while the solid blue line reports the curve when type 1 workers have $\gamma_1 = 1.5$ and type 2 workers have $\gamma_2 = 0.5$, both for all occupations k. The red solid line is the aggregate labor supply curve when $w_1 = w_2$, and Γ exhibits comparative advantage with $\gamma_{jk} = 1.5$ if $k = j$ and 0.5 otherwise. The blue dashed line is the curve when w_2 is fixed to 0.5, and Γ exhibits comparative advantage.

The red line recreates the labor supply curve from panel A under comparative advantage, while the blue dashed line shows the aggregate labor supply curve after fixing w_2 at 0.5. Fixing the wage in occupation 2 makes it appear as though the aggregate labor supply curve has shifted inward and become more inelastic. This is because, in order to induce type 2 workers to enter the labor force, one would require 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. The shifting labor supply curve generates a wedge between the "true" labor supply curve, and that implied by a model which fails to account for dispersion in occupation prices. This wedge is depicted on the figure by τ and may be interpreted in wedge accounting frameworks as a labor wedge (Chari et al., 2007), or in frictionless models as a shock to labor supply (Beraja et al., 2019).

This section makes clear that skill heterogeneity has a profound impact on the behavior of aggregate labor supply. In the long run, changes in the distribution of skills has changed would induce the aggregate labor supply curve to look quite different, even if all shocks are aggregate in nature. Additionally, the inferred aggregate labor supply curve may also move in the short run as a result of occupational price dispersion if workers have comparative advantage in particular tasks. Estimating the distribution of skills in the economy, as well as the medium-run behavior of the price of occupational labor, is therefore a task of paramount importance to which I turn next.

3 Model Estimation

This section describes the procedure used to estimate the labor supply side of the model, including a description of the data used. Next, I outline the approach to calibrating the additional parameters of the model, including the construction of an industry-level TFP series which corrects for unobservable selection in the human capital of employed workers.

3.1 Estimating the Skill Distribution

The identification and estimation of the skill distribution follows closely the distributional framework for employer-employee matched data developed by Bonhomme et al. (2019), who show that a model in which workers sort into firms according to their type is identified given short panel data on job-switchers' employment and wages and a set of key identification assumptions. However wages in the data are not fully determined by the J worker types' mean skill levels. I assume that individual wages in period t are observed with multiplicative measurement error ϵ_{it},[12] 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. Observed wages ω_{it} are then

$$\omega_{it} = \gamma_{j(i)k_t(i)} w_{k_t(i)} \epsilon_{it}.$$

This model of earnings is similar to that of Bonhomme et al. (2019), with two primary differences. First, while Bonhomme et al. (2019) study firm and worker sorting, I study the sorting of workers to occupations, and assign an economic meaning to the wage differences of two workers employed in the same occupation - namely, occupation skill. Second, while Bonhomme et al. (2019) treat the probability that workers switch between each firm type as additional unrestricted parameters to be estimated, I impose a Roy model of occupational choice, so that workers will select into jobs for which they are better suited. This economic model improves the power of my estimation routine by utilizing both wage *and* occupation choice information to estimate the skill vector of each type, rather than just wage information as in Bonhomme et al. (2019).

To fix notation, let m_{it} be an indicator for whether worker i switches occupations between period $t-1$ and t: $m_{it} = \mathbf{1}\{k_t(i) \neq k_{t-1}(i)\}$. Let the history of realization of a random variable Z up to period t be given by $Z^t = \{Z_{i1}, \ldots, Z_{it}\}$. Throughout, following Bonhomme et al. (2019), I maintain Assumption 1 below:

Assumption 1. ***Identification Assumptions***

1. *(Mobility Determinants) - The realization of mobility m_{it+1} and the choice of occupation in period $t+1$, $k_{t+1}(i)$ is independent of the history of measurement error in a worker's wage ϵ_i^t, conditional on the worker's type $j(i)$, and their history of moves and occupation choices $k^t(i), m_i^t$.*

2. *(Serial independence) - The realization of period $t+1$'s measurement error for worker i, ϵ_{it+1} is independent of the history of disturbances ϵ_i^t and occupation choices $k^t(i), m_i^t$, conditional on the worker's current occupation choice $k_{t+1}(i)$, type $j(i)$ and worker mobility decision m_{it+1}.*

3. *(Connecting Cycles) - For any two occupations k and $k' \in \{0, \ldots, K\}$, there exists a connecting cycle $(k_1, \ldots, k_R), (\tilde{k}_1, \ldots, \tilde{k}_R)$ such that $k_1 = k$ and $k_r = k'$ for some r, and such that the scalars $a(1), \ldots, a(J)$ are all distinct where*

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

[12] 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.

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

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 $(\tilde{\omega}_1, \tilde{\omega}_2)$ element*

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

Before unpacking the content of Assumption 1, it is worth noting what this assumption provides. Maintaining this assumption permits the formulation a simple likelihood function to be described below, which can be estimated using two-period panel data. Assumption 1 may be relaxed at the expense of greater data requirements. Unfortunately, the set of long-run panel datasets containing information on occupation and wages is small, requiring the use of panel data with just two periods.

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 may be reformulated to state that the idiosyncratic preference shock ζ_{ikt} is orthogonal to the measurement error ϵ_{it}.[13]

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 may have nearly i.i.d. fluctuations around a mean wage. Similarly, upper executives may have roughly i.i.d. fluctuations in their earnings as a result of random stock performance. However, this assumption will be violated for workers for whom there is strong backloading in wage contracts, or if the discrete type space poorly captures true worker heterogeneity.

The assumption of serial independence of idiosyncratic wages may be relaxed with additional structure and data. Bonhomme et al. (2019) show that first-order Markov processes for wages may be accommodated with four-period panel data. The crux of 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 are a high type individual, or because idiosyncratic wage draws are highly persistent. As a result, I maintain Assumption 1, and attribute all persistently high wages to differences in types, rather than as a result of persistent idiosyncratic shocks.

[13]In essence, this amounts to a timing assumption - although I may have decided to pursue a career in academic economics, I do not know the precise wage draw I will receive at the end of the job market, even if I can anticipate an expected wage for candidates of my type.

Figure 2: Connecting Cycles Illustration

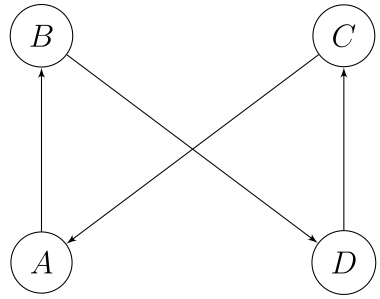

The third item of Assumption 1 requires that any two occupations belong to a connecting cycle for every type of worker. This does *not* require that every worker type must flow between every pair of occupations (k, k') bilaterally. Rather, it imposes graph connectedness in the sense of Abowd et al. (1999). For instance, suppose that there are four occupations in the economy: $K = 4$, as depicted in Figure 2 below. It is not necessary for there to be flows between every pair of occupations, so long as the flows form a cycle as depicted in the figure. This will always hold if the model is true, given the distributional assumptions on the idiosyncratic preference shocks. In addition, it must be that workers of different types flow in different ways – the scalars $a(j)$ must be distinct for each j. This imposes non-random mobility, which will be the case so long as the γ_{jk} differ by worker type.

Finally, the fourth item in Assumption 1 is a standard 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.

Assumption 1 implies that the parameters of the labor supply model are identified and may be estimated by maximizing the likelihood of observing workers' job mobility and wage patterns. Since the identification argument follows that of Bonhomme et al. (2019) almost exactly, I relegate it to Appendix B.

To construct the likelihood of the data, consider the likelihood of observing a single worker i who is observed 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 will include $\gamma_{jk} w_k, \xi_k$ and the parameters governing the idiosyncratic taste shocks ζ_{ikt} and measurement error θ_ϵ. Let $\psi(\omega|k, j, \theta)$ be the density of idiosyncratic wages implied for a type j worker in occupation k. 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, and unemployed workers' wage density has mass 1 and does not affect the likelihood function. 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:

$$\mathcal{L}(\theta) = \sum_i \sum_{k=0}^{K} \sum_{k'=0}^{K} \mathbf{1}\{k_1(i) = k\}\mathbf{1}\{k_2(i) = k'\} \ln l_i(k, k', \omega_{i1}, \omega_{i2}|\theta) \tag{11}$$

In order to maximize this likelihood function, I make the following distributional assumptions and normalizations:

Assumption 2. ***Distributional Assumptions***

1. *The log of measurement error in wages $\ln \epsilon_{it}$ is normally distributed with mean 0 and standard deviation σ_{jk} for a worker of type j in occupation k.*

2. *Idiosyncratic taste shocks ζ_{ikt} are drawn independently over time and across occupations*

3. *The matrix of γ_{jk} is fixed within each estimation window, and normalized to have*

$$\sum_{j=1}^{J} m_j \gamma_{jk} = 1$$

Item 1 of Assumption 2 assumes that wages follow a log-normal distribution which is type-occupation specific, following Bonhomme et al. (2019). Item 2 of the assumption places a restriction on the distribution of taste shocks. The assumption that taste shocks are independent through time is strong, as it generates close to random mobility. Stickiness in occupation choices therefore loads into small variance in ζ_{ikt} and a high within-type variance advantage. To address this concern, Appendix D outlines an approach to relax this assumption by allowing the idiosyncratic preference shocks to be correlated through time.[14] The likelihood function of equation (11) is numerically maximized as described in detail in Appendix D.

Finally, the third item of Assumption 2 normalizes the γ_{jk} to have unit mean within an occupation. This assumption disentangles the variation in mean occupation wages that arises from the price of occupation services w_k and the workers' ability γ_{jk}. I assume that labor prices w_k are fixed within estimation windows.

Intuitively, identification is achieved through occupation switchers. When a worker switches occupations, her type j is fixed across that move. As a result, the distribution of wage changes for workers switching from occupation k to k' informs the parameters of the type-occupation-specific wage distributions, namely $\gamma_{jk}w_k$ and σ_{jk}. In addition, the frequency of moves from occupation k to k' further pin down the relationship between γ_{jk} and $\gamma_{jk'}$. The likelihood that a worker chooses low expected utility jobs is determined by the variance ν of the idiosyncratic taste shocks. Finally, 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, such as engineers, 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 the annualized cost of maintaining engineering knowledge is high.

3.2 Data and Implementation

A key assumption for identification is that every unobserved worker type will form a connecting cycle across occupations. As the number of occupations K increases, this restriction becomes increasingly difficult to satisfy. 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

[14] Specifically, I assume that the joint distribution of taste shocks in period t and $t+1$ is given by applying the Gumbel copula to the marginal distributions of taste shocks in periods t and $t+1$. This loads the stickiness of occupation choices onto one parameter which governs the serial correlation of taste shocks through time. One may then numerically calculate the probability of choosing any pair of occupations (k, k') using properties of the type 1 extreme value distribution.

data from the Current Population Survey (CPS).[15] 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. A sample questionnaire from O*NET is reproduced in Figure A1. 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. For instance, a 2 on the scale for engineering/technology knowledge corresponds to the ability to install a door lock, while a 6 would be chosen by workers who plan for the impact of weather in bridge design or perform similarly complex tasks.

These data have been heavily employed in the existing literature on skill specificity with numerous studies building indices of skill relatedness using the responses to these surveys (Gathmann and Schönberg, 2010; Neffke and Henning, 2013). Within each education tercile, I cluster occupations into five groups according to their required level of knowledge and skills using a k-means algorithm. Specifically, let the number of SOC occupations be given by C, and index each SOC code by c. Suppose there are S distinct skills, indexed by s, and let the level of skill m required by occupation c be given by $h_{c,s}$. The goal is to define a set of K clusters, with required skill vector $\mathbf{H}_k$, and a mapping $k(c)$ assigning each SOC occupation c to a cluster k, so as to minimize the total distance between the SOC occupations' skill vectors, and the skill vector of their clustered occupation. Mathematically, this amounts to solving, within each tercile,

$$\min_{k(1),\ldots,k(C),\mathbf{H}_1,\ldots,\mathbf{H}_K} \sum_{c=1}^{C} \left(\sum_{s=1}^{S} \left[h_{c,s} - H_{k(c),s} \right]^2 \right)^{\frac{1}{2}} \tag{12}$$

A brief overview of the clustered occupations is provided in Table 1, with a fuller picture provided in Appendix C. Clusters are ordered according to their mean annual income in the period 2002-2006, as implied by data from the Bureau of Labor Statistics' Occupational Employment Statistics (OES). The occupation clustering is intuitive, with similar occupations being paired into the same cluster. Within each cluster, there remains a variety of occupations. For instance, cluster 12 pairs nurses together with surgeons. It is natural that

[15] 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. More data processing details are provided in Appendix C.

Table 1: Summary of k-means clustered occupations

#	Broad Category	Sample Occupations
1	Routine	Cashiers, Stock Clerks, Maids and Housekeeping Cleaners, Truck Drivers
2	Low-Skill Service	Waiters and Waitresses, Receptionists, Hairdressers, Counter Clerks
3	Manual Laborers	Painting Workers, Stock and Material Movers, Helpers-Production Workers
4	Salespeople	Retail Salespeople, Bartenders, Hotel Desk Clerks
5	Production	Machinists, Operating Engineers, Welders
6	Clerical	Secretaries, Office Clerks, Tellers, Bookkeepers
7	Construction	First-Line Supervisors of Construction Trades, Construction Laborers
8	Tradespeople	Carpenters, Plumbers, HVAC workers, Mechanics
9	Supervisors	First-Line Supervisors of Sales Workers/Food Prep Workers/Mechanics
10	Technicians	Electricians, Engineering Technicians, Telecom Line Installers
11	Social Skilled	Teachers, Lawyers, HR Workers
12	Medical	Registered Nurses, Physicians, Surgeons, Pharmacists, Counselors
13	Computing	Computer Support Specialists, Software Developers, Database Administrators
14	Engineers	Mechanical Engineers, Electrical Engineers, Architects
15	Business Services	Accountaints, General Managers, Financial Analysts

Notes: Table reports examples of occupations within each occupation cluster. Clusters are ordered according to their mean wages in the OES data in 2013. Broad categories are labels provided by the author. Occupation clustering proceeds in two steps: first occupations are grouped into terciles of educational attainment, measured by share with at least some college, then clustered according to a k-means clustering algorithm within each tercile using the Skill and Knowledge vectors implied by O*NET data.

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} - the worker types with high γ_{jk} for medical occupations may be thought of as the surgeons, while those with lower γ_{jk} may be the nurses.

With the occupation clusters in hand, I turn to the estimation of the Γ matrix. I use the March Supplement of the Current Population Survey going back to 1984, focusing on workers, both male and female, aged between 21 and 60 years old. The CPS is a rotating panel survey conducted by the BLS in cooperation with the Census Bureau designed to be representative of the US population. Households in the CPS are surveyed for four consecutive months, before an eight month hiatus, and a subsequent additional four month survey. Each month, it asks respondents about their employment status, including the occupation and industry in which they are employed. In addition, every March, the Annual Social and Economic Supplement (ASEC) is administered, which asks numerous additional questions regarding workers' annual income and hours worked. Given the rotating panel structure of the CPS, workers included in the ASEC will appear for two consecutive years.[16] My measure of worker

[16]Linking the ASEC to the basic CPS files is not a trivial task. I follow the IPUMS methodology of Flood

earnings ω_{it} is 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 multiple years together in order to minimize sampling noise. Specifically, I estimate the model for the period immediately before the Great Recession (2002-2006) and before the recession of 1990-91 (1984-1989).

3.3 Skill Estimation: Discussion

The estimation framework employed here has the large benefit of providing cardinal measures of skill transferability. Rather than relying entirely on potentially noisy survey answers about the importance of skills a particular occupation, this framework assumes that skill affects economically meaningful objects: the price and quantity of labor. This permits robust counterfactual analyses which have hitherto been rare.

The framework has the additional benefit of being estimable using publicly-available short panel data, such as the CPS. Such datasets have existed for long periods in many developed countries. As a result, this framework is portable to multiple settings and multiple time periods. Indeed, it may be applied to study firm- or industry-specific human capital, so long as Assumption 1 is satisfied.

However, it is not without its limitations. By assuming a Roy model of occupation choice, the framework abstracts from meaningful changes in the bundles of tasks that occupations employ. Instead, the matrix Γ must be thought of as a reduced form representation of the skills needed for each occupation. The model therefore cannot tell us whether the Γ matrix changes due to changes in the skills of workers or from changes in the required task content employed in each occupation cluster.

In addition, the requirement of connecting cycles imposes that the number of worker types J and occupations K may not grow too large lest the empirical probability of a type j worker moving between any two occupations shrink to 0. This necessitates the clustering of occupations described above. The estimated Γ matrix will naturally be sensitive to the choice of cluster, and the exogenously-imposed number of worker types J. What's more, clustering assumes that skills are perfectly transferable within cluster. In reality, the degree

and Pacas (2008) to generate consistent panel identifiers in the March supplement. This approach is detailed in Appendix C.

[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 workers' human capital.

of specificity of skills within cluster may have changed over time as well. If within-cluster skills have gotten more specific, then the trends presented below will understate the degree to which skills have become more specific in the economy.

Finally, the framework presented here is fundamentally static in nature. Workers do not make irreversible investments in specific human capital, nor is their occupation choice forward-looking. This is done for tractability. Were there irreversibility in workers' occupation choices, workers would need to know the process underlying the labor demand of each occupation, as well as the existing mass of each type of worker in each occupation in order to forecast the path of their value in each occupation. This renders estimation infeasible, as the dimensionality of the state space rises quickly. The extension outlined in Appendix D addresses this concern by loading the forward-looking nature of occupation choices onto the process of idiosyncratic preference shocks ζ, which maintains the static optimization problem of equation equation (2) while standing in for explicit costs of switching occupations.

The lack of investment in human capital implies that this framework should not be used to estimate long-run responses to structural shifts in the economy. Rather, it is suited for studying the impact of a *fixed* skill distribution on the economy's responsiveness to short-run shocks. This is appropriate in the application of this paper, but would be inappropriate for studies seeking to understand how the long-run decline in the labor share affects workers' reallocation across occupations in the last 40 years, for instance. Developing frameworks to estimate a dynamic skill distribution is a fertile area for future research.

3.4 Externally Calibrated Parameters

Table 2 summarizes the model's calibration. The parameters governing labor supply – the distribution of skills and types Γ, m_j, as well as the variance of the idiosyncratic wage draws σ_{jk}, fixed non-pecuniary benefits of each occupation ξ_k, and the variance of the idiosyncratic preference shocks ν – are estimated using the maximum likelihood approach outlined above.

There remain multiple parameters to input to the model. First, I choose the number of industries S to match the number of 3-digit NAICS industries. I assume that the number of worker types J is 8, and that the elasticity of substitution η between intermediate industries in the production of the final good is 4, following Broda and Weinstein (2006).[18]

[18] Broda and Weinstein (2006) estimate the mean elasticity of substitution across 3-digit SITC products, rather than industries. The true elasticity of substitution across 3-digit industries may therefore be somewhat lower than 4. Reducing the elasticity of substitution across industries would have the effect of reducing the dispersion of labor demand shocks for each occupation, as a shock to a particular industry would be partially capitalized into the price of that industry's output. As argued above, this would increase the importance of absolute advantage for employment elasticities, but has little qualitative effect on the model's ability to match the countercyclical wage growth of the 2009 recession.

Table 2: 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
ρ	Serial Correlation of T1EV shocks	Maximum Likelihood
External Calibration		
S	Number of Industries	57 (# 3-Digit NAICS)
J	Number of types	8
K	Number of occupations	15
η	Elast. of Subs. Between Industries	4
$F^{(s)}(l_{s1}, \ldots, l_{sK})$	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 Industry s	BEA Labor Share
α_{sk}	Share of Occupation k in Industry s	OES Share in Wage Bill
z_{st}	TFP series for industry s	Adjusted VA/Worker

I 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 industry s, x_s, will be equal to labor's share of value added in industry s, while α_{sk} will be the share of industry 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 industry, 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 assumed to be fixed to the average share in each industry over the period 2002-2006.

[19]If one were to instead impose a CES production function, one would need to estimate the elasticity of substitution across occupations at the industry level, which is outside of the scope of this paper. A CES production function could increase or decrease cross-industry 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 such as skilled engineers.

3.4.1 Estimating Industry-Level TFP Series

The traditional method for calculating industry TFP in a model with Cobb-Douglas production is to note that

$$\ln z_{st} = \ln \text{Value Added}_{st} - x_s \ln(\text{Labor Input}) - (1 - x_s) \ln(\text{Non-Labor Input}).$$

Therefore, given data on value added, the labor share of production, and production inputs, one may calculate an industry's TFP. A challenge arises when there is selection on unobservable quality in labor inputs. A standard approach to remedy this is to use the total wage bill of each industry under the assumption that highly-skilled workers are remunerated according to their human capital. However, the wage bill reflects both the quality of workers and the price of labor. Increases in TFP increase labor demand, which in turn increases the price of labor and the wage bill, inducing an endogeneity problem to the traditional estimation.

Through the lens of my model, one may think of the problem as arising because $\bar{\gamma}_{kt}$ fluctuates over the cycle. Specifically, let E^s_{kt} denote the number of workers employed in occupation k in industry s. Because workers are indifferent over industries conditional on their occupation, the total labor units employed in occupation k industry n are

$$l_{skt} = \bar{\gamma}_{kt} E^s_{kt}.$$

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

$$\ln z_{st} = \ln \text{Value Added}_{st} - x_s \sum_{k=1}^{K} \alpha_{sk} \ln(\bar{\gamma}_{kt} E^s_{kt}) - (1 - x_s) \ln(\text{Non-Labor Input}). \tag{13}$$

To calculate $\bar{\gamma}_{kt}$, I estimate the labor supply parameters – $\Gamma, \xi_k, m_j, \sigma_{jk}, \nu$, 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. Running these parameters through the Roy model of equation (2) yields an estimate of the mean human capital of workers employed in every occupation in every year.

With the estimated $\bar{\gamma}_{kt}$ in hand, I then estimate industry-level TFP series adjusted for selection on unobservable human capital. To do so, I employ data from the BLS' KLEMS Multifactor Productivity Series to calculate the value added and non-labor inputs in each industry every year. To calculate the employment of each occupation in each industry, I combine data from the Quarterly Census of Employment and Wages (QCEW) with data from the CPS. The QCEW provides the total employment and wages by industry and locale using administrative data derived from tax records. Using the CPS, I calculate the share of employment in each 3-digit NAICS industry 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.[20] Finally, I use equation (13) to estimate industry TFP series.

This adjustment is meaningful. Table 3 describes the annual percentage changes in implied total factor productivity for the largest sectors in the 1990-91 and 2008-2009 recession. The table excludes the 15 sectors which were among the 20 sectors with the smallest value added in both 1990 and 2008. Whereas the BLS series shows no drop in productivity in the Construction sector in 2009, despite large layoffs and declines in value added, the series adjusted for human capital selection shows a 6 percentage point decline. The same is true for miscellaneous manufacturing sectors, which saw a productivity increase of 2.4% in the BLS series, but a 4.3% decline after adjusting for worker composition. In some sectors, however, the adjustment has little bite. For example, in the hospital and residential care facilities sector, both series show a 1.3% increase in productivity from 2008 to 2009. The fact that selection is unimportant in this sector is intuitive given the specialized nature of medical care. Aggregating sectoral TFP series according to their 2008 shares of aggregate value added, the adjusted TFP series shows a decline in aggregate productivity of 5.9%, compared to a 4.2% decline in the unadjusted BLS series.

4 Estimation Results

This section describes the results of the maximum likelihood estimation. I first detail the estimated skill distribution for the periods before the 1991 and 2008 recessions, with a particular focus on a few key moments of the Γ matrix which represent workers' absolute and comparative advantage. I then use the estimated skill distributions to explore how occupation-specific labor supply elasticities have evolved. Finally I highlight changes in the cyclicality of selection on human capital. In Appendix A, I report estimation validation exercises, including the model's ability to match employment shares and earnings distributions, as well as the relationship of the estimated with skill distribution with some reduced form skill relatedness measures based on O*NET survey responses.

4.1 Estimated Skill Distributions

Table 4 reports the estimated matrix Γ, along with the mass of each type of worker m_j fir tge period 1984-1989. Each column reports the γ_{jk} vector for a given worker type j, while

[20] The BLS' OES data provide data on the occupation wage bill for each industry. However, it only provides annual information going back to 1997, and thus cannot be used to study the period in which wages were highly cyclical.

Table 3: 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	Industry 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: Data processing and other information services includes NAICS codes 516, 218, and 519. Aggregate TFP constructed as the mean of industry TFP series, weighted by value-added in each industry. 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 13.

Table 4: 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	–
$\frac{max_k(\gamma_{jk})}{\min_k(\gamma_{jk})}$	20.2	2.68	9.68	19.7	2.06	11.3	1.23	13.8	–

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 four 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, and the ratio of the type's skill in her best occupation relative to her worst occupation. Estimation procedure laid out in Section 3, and carried out using data from 1984-1989 in the CPS.

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 rows report the variance and geometric range of each column vector. The corresponding table for the 2002-2006 table is reported in Appendix A.

The table shows, for instance, that a type 1 worker supplies 0.855 units of human capital to routine occupations (cashiers, security guards etc.), but only 0.06 units of human capital to skilled business services occupations (such as financial analysts or management consultants). In contrast, type 8 workers supply 4.95 units of human capital to business services occupations, but only 0.48 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.

It is useful to consider a subset of meaningful moments of the estimate human capital distribution. A natural measure of a worker's absolute advantage is the mean level of human capital of each worker type $\mathbb{E}_k[\gamma_{jk}]$. In the period before the 1991 recession, the best 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. In recent periods, 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 total variance of skills in the economy indicates the deviation from a representative agent framework. In the late 1980s, the standard deviation of skills, weighted by the mass of types, was 0.77, while in the mid-2000s, this standard deviation had increased to 0.86. Given the mean of the Γ matrix is normalized to 1 within each occupation, this may be interpreted as the standard percentage deviation from mean workers in mean occupations. That variance of skills has increased 25% over the course of this 20 year period indicates that the quality of the representative agent approximation of skills has declined, and reflects increases in both within and across occupation variance in earnings.

The variance in skills may be decomposed into a within-type and an across-type variance. The across-type variance is informative about the difference in level of skill for various workers. If this variance is high, then some workers have a substantially higher mean level of skill than other workers. Meanwhile, the within-type variance informs us about the gains to workers of allocating themselves to their best occupations. If the within-type variance is high, there is great dispersion in workers' skills across occupations. Mathematically, we may consider the between and within variance as

$$Var^{BTWN} := \sum_{j=1}^{J} m_j(\mathbb{E}_k[\gamma_{jk}] - 1)^2; \qquad Var^{WTHN} := \sum_{j=1}^{J} m_j Var_k(\gamma_{jk}), \tag{14}$$

respectively, where we use that the weighted mean γ_{jk} is equal to 1.

Figure 3 plots the within and between variance of skills in the economy prior to the 1991 and 2008 recessions. 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. The figure shows that the cross-type variance of γ_{jk}s has increased from 0.50 to 0.56, an increase of 10.4% in the 20 years leading up to the Great Recession. There is an even larger increase in within-type variance, while the mean variance of the γ_{jk} vectors was 0.15 before the 1991 recession, it

Figure 3: Absolute and Comparative Advantage: 1984-1989 and 2002-2006

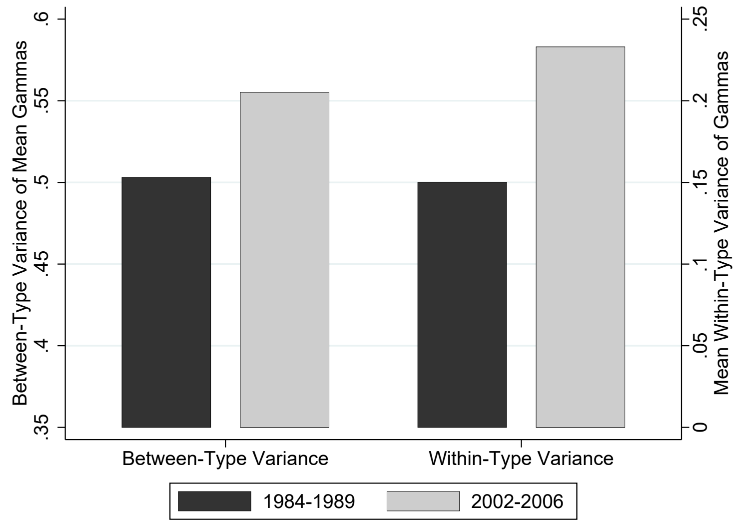

Notes: Figure plots the estimated within and between type variance of skills in the economy, captured by the Γ matrix of Table 4 and A2. Estimation follows the procedure outlined in Section 3, and carried out separately in the CPS March Supplement for the periods 1984-89 (gray bars) and 2002-2006 (black bars). Within and between variance defined as in equation 14.

was 0.23 prior to the 2008 recession, an increase of 55.2%. This suggests that skills have become more specific over time and that the gap between the best and workers has grown.

However, 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.

Heuristically, this result arises from two moments in the data. The increase in within-type variance owes to an increase in the variance of wage changes on occupation switches. As the between-occupation variance increase, the more one infers that individual workers' skills are better tailored to particular applications. Meanwhile, the increase in cross-type variance arises 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.

The degree to which skills are transferable across shocked sectors will similarly affect aggregate wage dynamics by dictating the size of labor supply spillovers as workers reallocate from declining occupations to growing occupations.[21] The degree of skill transferability

[21] In Appendix F, I provide reduced form evidence that isolated labor demand shocks generate labor supply

between any two occupation may be captured as the correlation of the row vectors of the Γ matrix. 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, such as stock and material movers, may easily transfer their skills to production occupations to be serviceable welders or machinists.

Figure 4 plots a correlogram of 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. Indeed, as a validation check, Appendix A 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.

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 does 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

spillovers in sectors with related skills. Following the rapid decline of the mining sector from 2014-2016, tradable goods sectors which employ skills related to mining saw increased employment and reduced wages relative to sectors which employ skills unrelated to mining, suggesting the existence of such labor supply spillovers. Horton and Tambe (2019) further presents a case study in which workers with skills in Adobe Flash quickly transitioned to related tasks upon the announcement that Apple would no longer support Flash for its applications.

Figure 4: Correlation of Occupation Skills, 1984-1989 and 2002-2006

Panel A: 1984-1989

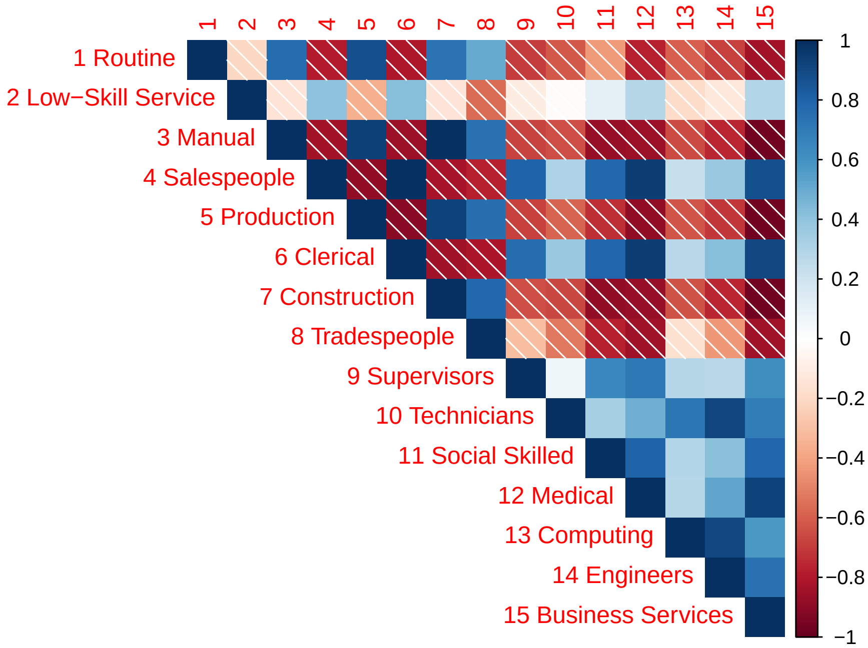

Panel B: 2002-2006

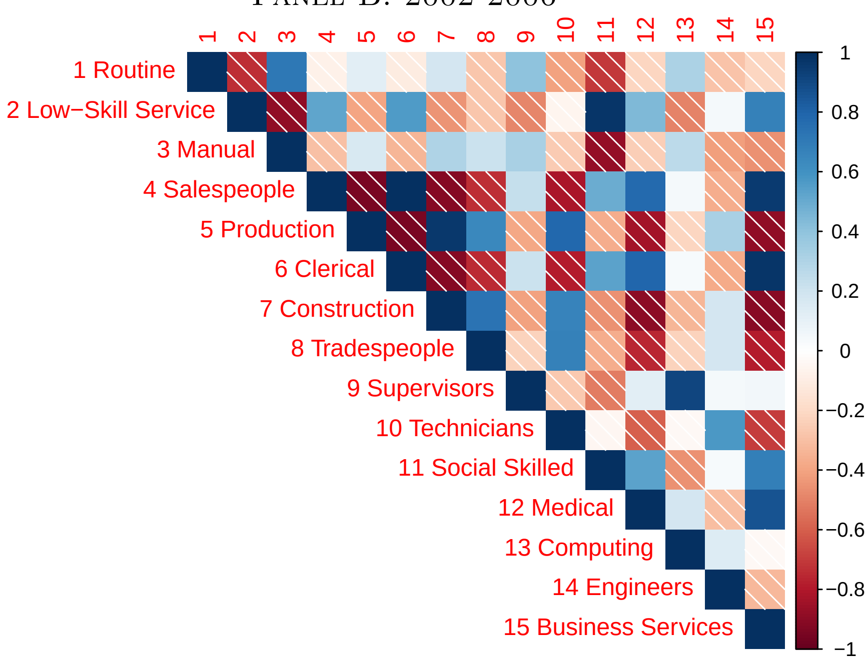

Notes: Figure plots the correlation of the row vectors of the estimated Γ, normalized by workers' mean skill in each occupation. Estimation follows procedure outlined in Section 3, and carried out separately in the CPS March Supplement for the periods 1984-89 (Panel A) and 2002-2006 (Panel B). Blue squares indicate that the correlation of skills between occupations is positive, while red checked squares indicate a negative correlation. Deeper colors indicate that the magnitude of the correlation is closer to 1.

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 period leading up to the 1991 recession, 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. 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 construction workers were displaced by declines in construction demand, they 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.

In total, the results presented in this section suggest that the labor market has moved further away from a representative agent framework in which all workers have interchangeable skills. Absolute advantage has increased, suggesting that the gap between the best workers and the least skilled workers in the economy has risen. Comparative advantage has similarly risen, which implies that workers have become more specialized over the last twenty years. Finally, the transferability of skills has generally declined, both 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 a fertile ground for future research. Section 2.3.2 suggests that the increasing specificity of skills has likely destablized the aggregate labor supply curve.

4.2 Shifting Occupational Labor Supply Elasticities

The deviation from representative agent skill distributions has implications for estimated labor supply elasticities. To illustrate this, consider the effect of unilateral increases in the price of each occupation w_k. Increasing these prices will induce flows out of non-employment, implying a labor supply elasticity of non-employment to the price of each occupation. Figure 5 plots these implied elasticities for each occupation. The gray bars plot the elasticities for

Figure 5: Estimated Labor Supply Elasticities for Each Occupation, 1984-1989 and 2002-2006

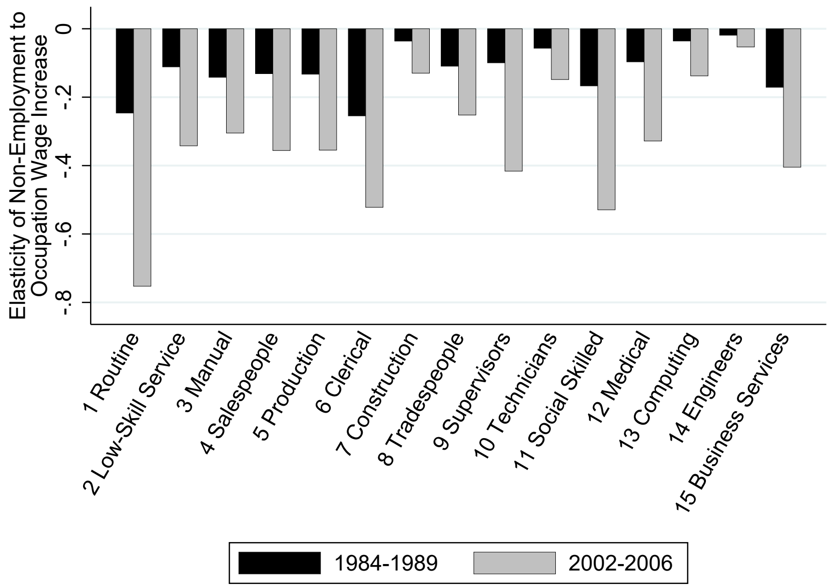

Notes: Figure reports the estimated model-implied elasticity of non-employment to a change in the price of each occupation's price of labor w_k. Estimation procedure outlined in Section 3, and carried out separately in the CPS March Supplement for the periods 1984-1989 (black bars) and 2002-2006 (gray bars). Elasticity calculated by calculating the percentage change in non-employment rates in response to a unilateral 1% change in the price of labor in each occupation.

the 1984-1989 period, while the black bars plot the elasticities for the 2002-2006 period.

The figure shows substantial variation in the elasticity of non-employment to changes in occupation prices. There is no single "aggregate labor supply curve." Rather, the aggregate labor supply curve arises by aggregating movements along each of these primitive occupation-specific labor supply curves. 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 labor supply curves. In many models with a representative agent, this will look as though workers are subject to labor supply shocks.

Figure 5 shows some systematic patterns to labor supply elasticities. In both periods, the occupation cluster with the highest non-employment elasticity is the set of routine occupations. Low-wage occupations generally have higher non-employment elasticities than do high wage occupations, such as engineering. This is intuitive, and results from the fact that the workers most on the margin of non-employment are the type $j \in \{1, 2\}$ workers, who are highly sensitive to fluctuations in the price of routine and other lower-skill occupations.

The figure additionally 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 absolute value in 2002-2006. This change in the elasticity of labor supply primarily results from three forces. First, the degree of absolute advantage has risen, implying that more low skill workers are marginal to small changes in the price of labor. Second, the transferability of skills has fallen due to the rise in comparative advantage. As a result, a decline in the price of a particular occupation does not allow workers to move across occupations as easily as it once did, particularly amongst low skill workers. Finally, the standard deviation of idiosyncratic preference shocks ν is estimated to have declined from 0.60 to 0.29 so that workers have become more responsive to changes in expected utility when making occupation choices. This is another, perhaps more primitive, reason for shifts in the aggregate labor supply elasticity.

4.3 Human Capital Selection in the Employed Pool of Workers

The presence of multiple worker skill types opens up the possibility for the selection of workers to vary over the cycle. Figure 6 plots the time series of estimated mean human capital level of employed workers $\bar{\gamma}_{kt}$ for each of the 15 occupation clusters, as well as the aggregate mean human capital level of employed workers. To calculate these mean human capital levels, I re-estimate the maximum likelihood function in every two-year period of the CPS, and then estimate the choice probabilities $\mathbb{P}_{kt}(j)$ for each worker type and occupation according to equation (3).

The figure shows that the cyclical patterns of selection have changed for many occupations. For example, although it has always been the case that the selection of production workers, construction workers, and tradespeople improves in recessions, this was especially strong during the Great Recession. Whereas the mean human capital of production workers increased by 14% in the 1990 recession, the efficiency of production workers improved by 50% between 2008 and 2009. Similarly the selection of employed construction, tradespeople, and engineers were relatively flat during the 1990 recession, but increased by 42%, 40%, and 22%, respectively during the 2009 recession. However, some occupations, such as medical occupations, exhibit little cyclical selection patterns.

In aggregate, the mean human capital of employed workers rose by 10% from 2008-2009, but only 4% in 1991. Given that wage growth during the Great Recession was approximately 2% and wages declined by about 2.3% in the 1990 recession, this change in the cyclicality of selection on human capital can account for greater than 100% of the change in wage cyclicality. The model implies that, absent this selection force, real wages would have fallen in 2008-09 by more than they did in 1990-91.

To summarize, the data show that the variance of earnings both within and between

Figure 6: Time Series of Estimated Mean Human Capital of Employed Workers $\bar{\gamma}_{kt}$

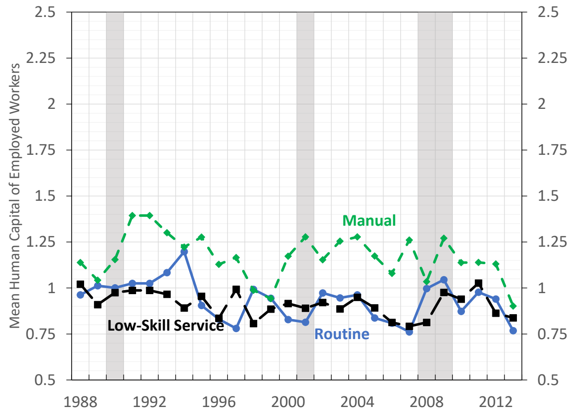

Panel A: Occupations 1 to 3

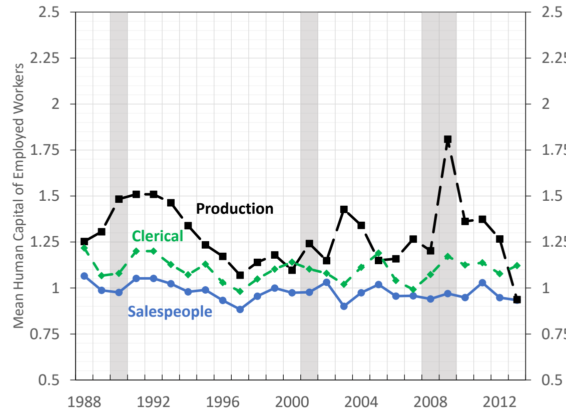

Panel B: Occupations 4 to 6

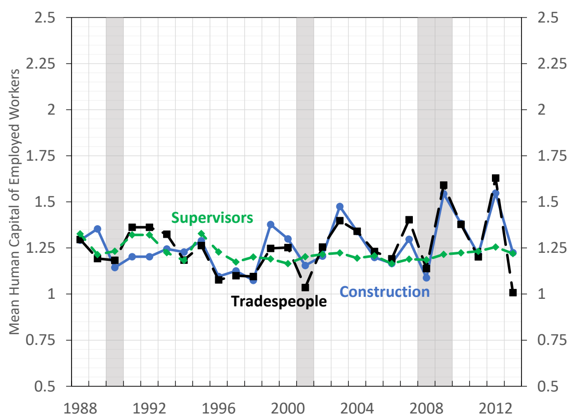

Panel C: Occupations 7 to 9

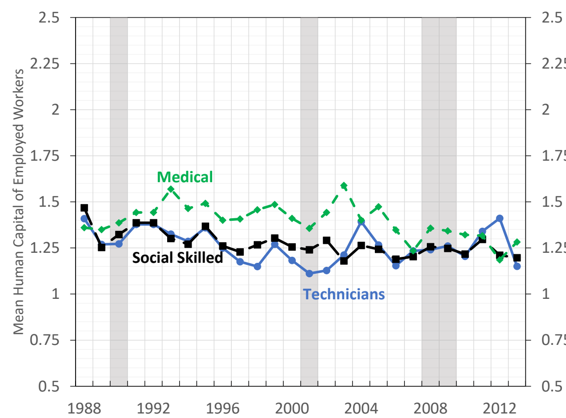

Panel D: Occupations 10 to 12

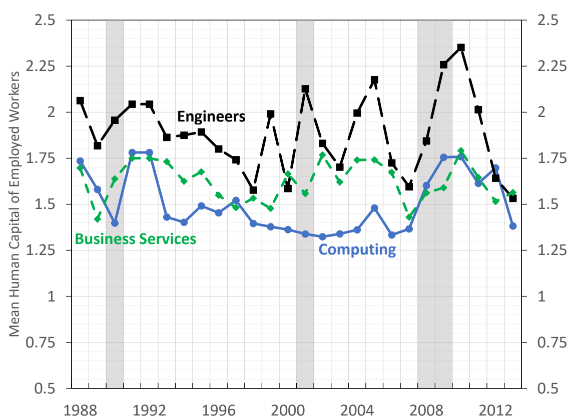

Panel E: Occupations 13 to 15

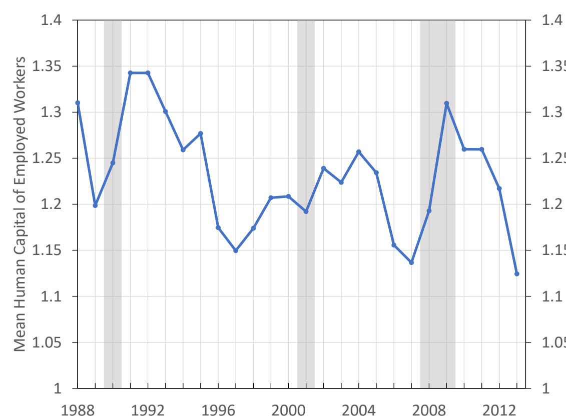

Panel F: Aggregate

Notes: Figure plots the time series of the estimated mean human capital level of employed workers in each of the 15 occupation categories (Panels A-E) and in the aggregate economy (Panel F). Estimation is based on the approach detailed in Section 3 using 2-year rolling panels in the CPS.

occupations has grown, which is interpreted through by the model as an increase in between- and within-type variance, respectively. The increased variance has been coupled with a decline in skill transferability. As a result, when labor demand shocks arise, workers are less able to reallocate themselves to other occupations, reducing the extent to which isolated shocks exert downward wage pressure elsewhere in the economy. This led to a greater elasticity of non-employment to occupation-specific price shocks. Due to the increase in skill variance, selection on unobservable human capital levels in the employed pool of workers has become more countercyclical: now more than ever, the low-skill workers leave employment in a downturn. This selection on unobservables can account for more than all of the change in wage cyclicality. Estimating whether the change in aggregate employment and wage dynamics is principally due to changes in the composition of industry shocks or underlying labor supply, is the focus of the next section of the paper.

5 Model-Based Quantitative Exercises

I shock the model with the selection-adjusted industry TFP series constructed in section 3.4.1. For the period around the 1990 recession, I calibrate the skill distribution to match the 1984-1989 estimation, while around the 2008 recession, I use the estimated skill distribution from the 2002-2006 period. The remainder of the parameters are calibrated as described in section 3.4 and held fixed through the period.

The model is able to replicate the increase in average wages in 2009, followed by a decline in average wages in the recovery, as well as a steep drop in employment. Figure 7 plots the aggregate labor market dynamics implied by the model during the 1990 and 2008 recession. The figure plots the level of mean average earnings of employed workers and the measure of workers employed, relative to the pre-recession peak, that is, relative to 1990 and 2008, respectively. The blue solid line plots the evolution for the model calibrated to the 2008-09 recession, while the green dashed line plots it for the 1990-91 recession. The model implies a wage increase of 2.2% and employment decline of 10%; in the data these numbers are 1.6% and 5%, respectively.[22] The implied elasticity of employment to wages is 4.5, while in the data it was 3.1, In addition, the model shows a decline in wages in the 1990 recession with a more muted decline in employment. The model implies a wage decline of 6%, compared with 2.2% in the data, and an employment decline of 0.4%, compared with 1.4% in the data. The model-implied elasticity is 0.07, compared with 0.64 in the data.

As highlighted throughout the paper, this change could be due to either a change in the

[22]The data numbers consider year over year changes from March 2008 to March 2009, and reflect employment changes, rather than aggregate hours changes.

Figure 7: The Effect of the Changing Skill Distribution and Nature of Demand Shocks on the Wage Cyclicality of 2008-9

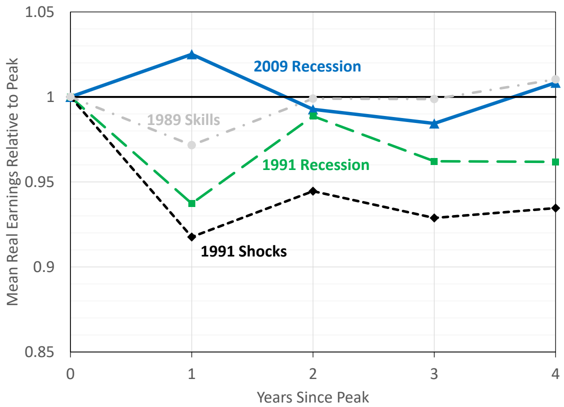

Panel A: Wages

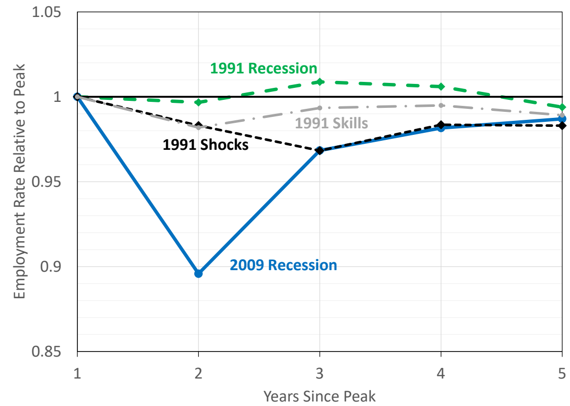

Panel B: Employment

Notes: Figure plots the model-implied aggregate behavior of wages (panel A) and employment (panel B) in response to recessionary shocks. The blue solid line plots the behavior for the 2008-09 recession, while the green long-dashed line plots the behavior for the 1990-91 recession. The gray dash-dotted line reports the response of an economy with the labor supply parameters estimated on the 1984-89 data to the set of 2008-09 industry shocks. The black short-dashed line reports the response of an economy with the labor supply parameters estimated on the 2002-06 data to the set of 1990-91 industry shocks.

distribution of skills or the selection of demand shocks that hit the economy. To decompose these forces, I consider the following exercises. First, I hold fixed the distribution of skills to the level before the Great Recession, and study the movements of aggregate employment and wages in response to the sequence of technology shocks observed during the 1990-91 recession. This studies the effect of the change in the structure of demand shocks between 1990 and 2008 for determining the movements in 2008. Second, I suppose that the shocks hitting the economy were as they were in the Great Recession, but the skill distribution were as it was prior to the 1990-91 recession. This studies the effect of the changing skill distribution on the changing cyclicality of employment and wages.

The black dashed line in Figure 7 shows the implied labor market response when the 2002-2006 skill distribution is shocked by the 1990-91 TFP shocks, while the dash-dotted gray line shows the economy's response under the 1984-1989 skill distribution. In both cases, the recession would have seen declining wages. The countercyclical wages during the 2008 recession were therefore a result of a confluence of unique factors. Without the shift in the skill distribution, the model shows declining wages, and a smaller employment response. With the 1991 skill distribution, the model implies that wages would have declined 3%, with a 2% employment decline. This is for two reasons. First, the higher transferability of human capital in the 1990 calibration implies that low type workers were more able to move to other occupations, rather than exiting the workforce. What's more, the gap between low and high skill workers fell, muting the strength of the pro-cyclical selection effect. Second, the increased transferability exerted larger labor supply spillovers from declining occupations to other occupations, increasing the downward wage pressure throughout the economy.

Although the change in the skill distribution is important, changing the composition of shocks that hit the economy during 2008 generates a larger decline in aggregate wages. Were the shocks in 2009 the same as those in the 1990 recession, the model implies that aggregate wages would have fallen by approximately 8%, with employment declines of 2%. The 2008 shocks hit all sectors to which low-skill workers tend to sell their labor. As a result, labor demand for such workers fell dramatically. This led to a great deal of selection in the employed pool of worker, as only the most skilled workers remained employed. Because all of the industries receiving negative shocks employed the same particular low-skill workers, the price of occupation services for those workers declined, pushing them out of the labor force. Meanwhile, the wage of certain high skill workers, particularly medical, software, and skilled business service occupations, rose. The divergence between occupation labor prices led to a shift in the aggregate labor supply curve, as suggested by Panel B of Figure 1.

The heterogeneity of skills is central to this result. In a representative agent framework in which Γ is a matrix of ones, the model implies wage declines on the order of 4% in the 2007-09 recession, while employment declines by just 1%, implying an aggregate employment elasticity of 0.27.[23] This is substantially lower than the 1-2 range used in much of the macro literature, but more in line with some of the micro estimates of labor supply elasticities.[24] This is, in essence, a restatement of the Shimer (2005) puzzle - under reasonable calibrations, most models of the labor market lead to large wage movements and only muted movements in employment. Aggregate elasticities must account for the selection in the workers employed over the cycle, which mutes wage movements and spurs larger aggregate elasticities. In my model, when this selection effect is shut down by assuming workers are identical, the aggregate elasticity of labor supply returns to more plausible micro labor supply elasticities.

This section shows that a reasonably-calibrated model with workers of heterogeneous skill types, firms employing heterogeneous task content, and imperfect transferability of skills is able to replicate the shifts in aggregate employment and wage dynamics, to the extent that the relationship between wage and employment growth turned negative during the Great Recession. Central to this result is the fact that the sectors which observed

[23]For this exercise, the standard deviation of idiosyncratic preference shocks ν is fixed at its estimated level for 2002-2006: 0.29.

[24]See Saez et al. (2012) for a survey of these elasticities, and Chetty (2012) for an alternative approach to reconciling micro and macro labor supply elasticities.

large downturns all employed similar skills, limiting low-skill workers' ability to reallocate to other occupations. This both limited the downward wage pressure that these sectors exerted throughout the economy, and induced large positive selection in the pool of employed workers.

6 Model-Implied Selection Corrections

Economists have long recognized that the composition of workers employed varies over the business cycle. This has prompted a number of attempts to correct aggregate wage series to account for these changing worker composition (Daly and Hobijn, 2014). For example, Solon et al. (1994) assume the following statistical model

$$\ln \omega_{it} = \alpha_i + \beta_1(U_t - \delta_1 - \delta_2 \cdot t - \delta_3 \cdot t^2) + \beta_2 \cdot t + \beta_3 \cdot t^2 + \beta_4 X_{it} + \beta_5 X_{it}^2 + \epsilon_{it} \tag{15}$$

where U_t is the contemporaneous aggregate unemployment rate, 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 β_1 will be biased upward in aggregate data. By estimating equation (15) in first differences, one can hold fixed characteristics of a worker which are fixed over time, such as the workers' education, race, sex, and fixed unobserved ability. Therefore estimating

$$\Delta \ln \omega_{it} = \eta_0 + \beta_1 \Delta U_t + \eta_1 \cdot t + \eta_2 \cdot X_{it} + \nu_{it}, \tag{16}$$

where ΔZ_t represents the change in a variable Z between $t-1$ and t, yields a consistent estimate of the true cyclicality of wages β_1.

However, equation (15) implicitly assumes that workers are vertically differentiated: some workers are high type, while others are low type. My framework suggests an alternative challenge for these methods of composition adjustment. Even if the same set of workers remain employed throughout the cycle, they may be reallocated to tasks in which they have different human capital levels. This induces fluctuations in workers' wages that reflect changes in the allocation of employed workers to tasks rather than in the price of labor.

This bias could either inflate or deflate the measured cyclicality of wages. If workers move to tasks to which they are less well-suited during downturns, then the Solon-Barsky-Parker (SBP) correction would overstate the cyclicality of wages that arises purely from price effects. For instance, if middle-class manufacturing employees become janitors in a recession, they may see large reductions in wages even if the price of manufacturing or janitorial labor does not fall substantially. The reverse would be true if workers are employed in tasks to which they are poorly suited in booms. For example, if workers with little mining skill began work in North Dakota's oil industry during its oil boom, they would see smaller increases in wages than would be implied by the pure price of oil extraction labor.

To address this reallocation concern, recall that the aggregate wage may be written as

$$\bar{\omega}_t = \sum_{k=1}^{K} \left(\frac{E_{kt}}{E_t}\right) \bar{\gamma}_{kt} w_{kt}. \tag{17}$$

Composition bias arises from cyclical movements in either the selection of workers along some observed dimension E_{kt}/E_t, or from the unobserved quality of employed workers $\bar{\gamma}_{kt}$. The unobserved nature of $\bar{\gamma}_{kt}$ confounds conventional methods to control for composition bias by simply reweighting the data along observable dimensions.

The model suggests two ways to control for this unobserved selection. The first is simply to estimate the distribution of γ_{jk} each year using the estimation approach of Section 3. I present the results of this approach in section 4.3 above. However, the data requirements for this method can be large, as one requires each type of worker to have a connecting cycle of occupation mobility in a given year. This limits the usefulness of the approach if one wished to study the selection of unobserved quality in workers at a high frequency within some high-dimensional partition of the economy. For instance, were a researcher interested in studying the cyclical selection patterns of workers within 4-digit NAICS codes or even at the firm level, it would be infeasible to estimate the full model each year.

The model proposes an additional reduced form method to correct for selection in unobserved human capital. Consider the change in log wages for a worker i who works in occupation k in period t and k' in period $t-1$. Suppose that worker i has human capital level γ_{ikt} in occupation k in period t. The model suggests that the worker's log wage change may be written as the sum of the change in her log human capital level for the two occupations, and the change in log labor prices in each occupation:

$$\Delta \ln \omega_{it} = (\ln \gamma_{ikt} - \ln \gamma_{ik't-1}) + (\ln w_{kt} - \ln w_{k't-1}). \tag{18}$$

If workers' human capital is fixed in the short run, then the change in human capital levels for occupation stayers is zero. In this case, the mean wage change of occupation-stayers only reflects the change in the price of labor in occupation k between periods t and $t+1$. Therefore, one may estimate the log change in the price of each occupation's labor by calculating the mean log wage change of workers who stay in that occupation:

$$\Delta \ln w_{kt} = \mathbb{E}[\Delta \ln \omega_{it} | k_{t-1}(i) = k_t(i) = k]$$

This yields a method to estimate the degree of selection in occupation k by noting

$$\Delta \ln \bar{\gamma}_{kt} = \Delta \ln \bar{\omega}_{kt} - \mathbb{E}[\Delta \ln w_{ikt} | k_t(i) = k_{t-1}(i) = k].$$

This approach relies on two basic assumptions. First, workers' human capital levels must not vary at a high frequency. This assumption is implicitly maintained in most existing composition adjustment procedures. Second, it is necessary that changes in workers' log wages solely reflect changes in their marginal product or the price of labor. If workers' wages reflect a *constant* markdown on their marginal product - whether it be from employer monopsony power or search frictions - this decomposition will remain valid, as these markdowns will be differenced out. However, if high-frequency wage movements reflect movements in factors not related to the price of labor or the human capital of the workers, this assumption will be violated. This concern is reasonable. There is a growing literature arguing that labor market monopsony power is rising, and it is well-known that wage contracts are often backloaded (Burdett and Coles, 2003). Therefore the exercise presented here should be viewed as a complement to rather than replacement of the existing literature. The current approaches to composition adjustment do not isolate movements in the price of labor from cyclical job-downgrading but remain model-free. In contrast, my approach imposes assumptions on what drives wage fluctuations and, in return, is able to account for cyclical changes in the allocation of workers to jobs.

To partially account for these concern, I first residualize the wage changes of occupation-stayers against an occupation-specific age-earnings profile, and a linear trend. The mean residuals from this regression then serve as my proxy of occupation-specific growth in labor prices. One may build a chain-weighted index of occupation prices by noting that

$$\ln \hat{w}_{kt} = \ln w_{kt_0} + \sum_{\tau=t_0}^{t} \Delta \ln w_{k\tau}$$

for some reference year t_0. Aggregating with occupational-employment shares in t_0 yields the selection-corrected wage series.

I implement this approach using data from the CPS March Supplement between 1979 and 2018, partitioning occupations according to 2-digit SOC codes. Throughout, I only include private wage and salaried workers between the ages of 21 and 60. Hourly wages are

Figure 8: Selection-Corrected Aggregate Wage Series, 1980-2018

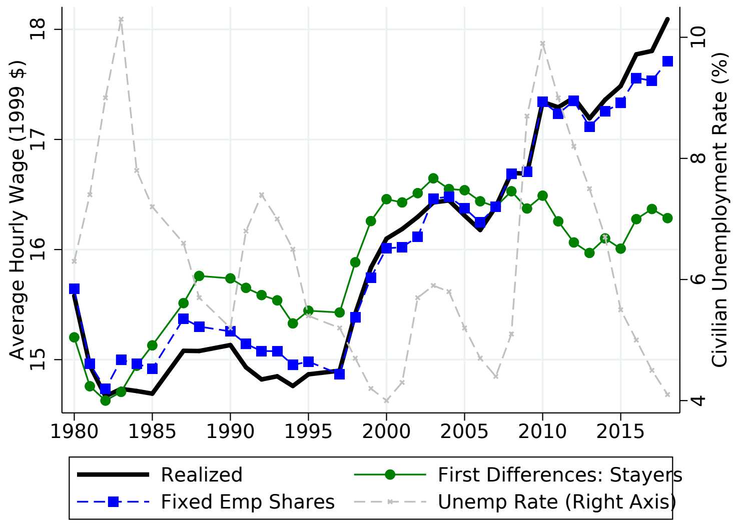

Notes: Figure plots the time series of various aggregate wage series. The solid black line plots the realized aggregate wage series. The dashed blue line with square markers fixes the employment shares of each occupation at its 2007 level. The green line with circle markers corrects for the selection of workers employed by producing a chain-weighted aggregate wage series using the wage changes of occupation-stayers. The gray dashed line plots the civilian unemployment rate against the right axis. Data come from the March Supplement of the CPS.

calculated as the ratio of total annual income, deflated by the CPI, to total annual hours. The data are weighted by the ASEC person weight when calculating aggregate wage series. The reference year t_0 is chosen to be 2007.

Figure 8 presents various aggregate wage series implied by the CPS. The solid black line corresponds to the realized real average hourly earnings.[25] For comparison, the civilian unemployment rate is plotted by the gray dashed lines against the right axis. As has been well established in the literature, the realized aggregate wage series exhibits mild procyclicality in the 1980s and early 1990s. However, beginning around the mid-1990s, the cyclicality of the aggregate wage series decline drastically, with no observable decline in aggregate mean wages in either the 2001 or 2008-9 recessions.

The blue dashed line with square markers plots an aggregate wage series holding fixed the occupational composition of the economy, E_{kt}/E_t, fixed to its 2007 level, while allowing both the labor price w_{kt} and selection of workers $\bar{\gamma}_{kt}$ to vary as in the data. Holding fixed the occupation shares of employment has little impact on aggregate wages, suggesting that

[25]Note that the wage cyclicality in Figure 8 does not align exactly with those reported in Table A1. This is due to a difference in data sources - while Figure 8 constructs the mean employment-weighted wage for prime age workers using the CPS microdata, Table A1 uses data from the BLS' CES data, which reports hours-weighted wages.

the majority of aggregate wage movements occur through within occupation effects.

The green line with circle markers plots the full selection corrected series, which isolates movements in aggregate wages arising solely from the wage movements of occupation stayers, after controlling for worker age and occupation effects. The selection corrected series diverges sharply from the uncorrected series after 2007. Whereas realized wages continued to rise during and after the recession, the corrected series shows mild declines during the recession, with accelerating wage growth as unemployment declines.

To assess the importance of occupational reallocation for wage cyclicality, I estimate the cyclicality of the various aggregate real wage series, and compare them to the Solon et al. (1994) estimates of selection-corrected wage cyclicality. I estimate the cyclicality of aggregate wage series by estimating regressions of the form

$$\Delta \ln \bar{\omega}_t = \beta_0 + \beta_1 \Delta U_t + \epsilon_t \tag{19}$$

which is the aggregate version of equation (15), for the period 1980-2018. Table 5 reports the estimated semi-elasticity of wages to the cycle β_1. Column 1 reports the cyclicality of the realized aggregate wage series from the CPS. The estimate implies that a one percentage point increase in the unemployment rate increases decreases aggregate real wages by a statistically insignificant 12 basis points, indicating muted procyclicality of wages. Holding fixed the occupation shares of employment removes all procyclicailty of wages, as shown in column 2. These findings mirror those of the figure above.

Column 3 reports the cyclicality of selection-corrected wage series. The coefficient of -0.0033 indicates that when the unemployment rises by one percentage point, aggregate selection-corrected wages fall by a statistically significant 33 basis points. Therefore, once the allocation of workers to jobs is taken into account, we observe mild pro-cyclicality of wages.

Finally, columns 4 and 5 report the results of a Solon et al. (1994) selection-correction by estimating equation (16) in the CPS for the full set of workers and for occupation-stayers, respectively. The estimates in both columns 4 and 5 are statistically indistinguishable from the -0.0033 estimated on the selection-corrected aggregate wage series. This suggests that the reallocation of workers across occupations does not drastically alter the cyclicality of aggregate wages. This is not to say that reallocation across occupations is unimportant for labor market dynamics. As Chodorow-Reich and Wieland (2019) show, frictional reallocation across industries or occupations tends to induce aggregate employment fluctuations, particularly in recessionary period. However, the cyclical reallocation of employed workers across occupations does not systematically increase or decrease mean human capital levels.

Table 5: Cyclicality of Selection-Corrected Wage Series, 1980-2018

	Realized (1)	Fixed Emp Shares (2)	Selection Corrected (3)	Solon et al. (4)	SBP: Stayers (5)
Δ Unemp. rate (%)	-0.0012	0.0003	-0.0033**	-0.0035***	-0.0037**
	(0.0018)	(0.0019)	(0.0013)	(0.0013)	(0.0015)
Obs	32	32	32	363695	357333
R^2	0.20	0.12	0.09	0.00	0.00

Notes: Table reports the cyclicality of log wages. Columns 1-3 estimate first-difference regressions at the aggregate level following equation (19). Column 1 estimates the semi-elasticity of the aggregate wage series to the cycle. Column 2 fixes the employment shares of each occupation at its 2007 level. Column 3 corrects for the selection of workers employed by producing a chain-weighted aggregate wage series using the wage changes of occupation-stayers. Columns 4 and 5 estimate first-difference regressions at the micro level following equation (16), where column 4 includes all workers and column 5 restricts attention to occupation-stayers. All regressions include linear time trends. Standard errors reported in parentheses. Columns 1-3 use White heteroskedasticity robust standard errors, while columns 4 and 5 cluster standard errors at the year level. Data come from the CPS.

It is worth nothing that the procyclicality measured in both the aggregate wage series and from the equation (16) regressions is lower than that reported in much of the literature. This is primarily due to the difference in periods. Re-estimating aggregate equation (19) for the period 1980-1994 reveals that β_1 is -0.0044. Rerunning the first-difference specification at the micro level for this time periods reveals a cyclicality of individual wages which is -0.007. These are closer to the numbers of Solon et al. (1994), which are -0.006 and -0.014, respectively. However, omitting the 1970s, which had a strong negative correlation between real wages and unemployment rates reduces the measured cyclicality.

This is not a failure of the approach, however. As the model makes clear, one should not expect to have constant elasticities of employment to wages even after controlling for selection because the nature of labor supply spillovers change the extent to which the price of occupational services varies over the cycle. Using the selection-corrected wage series implies that the implied elasticity of aggregate employment to wages in 2008-09 was 5.6, and 2.6 in 1990-91. The reduced cyclicality of wages in the selection-corrected series in the last 20 years is as a result of the increased specificity of skills. Whereas in the past, shocks to a particular set of tasks would put downward price pressure throughout the economy, the specificity of skills in recent times has limited the strength of this spillover force, dampening the movements in even the selection-adjusted wage series.

7 Broader Implications

7.1 The Role of Nominal Wage Rigidity

The fact that aggregate wages are relatively acyclical is a well-known feature of the data. Many models incorporate this fact by assuming adjustment frictions in nominal wages (Erceg et al., 2000; Smets and Wouters, 2003; Christiano et al., 2005; Smets and Wouters, 2007). In such models, real wages adjust gradually to nominal spending shocks; the sluggishness of their response is dictated by the degree to which nominal wages are rigid, and the inflation rate of the economy. Is it possible that the shifts in aggregate labor market dynamics might be caused by changed in inflation regimes and nominal wage rigidity?

There is strong evidence that wages are rigid for job-stayers. Bewley (1999) surveys numerous business owners and reports that many managers are reluctant to cut wages for fear of its effect on morale. This birthed a long literature attempting to measure the rigidity of wages using survey data (Daly and Hobijn, 2014; Kahn, 1997; Barattieri et al., 2014), employer payroll records (Altonji and Devereux, 2000; Lebow et al., 2003), or the universe of online job boards (Hazell, 2019). Many of these studies find that changes in workers earnings per hour are common, but often suffer from measurement error in household surveys, or a lack of reliable hours information. More recently, Grigsby et al. (2019) use administrative payroll records from ADP to show that, although reductions in the base wages of job-stayers are infrequent, they become more common in recessions, and other forms of compensation, such as bonuses, provide important margins of adjustment for earnings per hour. Furthermore, job-changers often receive wage cuts – a fact highlighted by Bils (1985), and Gertler et al. (2016) among others – so that in aggregate, approximately one-in-five workers received a wage cut during the Great Recession. The relative frequency of cuts in earnings per hour – the relevant concept for measured aggregate wages – has been confirmed using administrative data from Washington state by Kurmann and McEntarfer (2019) and Jardim et al. (2019).

Taken as a whole, the base wages of job-stayers do appear rigid in the data, and may therefore have important allocative consequences if base wages are a better proxy of the user cost of labor (Kudlyak, 2014). However, bonuses and job-changers provide other important margins of adjustment for aggregate average hourly earnings.

This is not to say that a change in the inflation regime had no effect on the cyclicality of real wages. Core CPI inflation ranged between 6 and 13 percent during the 1969-70, 1973-75, and 1980-82 recessionary periods. As a result, real wages could fall substantially even if nominal wages did not. However, it does not appear to be the sole driver of the changes in dynamics of aggregate employment and wages. Theories solely relying on changes in real wage rigidity will struggle to match the increase in real wages observed during the last two recessions: simple models of wage rigidity generally reduce the magnitude of nominal wage movements, without shifting the sign of those movements.

Furthermore, the changes in inflation do not quantitatively account for the change in real wage behavior. In the 1990-91 recession inflation was approximately 5%, with real wages falling by 2%. For much of the 2007-09 recession, inflation remained anchored at roughly 2.5%, about 2.5 percentage points higher than during the 1990-91 recession. If inflation were 2.5 percentage points higher in 2007-09 but the rest of the economy operated identically, then real wages would have fallen by at most 0.5%, far less than any prior recession.

Overall, nominal wage rigidity may have important allocative consequences for the economy, and likely affects the cyclical movements of real wages. Indeed, whether the rigidity observed in the microdata is sufficient to generate the observed macro patterns is an area of active debate. The arguments I make above are by no means conclusive on this issue. However, theories relying solely on wage rigidity do not account for the cyclical changes in the composition of the workforce, and therefore cannot speak to the long literature highlighting the importance of this channel. The model presented here provides an intuitive alternative explanation for the variable dynamics of employment and wages over the medium run which relies on this composition channel.

7.2 The Role of Sectoral Shocks

A long literature has developed seeking to evaluate the importance of sectoral shocks to aggregate fluctuations. Lilien (1982) argues that the counter-cyclical dispersion in sectoral growth rates is evidence for an important role for sectoral shocks in aggregate fluctuations. Abraham and Katz (1986) point out that, if sectors are differentially sensitive to aggregate shocks, Lilien's findings may not imply a large role for sectoral shocks, and argues that the pro-cyclical behavior of vacancies suggests aggregate shocks are more important.

The existing literature on this topic has yet to arrive at a consensus estimate of the importance of sectoral shocks for aggregate fluctuations. One potential reason for this is simply that the effect of sectoral shocks should not be expected to be constant. The analysis above shows that different sectors will have different impacts on aggregate employment and wages depending on the degree to which workers may transfer their skills to other activities. Additionally, the covariance of sectoral shocks across industries employing similar skills will affect the extent to which any individual industry-level shock affects aggregate employment and wages, as will the underlying distribution of skills at the time of the shock. As a result, the quantitative importance of sectoral shocks for determining aggregate fluctuations is a complicated non-stationary object which is difficult to quantify. Indeed, the framework presented here suggests a reason as to why the role of sectoral shocks may have changed over time – whereas in the past, a declining sector may have had easily transferable skills to a growing sector, this may no longer be the case.

This might explain why the empirical literature finds a shifting importance of sectoral shocks. For instance, Garin et al. (2018) employ factor analysis on industrial production tables to argue that the importance of sectoral shocks has grown over time, while Foerster et al. (2019) confirms this fact and shows that it may lead to slower trend GDP growth in a model with production networks. Meanwhile Quah and Sargent (1993) suggests that aggregate shocks play a large role for determining aggregate employment, while Forni and Reichlin (1998) find the opposite for high-frequency fluctuations using structural VAR techniques.

This paper is not the first to show that the importance of sectoral shocks may vary with the state of the economy. Chodorow-Reich and Wieland (2019) show that reallocation across sectors has a large impact on unemployment during recessions, but little effect in expansions, and build a macro model with sector-level downward nominal wage rigidity to explain these findings. Acemoglu et al. (2012) build a model in which sectors are connected via input-output linkages, and shows that shocks to the most centrals nodes in the production network generates larger fluctuations in aggregate output. My paper offers another reason why the importance of sectoral shocks may have changed over time, namely that human capital specificity differs over time and across sectors. Further, it predicts and is able to estimate *which* sectors are likely to be most important for aggregate fluctuations. Estimating the contribution of each sector to aggregate fluctuations under different shock regimes is out of the scope of this paper, but is fertile ground for future research.

7.3 Measuring Human Capital Specificity

This paper additionally contributes to a long literature measuring the specificity of human capital. The existence of job-specific human capital was first proposed by Becker (1964), which birthed a long empirical literature seeking to measure the returns to this human capital, and understand its effects. The early literature on this topic showed large returns to job tenure (Topel, 1991; Dustmann and Meghir, 2005), which may in part be due to long-tenure workers having more general experience and being well-matched to their employers (Altonji and Shakotko, 1987). Neal (1995) shows that workers who have an exogenous layoff event have bigger wage declines if they switch industry, while Sullivan (2010) shows steep earnings profiles in occupation tenure, both suggesting a role for occupation- and industry-specific human capital. Shaw and Lazear (2008) show that worker output and wages both grow steeply in tenure using detailed individual-level data from an autoglass company. Kambourov and Manovskii (2009b) shows steep returns to occupational tenure and argue that occupation-specific human capital is a more salient feature of the data than industry- or firm-specific human capital, while Kambourov and Manovskii (2008) shows that occupational and industry mobility has increased in the US since the late 1960s.

In an important paper, Lazear (2009) argued that specific human capital may be considered in a "skill-weights" framework. In Lazear's set up, jobs are characterized by the weights that they place on a discrete mix of skills. Workers with high ability levels in the skills required by a particular job may be thought to have job-specific human capital. Following this idea, recent papers have developed measures of skill remoteness between occupations using surveys of the skills required to perform the tasks of an occupation, such as O*NET in the US (Guvenen et al., 2018) or the German Qualification and Career Survey (QCS) (Gathmann and Schönberg, 2010; Geel and Backes-Gellner, 2009). A consistent finding of this literature is that workers who move to more remote occupations realize larger wage declines (Poletaev and Robinson, 2008; Nedelkoska et al., 2015), while Cortes and Gallipoli (2018) estimate a gravity equation of worker flows to claim that task-independent occupation-specific factors account for most of the variation in transition costs between occupations.

These approaches are based on surveys which ask "how important is this skill in the performance of your job?" As a result, they do not provide cardinal measures of skill transferability. Although these studies provide compelling evidence for the existence of job-specific human capital, the subjectivity and measurement error inherent in responses to surveys of this sort limit their usefulness for counterfactual analyses. To partially deal with this problem, Neffke and Henning (2013) and Neffke et al. (2017) propose a measure of skill relatedness which is equal to the flow between two industries in excess of what would be predicted given the industries' sizes, growth rates, and wage levels. Using this measure, they show that firms are more likely to diversify into industries with more related skills.

One shortcoming of such measures is that they lack cardinal interpretations. As a result, it is challenging to use these measures to perform robust counterfactual analyses. The framework presented in this paper helps overcome this issue by estimating the economy's skill distribution with microdata on wages and employment. Although this comes at the cost of some assumptions on occupation mobility and earnings dynamics, it carries the substantial benefit of being able to use the estimates in economic models of the labor market.

8 Conclusion

What determines the joint dynamics of aggregate employment and wages? This paper argues that the degree of skill transferability out of declining industries determines the effect of sectoral shocks on the aggregate labor market. I propose a model in which workers differ in their skills for various occupations, and industries combine each occupation with different weights in order to produce differentiated output. When an industry declines, its workers reallocate to other activities. If those workers have highly transferable skills, they will find employment elsewhere in the economy, limiting the aggregate employment effects of the shock but exerting downward pressure on the price of labor. If, however, those workers have little human capital for other activities, they will drop out of the employed pool, buttressing measured mean wages.

I estimate the model using 2-period panel data from the CPS and show that the variance of skills in the economy - both within worker across occupations, and across workers - grew between the late 1980s and the mid 2000s. In addition, the correlation of worker skills across high education jobs fell during this period. As a result, primitive occupation-specific labor supply elasticities rose as workers became less able to transfer their skills to other occupations, and thus became more marginally attached to the employment pool.

I calibrate the model to the US economy around the 1990-91 and 2007-09 recessions using 3-digit industry-level TFP series which have been corrected for selection in the human capital of workers employed. Although there is always positive selection in the employed pool during recessions - the lowest skill workers tend to leave employment in downturns - the selection was particularly strong in the 2007-09 recession, especially in production, construction, and tradespeople occupations. Adjusting for this selection reveals much larger shocks for key sectors during the Great Recession; for instance, the Construction sector saw a 6% decline in productivity in the selection-corrected series, but no change according to the raw BLS multifactor productivity series.

The calibrated model reveals generates an increase in real wages of 2% during the 2007-09 recession, while generating real wage declines in prior recessions, in line with the data. The change in wage and employment cyclicality come from two sources. First, were the economy to have the pre-1990 skill distribution during the 2007-09 recession, real wages would have fallen by 3% in 2007-09, with aggregate employment falling 2%. Second, the composition of shocks in 2007-09 was such that several shocks employing related skills declined simultaneously. As a result, there were limited labor supply spillovers across the rest of the economy, generating small wage movements and large employment declines. The model implies that if the industry shocks that hit the primitive labor supply curves in 2007-09 resembled those of the 1990-91 recession, then wages would have fallen by about 8%, with an aggregate employment decline of 2%.

Recognizing that shocks differ in their impact on aggregate employment and wages depending on the skill transferability of the workers they displace has implications for a host of questions commonly debated in the literature. First, it implies that sectors will differ in their impact on aggregate employment based on the transferability of the human capital they employ to alternative tasks, which in turn will depend on the selection of shocks hitting other similar sectors. Economists studying particular labor demand shocks, such as the impact of trade liberalization with China (Autor et al., 2013), automation (Acemoglu and Restrepo, 2019), or artificial intelligence (Webb, 2019) wishing to estimate the aggregate impact of such shocks may wish to account for the labor supply spillovers that such shocks generate. Doing so is fertile ground for future research.

Although the framework presented here has several attractive features, including its tractability and ease of estimation, it is ill-suited for a variety of questions due to its short-run nature. Incorporating realistic dynamics into the model is a useful direction for future research, as it would permit the study of the economy's response to long run shocks. For instance, labor demand declines arising from changes in sectoral production functions, such as a decline in the labor share, will induce workers to seek employment elsewhere. How these workers retool themselves, and how policy can best direct human capital acquisition in the presence of unobserved worker skill types, are key questions for future research.

References

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

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

_ **and Lawrence F. Katz**, "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?," *Journal of Political Economy*, 1986, *94* (3), 507–522.

Acemoglu, Daron and Pascual Restrepo, "Robots and Jobs: Evidence from US local labor markets," *Journal of Political Economy*, 2019, *Forthcoming.*

_ **, Vasco M. Carvalho, Asuman E. Ozdaglar, and Alireza Tahbaz-Salehi**, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Adão, Rodrigo, "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil," *University of Chicago Booth School of Business, unpublished manuscript*, 2019.

Altonji, Joseph G. and Paul J. Devereux, "The Extent and Consequences of Downward Nominal Wage Rigidity," in Solomon Polachek and Konstantinos Tatsiramos, eds., *Research in Labor Economics Volume 19*, Emerald Group Publishing Limited, 2000, pp. 383–431.

— **and Robert A. Shakotko**, "Do Wages Rise with Job Seniority?," *The Review of Economic Studies*, 1987, *54* (3), 437.

Alvarez, Fernando and Robert Shimer, "Unemployment and Human Capital," *University of Chicago Mimeo*, 2012.

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg, "Benchmarking Global Optimizers," *NBER Working Paper Series*, 2019, (26340).

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk, "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics*, 2014, *6* (1), 70–101.

Becker, Gary S, *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education* 1964.

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

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, *25* (2), 242–262.

Bewley, Truman F., *Why Wages Don't Fall During a Recession*, Cambridge, MA: Harvard University Press, 1999.

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

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

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa, "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 2019, *87* (3), 699–738.

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

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

Burdett, Ken and Melvyn Coles, "Equilibrium Wage-Tenure Contracts," *Econometrica*, 2003, *71* (5).

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

Chetty, Raj, "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 2012, *80* (3), 969–1018.

Chodorow-Reich, Gabriel and Johannes Wieland, "Secular Labor Reallocation and Business Cycles," *Journal of Political Economy*, 2019.

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

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

Cortes, Guido Matias and Giovanni Gallipoli, "The Costs of Occupational Mobility: An Aggregate Analysis," *Journal of the European Economic Association*, 2018, *16* (2), 275–315.

Cosar, A. Kerem, "Adjusting to Trade Liberalization : Reallocation and Labor Market Policies," *University of Chicago Booth School of Business, unpublished manuscript*, 2013, pp. 1–43.

Daly, Mary C and Bart Hobijn, "Downward Nominal Wage Rigidities Bend the Phillips Curve," 2014.

Daly, Mary C., Bart Hobijn, and Theodore S. Wiles, "Dissecting Aggregate Real Wage Fluctuations: Individual Wage Growth and the Composition Effect," *Federal Reserve Bank of San Francisco, Working Paper Series*, 2011, (May), 1–50.

Devereux, Paul J., "The Cyclicality of Real Wages within Employer-Employee Matches," *ILR Review*, 2001, *54* (4), 835–850.

Dustmann, Christian and Costas Meghir, "Wages, Experience and Seniority," *Review of Economic Studies*, 2005, *72* (April), 77–108.

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

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, 2000, *46*, 281–313.

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

Foerster, Andrew, Andreas Hornstein, Pierre-Daniel Sarte, and Mark W Watson, "Aggregate Implications of Changing Sectoral Trends," *NBER Working Paper*, 2019, *25867*, 1–55.

Forni, Mario and Lucrezia Reichlin, "Let's Get Real: Factor to Analytical Approach to Disaggregated Business Cycle Dynamics," *The Review of Economic Studies*, 1998, *65* (3), 453–473.

Garin, Julio, Pries Michael, and Sims Eric, "The Relative Importance of Aggregate and Sectoral Shocks and the Changing Nature of Economic Fluctuations," *American Economic Journal: Macroeconomics*, 2018, *10* (1), 119–148.

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

Geel, Regula and Uschi Backes-Gellner, "Occupational Mobility Within and Between Skill Clusters: An Empirical Analysis Based on the Skill-Weights Approach," *Leading House Working Paper*, 2009, *47.*

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

_ , Christopher Huckfeldt, and Antonella Trigari, "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," 2016.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz, "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *NBER Working Paper*, 2019.

Guvenen, Fatih, Burhanettin Kuruscu, Satoshi Tanaka, and David G Wiczer, "Multidimensional Skill Mismatch," *NBER Working Paper*, 2018, *21376.*

Hagedorn, Marcus and Iourii Manovskii, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 2008, *98* (4), 1692–1706.

_ and _ , "Job Selection and Wages over the Business Cycle," *American Economic Review*, 2013, *103* (2), 771–803.

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

Hazell, Jonathon, "Downward Rigidity in the Wage for New Hires," *MIT mimeo*, 2019.

Horton, John J and Prasanna Tambe, "The Death of a Technical Skill," *Working Paper*, 2019.

Jardim, Ekaterina S., Gary Solon, and Jacob L. Vigdor, "How Prevalent Is Downward Rigidity in Nominal Wages? Evidence from Payroll Records in Washington State," *NBER Working Paper*, 2019, *25470.*

Kahn, Shulamit, "Evidence of Nominal Wage Stickiness from Microdata," *The American Economic Review*, 1997, *87* (5), 993–1008.

Kambourov, Gueorgui and Iourii Manovskii, "Rising Occupational and Industry Mobility in the United States: 1968-1993 Rising Occupational and Industry Mobility," *International Economic Review*, 2008, *49* (1).

_ **and** _ , "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 2009, *76*, 731–759.

_ **and** _ , "Occupational Specificity of Human Capital," *International Economic Review*, 2009, *50* (1), 63–115.

Kudlyak, Marianna, "The cyclicality of the user cost of labor," *Journal of Monetary Economics*, 2014, *68*, 53–67.

Kurmann, André and Erika McEntarfer, "Downward Nominal Wage Rigidity in the United States: New Evidence From Worker-Firm Linked Data," *Working Paper*, 2019.

Lazear, Edward P, "Firm-Specific Human Capital: A Skill-Weights Approach," *Journal of Political Economy*, 2009, *117* (5), 914–940.

Lebow, David E, Raven E Saks, and Beth Anne Wilson, "Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index," *Advances in Macroeconomics*, 2003, *3* (1).

Lilien, David M, "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy*, 1982, *90* (4), 777–793.

Longin, François and Bruno Sonik, "Extreme Correlation of International Equity Markets," *Journal of Finance*, 2001, *46* (2), 649–676.

McFadden, Daniel, "Conditional logit analysis of qualitative choice behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, Academic Press, 1974, chapter 4, pp. 105–140.

Neal, Derek, "Industry-Specific Human Capital: Evidence from Displaced Workers," *Journal of Labor Economics*, 1995, *13* (4), 653–677.

Nedelkoska, Ljubica, Frank Neffke, and Simon Wiederhold, "Skill Mismatch and the Costs of Job Displacement," *CESifo Area Conference on the Economics of Education Papers*, 2015.

Neffke, Frank and Martin Henning, "Skill Relatedness and Firm Diversification," *Strategic Management Journal*, 2013, *34*, 297–316.

Neffke, Frank M.H., Anne Otto, and Antje Weyh, "Inter-industry labor flows," *Journal of Economic Behavior and Organization*, 2017, *142*, 275–292.

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

Quah, Danny and Thomas J Sargent, "A Dynamic Index Model for Large Cross Sections," in James H. Stock and Mark W. Watson, eds., *Business Cycles, Indicators, and Forecasting*, number January, University of Chicago press, 1993, chapter 7, pp. 285–310.

Rogerson, Richard, Robert Shimer, and Randall Wright, "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 2005, *43*, 959–988.

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

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

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz, "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, 2012, *50*, 3–50.

Schmitt-Grohé, Stephanie and Martin Uribe, "The Making of a Great Contraction with a Liquidity Trap and a Jobless Recovery," 2012.

Shaw, Kathryn and Edward P. Lazear, "Tenure and output," *Labour Economics*, 2008, *15* (4), 704–723.

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

Smets, Frank and Raf Wouters, "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, 2003, *1* (5), 1123–1175.

_ **and Rafael Wouters**, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American Economic Review*, 2007, *97* (3), 586–606.

Smith, Bruce D., "Sectoral Employment and Cyclical Fluctuations in an Adverse Selection Model," *International Economic Review*, 1995, *36* (2), 261–281.

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

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

Sullivan, Paul, "Empirical evidence on occupation and industry specific human capital," *Labour Economics*, 2010, *17* (3), 567–580.

Topel, Robert, "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority," 1991.

Verboven, Frank, "Internatinoal Price Discrimination in the European Car Market," *The RAND Journal of Economics*, 1996, *27* (2), 240–268.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," *Stanford mimeo*, 2019.

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

Appendix

Appendix A Additional Tables and Figures

Table A1: Peak-to-Trough Hours, Employment, and Wage Changes in Recent Recessions, 1969-2009

	1969-70	1973-75	1980-82	1990-91	2001	2007-09
Peak-to-Trough Changes						
%Δ Real Average Hourly Earnings	-1.4	-3.1	-5.4	-1.9	0.4	2.2
%Δ Employment	-1.0	-1.2	-2.1	-1.2	-1.2	-5.3
%Δ Total Hours	-3.4	-5.6	-4.8	-2.9	-2.2	-9.2
%Δ Employment-to-Population Ratio	-2.1	-3.6	-4.7	-1.7	-2.0	-5.6
%Δ Core CPI	6.5	14.0	27.0	3.9	2.0	3.1

Notes: Table reports the behavior of earnings and employment in the United States over the past six recessions. The first five rows show the peak-to-trough percentage change in a host of labor market indicators, while the final three rows present the ratio of peak-to-trough changes in log employment measures to the peak-to-trough change in log real wages. Each column shows the change for a separate recession. Wage and employment data taken from the Current Employment Statistics (CES) provided by the BLS.

Figure A1: Sample O*NET Questionnaire

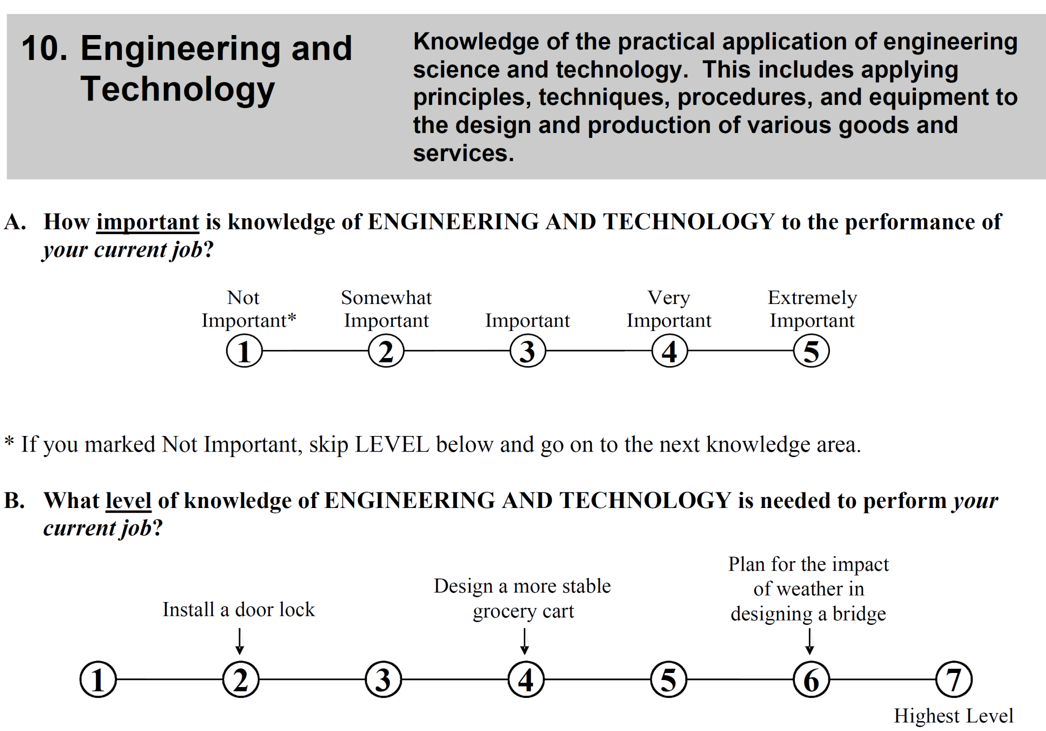

10. Engineering and Technology

Knowledge of the practical application of engineering science and technology. This includes applying principles, techniques, procedures, and equipment to the design and production of various goods and services.

A. How important is knowledge of ENGINEERING AND TECHNOLOGY to the performance of *your current job*?

Not Important* (1) — Somewhat Important (2) — Important (3) — Very Important (4) — Extremely Important (5)

* If you marked Not Important, skip LEVEL below and go on to the next knowledge area.

B. What level of knowledge of ENGINEERING AND TECHNOLOGY is needed to perform *your current job*?

(1) — (2) Install a door lock — (3) — (4) Design a more stable grocery cart — (5) — (6) Plan for the impact of weather in designing a bridge — (7) Highest Level

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

	Worker type j								
Occupation k	1	2	3	4	5	6	7	8	ξ_k
1	0.806	0.739	0.699	0.910	1.585	0.382	3.853	13.710	-2.01
2	0.040	0.777	0.704	1.010	1.672	2.889	4.063	3.806	-2.12
3	1.180	0.046	0.869	1.187	2.002	0.293	1.448	16.644	-2.45
4	0.036	0.778	0.674	0.980	1.564	2.774	3.800	12.819	-2.31
5	1.028	0.602	0.739	0.959	1.684	0.896	3.816	1.057	-2.74
6	0.034	0.798	0.656	1.019	1.565	2.773	3.735	12.375	-2.31
7	1.059	0.377	0.773	0.989	1.871	0.699	4.268	1.577	-2.92
8	1.064	0.035	0.769	1.039	1.781	2.740	3.853	1.929	-2.99
9	0.669	0.732	0.629	0.891	1.438	2.452	3.269	10.511	-2.68
10	0.865	0.718	0.627	0.943	1.539	2.381	3.197	1.206	-3.24
11	0.031	0.858	0.767	1.036	1.574	2.657	3.567	3.530	-2.88
12	0.028	0.920	0.771	1.085	1.588	2.652	1.142	10.787	-3.33
13	0.659	0.766	0.660	0.926	1.434	2.331	2.929	9.223	-3.53
14	0.731	0.905	0.719	0.125	1.677	2.662	3.392	3.601	-3.90
15	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	–
$\frac{\max_k(\gamma_{jk})}{\min_k(\gamma_{jk})}$	42.619	26.355	1.385	9.522	1.396	9.870	3.737	15.744	–

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 four 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, and the ratio of the type's skill in her best occupation relative to her worst occupation. Estimation procedure laid out in Section 3, and carried out using data from 1984-1989 in the CPS.

Table A3: In-Sample Model Fit, 1984-1989

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

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

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

	Emp. Shares		Mean Log Wage		SD Log Wage	
	Model	Data	Model	Data	Model	Data
	(1)	(2)	(3)	(4)	(5)	(6)
Non-Employed	26.04	27.72	–	–	–	–
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.

Figure A2: Correlation of Skill Relatedness in Γ with Euclidean Skill Distance in O*NET

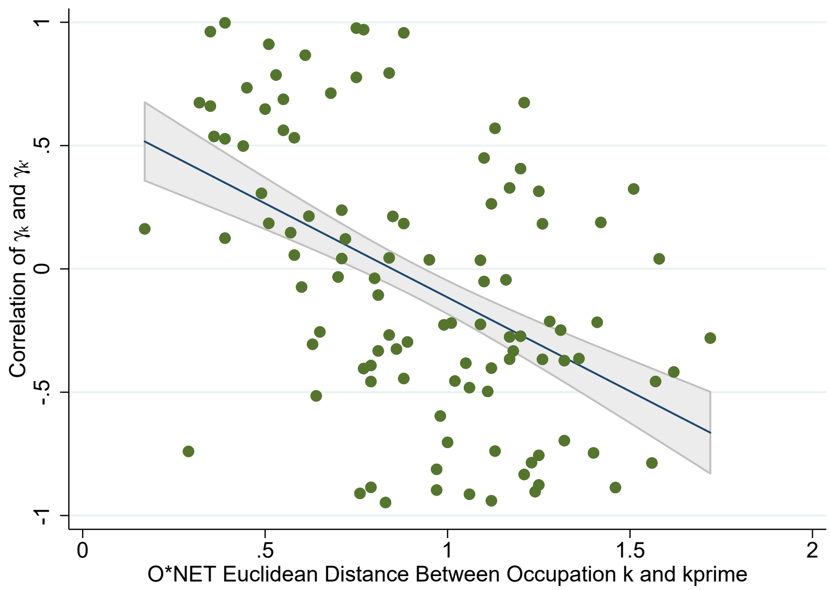

Notes: Figure compares the structurally-estimated skill transferability from the model to a common measure of skill relatedness from O*NET. Each dot corresponds to a pair of occupation clusters (k, k'). Occupations clustered by O*NET skill and knowledge vectors within terciles of the share with at least 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 4. Line of best fit reported, with shaded area representing 95% confidence interval using White heteroskedasticity robust standard errors.

Figure A3: Estimated Own- and Cross-Price Elasticities of Labor Supply by Occupation

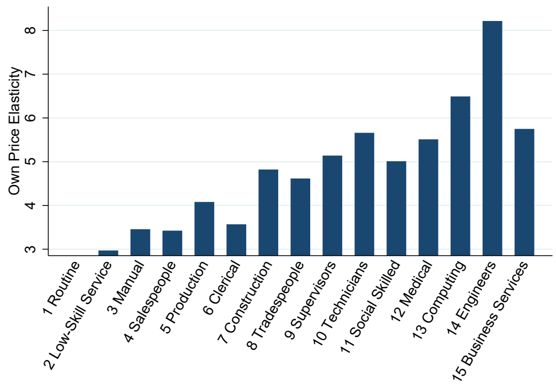

Panel A: Own Price Elasticities: 1984-1989

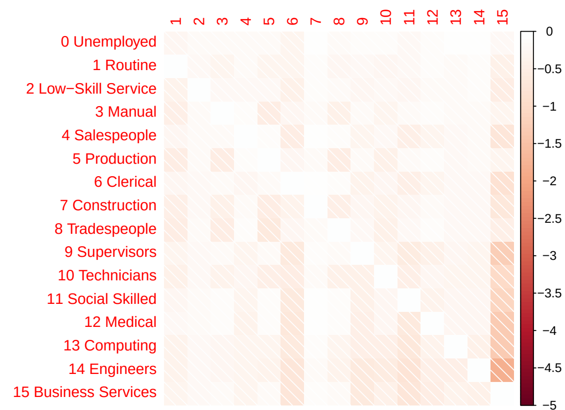

Panel B: Cross Price Elasticities: 1984-1989

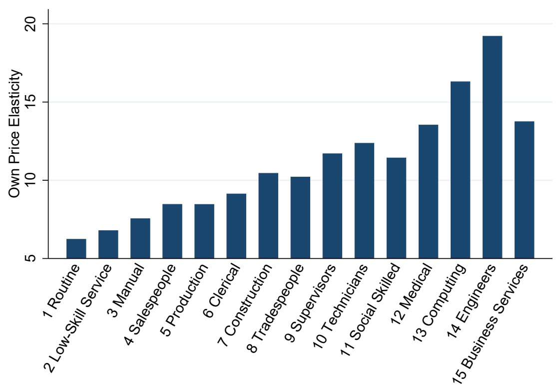

Panel C: Own Price Elasticities: 2002-2006

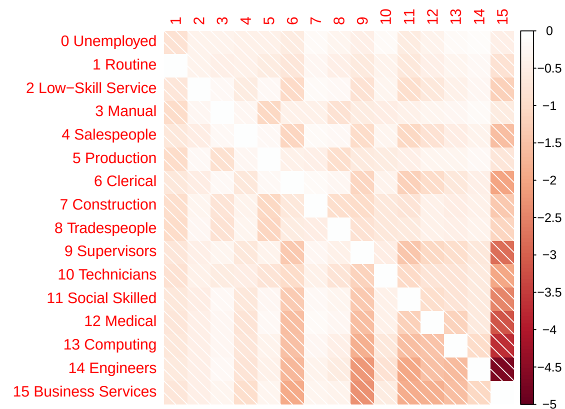

Panel D: Cross Price Elasticities: 2002-2006

Notes: Figure reports estimated own and cross price labor supply elasticities by occupation cluster. Panels A and B report the elasticities in the 1984-1989 estimation, while Panels C and D report elasticities in the 2002-2006 estimation. Panels A and C report own-price labor supply elasticities, calculated as the model-implied percentage change in employment rates in occupation k for a 1% increase in the price of labor in that occupation. Panels B and D report the matrix of cross-price elasticities. Each cell (k, k') of the figure reports the implied percentage change in the employment rate in occupation k to a 1% increase in the price of labor in occupation k'. Estimation proceeds as detailed in Section 3, using data from the CPS March Supplement.

Appendix B Identification Proof

This section proves that the vector of labor supply parameters - the mass of each type of worker m_j, distribution of wage draws for each worker type in each occupation $\Omega(\omega|k(i), j(i))$, the non-pecuniary benefits of each occupation ξ_k, and the parameters governing the distribution of type 1 extreme value shocks ν and ρ - is identified given 2-period panel data on occupations and wages. The argument presented here in fact proves non-parametric identification of the earnings distributions $F(\cdot)$ and choice probabilities $\mathbb{P}_{kk'}(j)$, following exactly the argument of Bonhomme et al. (2019). The identification and consistent estimation of the specific parameters of the model therefore follows under the assumption that the model is correctly specified, given standard arguments in maximum likelihood estimation.

Let $k \in \{1, \ldots, K\}$, and let $(k_1, \ldots, k_R), (\tilde{k}_1, \ldots, \tilde{k}_R)$ as in parts 3 and 4 of Assumption 1, with $k_1 = k$. We consider the joint cumulative distribution function of earnings in periods 1 and 2 for a given worker who moves occupations within the cycles. That is, consider workers who move from k_r to $\tilde{k}_{r'}$ for some $r \in \{1, \ldots, R\}$ and $r' \in \{r-1, r\}$. Given parts 1 and 2 of Assumption 1, the probability that a worker's wages are below $\tilde{\omega}_1$ in period 1 and below $\tilde{\omega}_2$ in period 2 is given by:

$$Pr\left\{\omega_{i1} \leq \tilde{\omega}_1, \omega_{i2} \leq \tilde{\omega}_2 | k_1(i) = k_r, k_2(i) = \tilde{k}_{r'}, m_{i1} = 1\right\} = \sum_{j=1}^{J} p_{k_r, \tilde{k}_{r'}}(j) \Omega(\tilde{\omega}_1 | k_r, j) \Omega(\tilde{\omega}_2 | \tilde{k}_{r'}, j) \tag{A1}$$

where

$$p_{k_r, \tilde{k}_{r'}}(j) = \frac{m_j \mathbb{P}_{k_r, \tilde{k}_{r'}}(j)}{\sum\limits_{j'=1}^{J} m_{j'} \mathbb{P}_{k_r, \tilde{k}_{r'}}(j')}$$

is the probability that a worker is type j given that she chooses occupation k_r in period 1 and $\tilde{k}_{r'}$ in period 2.

Now consider M sets of values for $\tilde{\omega}_1$ and $\tilde{\omega}_2$ that satisfy part 4 of Assumption 1. Note that one can augment these sets of values with a finite number of other values, including $+\infty$, while preserving the rank condition in part 4 of Assumption 1. Then, writing A1 in matrix form, we have:

$$A(k_r, \tilde{k}_{r'}) = \Omega(k_r) D(k_r, \tilde{k}_{r'}) \Omega(\tilde{k}_{r'})^T \tag{A2}$$

where $A(k_r, \tilde{k}_{r'})$ is an $M \times M$ matrix with element

$$Pr\left\{\omega_{i1} \leq \tilde{\omega}_1, \omega_{i2} \leq \tilde{\omega}_2 | k_1(i) = k_r, k_2(i) = \tilde{k}_{r'}, m_{i1} = 1\right\},$$

$\Omega(k_r)$ is an $M \times J$ matrix with element $\Omega(\tilde{\omega}_1|k_r, j)$, $\Omega(\tilde{k}_{r'})$ is similarly $M \times J$ with element $\Omega(\tilde{\omega}_2|\tilde{k}_{r'}, j)$, and $D(k_r, \tilde{k}_{r'})$ is an $L \times L$ diagonal matrix with element $p_{k_r,\tilde{k}_{r'}}(j)$. A matrix X^T denotes the transpose of the matrix X.

Note that $A(k_r, \tilde{k}_{r'})$ is observed in the data – it is simply the joint distribution of earnings in periods 1 and 2 for movers between k_r and $\tilde{k}_{r'}$ – and has rank J by Assumption 1 (4). Consider, then, the singular value decomposition of $A(k_1, \tilde{k}_1)$:

$$A(k_1, \tilde{k}_1) = U\Sigma V^T$$

where Σ is a non-singular $J \times J$ diagonal matrix, and U and V have orthonormal columns. Since $A(k_1, \tilde{k}_1)$ is observed in the data, so too can U, Σ, and V be computed. Therefore, define two further matrices:

$$\begin{aligned} B(k_r, \tilde{k}_{r'}) &= S^{-\frac{1}{2}} U^T A(k_r, \tilde{k}_{r'}) V S^{-\frac{1}{2}} \\ C(k_r) &= S^{-\frac{1}{2}} U^T \Omega(k_r) \end{aligned}$$

Note that $B(k_r, \tilde{k}_{r'})$ and $Q(k_r)$ are non-singular by Assumption 1 (4), and further that $B(k_r, \tilde{k}_{r'})$ may be constructed purely out of data objects. Moreover, for all $r \in \{1, \ldots, R\}$:

$$\begin{aligned} B(k_r, \tilde{k}_r)B(k_{r+1}, \tilde{k}_r)^{-1} &= S^{-\frac{1}{2}} U^T A(k_r, \tilde{k}_r) V S^{-\frac{1}{2}} \left(S^{-\frac{1}{2}} U^T A(k_{r+1}, \tilde{k}_r) V S^{-\frac{1}{2}}\right)^{-1} \\ &= S^{-\frac{1}{2}} U^T \Omega(k_r) D(k_r, \tilde{k}_r) \left(S^{-\frac{1}{2}} U^T \Omega(k_{r+1}) D(k_{r+1}, \tilde{k}_r)\right)^{-1} \\ &= C(k_r) D(k_r, \tilde{k}_r) D(k_{r+1}, \tilde{k}_r)^{-1} C(k_{r+1})^{-1} \end{aligned}$$

where the first equality uses the definition of $B(\cdot, \cdot)$, the second substitutes in for the definition of $A(\cdot, \cdot)$ with equation A2, and the third uses the definition of $C(\cdot)$. Letting $E_r = B(k_r, \tilde{k}_r)B(k_{r+1}, \tilde{k}_r)^{-1}$, we have

$$E_1 E_2 \ldots E_R = C(k_1) D(k_1, \tilde{k}_1) D(k_2, \tilde{k}_1)^{-1} \ldots D(k_R, \tilde{k}_R) D(k_1, \tilde{k}_R)^{-1} C(k_1)^{-1}$$

By the third part of Assumption 1, the eigenvalues of this matrix are all distinct, so that, since E_r is constructed of data objects for all r, $C(k_1) = C(k)$ is identified up to right-multiplication by a diagonal matrix and permutation of its columns.

Now note, by the properties of the singular value decomposition, that $\Omega(k) = UU^T\Omega(k)$ so that

$$\Omega(k) = US^{\frac{1}{2}}C(k)$$

is identified up to right-multiplication by a diagonal matrix and permutation of its columns. Hence, the quantity $\Omega(\tilde{\omega}_1|k,j)\lambda_j$ is identified, where $\lambda_j \neq 0$ is a scaling factor. Adding ∞ to the choice of $\tilde{\omega}_1$ values identifies λ_j and therefore $\Omega(\tilde{\omega}_1|k,j)$, as $\Omega(\infty|k,j) = 1$ for all k, j. As a result, we have identified the distribution of earnings for every type-occupation pair, up to a relabeling of types, for the set of M values chosen. Adding additional $\tilde{\omega}_1$ values to the set of M – which maintains the rank assumption – identifies the full distribution.

It remains to identify the choice probabilities of each type, as well as the distribution of types in the economy. To do so, consider $k' \neq k$, and let $(k_1, \ldots, k_R), (\tilde{k}_1, \ldots, \tilde{k}_R)$ be a connecting cycles such that $k_1 = k$ and $k' = k_r$ for some r. We have

$$A(k, \tilde{k}_1) = \Omega(k)D(k, \tilde{k}_1)\Omega(\tilde{k}_1)^T$$

Since $\Omega(k)$ and $\Omega(\tilde{k}_1)$ are identified and has rank J by the above arguments, the choice probability matrix $D(k, \tilde{k}_1)$ is identified as

$$D(k, \tilde{k}_1) = \Omega(k)^{-1}A(k, \tilde{k}_1)(\Omega(\tilde{k}_1)^T)^{-1}$$

One may apply a similar argument to $A(k_2, \tilde{k}_1)$ to show that $D(k_2, \tilde{k}_1)$ is identified. Therefore, by induction, $p_{k_r, \tilde{k}_{r'}}$ is identified, up to a labeling of types, for all r and $r' \in \{r-1, r\}$.

All that remains is to identify the distribution of types m_j. To do so, note that the marginal distribution of earnings in occupation k in period 1 may be written

$$Pr\{\omega_{i1} \leq \tilde{\omega}_1 | k_1 = k\} = \sum_{j=1}^{J} q_k(j)\Omega(\tilde{\omega}_1|k_1, j)$$

for $q_k(j)$ the probability that worker choosing occupation k is a type j given by

$$q_k(j) = \frac{m_j \mathbb{P}_k(j)}{\sum_{j'=1}^{J} m_{j'}\mathbb{P}_k(j')} \tag{A3}$$

Writing this marginal distribution in matrix form yields

$$H(k) = \Omega(k)Q(k)$$

where $H(k)$ has element $Pr\{\omega_{i1} \leq \tilde{\omega}_1 | k_1 = k\}$, and the $J \times 1$ vector $Q(k)$ has element $Q_k(j)$. Since $\Omega(k)$ is identified and has rank J, $Q(k)$ is similarly identified as

$$Q(k) = [\Omega(k)^T \Omega(k)]^{-1} \Omega(k)^T H(k).$$

Finally, m_j is identified by inverting equation A3 to arrive at

$$m_j = \frac{q_k(j)\mathbb{P}_k(j)}{Pr\{k_1 = k\}}$$

where $\mathbb{P}_k(j)$ may be treated as identified given knowledge of $p_{kk'}(j)$. Finally, the consistency of the maximum likelihood estimator, given a set of occupation clusters k, under correct specification is well-established, and yields estimates of the parameters of the model $\Gamma, \nu, \rho, \{\sigma_{jk}\}, \xi_k$ and m_j.

Appendix C Data Appendix

This section contains additional details of the data cleaning process employed in the paper. I primarily use the March Supplement of the IPUMS Current Population Survey (CPS) data. The CPS is designed to be 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. In addition, every March, a supplemental survey - the Annual Social and Economic Supplement - is administered which solicits additional information on respondents income sources and hours. 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.[26] 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 additionally drop individuals living in group quarters, retired workers, those serving in the armed forces, or employed workers with missing wage information.

I harmonize all industry codes to the 2010 NAICS coding using the crosswalks of provided by the Census bureau, and available at https://www.census.gov/topics/employment/

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

industry-occupation/guidance/code-lists.html. Similarly, 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 3 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-industry cells, which is an important input into the estimation of industry-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 industry-level employment provided by the Bureau of Economic Analysis (BEA) to construct an estimate for the total employment in each occupation-industry cell for every year.

I use the Occupation Employment Statistics (OES) to calculate the share of industry 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 A5-A7 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.

Table A5: 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 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 A6: 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 A7: Largest Employment SOC Codes within Occupation Clusters, Set 3

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

Appendix D Computational Appendix

This section outlines the computational approach taken to maximizing the log-likelihood of the data. Consider the likelihood of observing a worker with earnings ω_{i1} and ω_{i2}, and occupations k_1 and k_2. The probability of observing this worker may be written

$$l_i = \sum_{j=1}^{J} Pr\{j(i) = j\} \cdot Pr\{k_1(i) = k_1, k_2(i) = k_2 | j(i) = j\} \cdot \\ \psi(\omega_{i1} | k_1(i) = k, j(i) = j) \cdot \psi(\omega_{i2} | k_2(i) = k, j(i) = j) \tag{A4}$$

where $\psi(\omega|k, j)$ is the density of the earnings distribution for a type j worker in occupation k, evaluated at ω. Given the assumption of log-normal measurement error in wages, this distribution $\psi(\cdot)$ is log-normally distributed with a different mean μ_{jk} and standard deviation σ_{jk} for every worker type-occupation pair. Summing the log of these l_is over individuals yields the log-likelihood of the data.

To construct the probability of choosing a pair of occupations (k_1, k_2) in period 1 and 2, respectively, recall our model of occupation choice. Workers decide which occupation to pursue by maximizing their utility of doing so. Their utility is given by

$$u_{ikt} = \bar{u}_{j(i)kt} + \zeta_{ikt}$$

In order to make progress, it is necessary to assume some process for the ζ_{ikt} shocks. In particular, I assume that the ζ_{ikt} follow a Markov process, so that the distribution of idiosyncratic preferences in period $t + 1$ may depend on the realizations of those shocks in period t. Additionally, I assume that the marginal distribution of these preference shocks are distributed according to a Type 1 Extreme Value distribution with standard deviation ν. That is, the marginal distribution of ζ_{ikt} may be expressed as

$$G(\zeta) = \exp(-\exp(-\zeta/\nu))$$

for all k and t. To build the joint distribution of $\{(\zeta_{ikt}, \zeta_{ikt+1})\}_k$, we employ copula theory. In particular, we assume that the draws of ζ_{ikt} and $\zeta_{ik't}$ are independent, so that having idiosyncratic preferences for occupation k does not inform us about the preferences for occupation k'. Although strong, this assumption is standard in the literature on discrete choice (McFadden, 1974), and grants a great deal of tractability.

In addition, we allow for correlation over time of the ζ_{ikt} in a sparsely parameterized manner. In particular, we assume that the joint distribution of $(\zeta_{ikt}, \zeta_{ikt+1})$ may be described by the two marginal distributions and the Gumbel copula. The Gumbel copula is a convenient Archimedean copula, commonly employed in quantitative finance (Longin and Sonik, 2001). The Gumbel copula asserts that, if the CDF of two random variables X and Y evaluated at x and y are p_x and p_y, respectively, the CDF of the joint distribution may be given by

$$Pr\{X \leq x, Y \leq y\} = \exp\left(-\left[(-\log p_x)^{\frac{1}{\tilde{\rho}}} + (-\log p_y)^{\frac{1}{\tilde{\rho}}}\right]^{\tilde{\rho}}\right)$$

where $\tilde{\rho}$ is a parameter between 0 and 1 which pins down the correlation between the random variables X and Y. In fact, the correlation between X and Y is given by $\rho := 1 - \tilde{\rho}^2$. Applying this transformation to the marginal distributions of ζ_{ikt} and ζ_{ikt+1} implies that the joint distribution of $(\zeta_{ikt}, \zeta_{ikt+1})$ is given by

$$\exp\left(-\left[e^{-\frac{\zeta_{ikt}}{\nu\tilde{\rho}}} + e^{-\frac{\zeta_{ikt+1}}{\nu\tilde{\rho}}}\right]^{\tilde{\rho}}\right)$$

Finally, given the assumption that ζ_{ikt} is independent of $\zeta_{ik't}$, we may express the joint CDF of all $\{\zeta_{ikt}\}_{k,t}$ as

$$G(\{\zeta_{ikt}, \zeta_{ikt+1}\}_k) = \exp\left(-\sum_{k=0}^{K}\left[e^{-\frac{\zeta_{ikt}}{\tilde{\rho}}} + e^{-\frac{\zeta_{ikt+1}}{\tilde{\rho}}}\right]^{\tilde{\rho}}\right) \tag{A5}$$

Observe that this is exactly the distribution of taste shocks assumed for applications of nested logit demand functions, commonly employed in the industrial organization literature (Berry, 1994; Verboven, 1996). However, in this context, one may not simply use the standard functional forms for nested logit choice probabilities, as workers are making two choices: their occupation in period t and $t+1$. As a result, one must approximate the choice probabilities numerically.

To do so, note that the probability that a type j individual chooses occupation k in period t and k' in period $t+1$ may be expressed as

$$\mathbb{P}_{kk'}(j) = \begin{cases} Pr\{k_t = k|j\} \cdot Pr\{k_{t+1} = k|k_t = k, j\} & \text{if } k = k' \\ Pr\{k_t = k|j\} \cdot Pr\{k_{t+1} \neq k|k_t = k, j\} \cdot Pr\{k_{t+1} = k'|j, k_{t+1} \neq k, k_t = k\} & \text{if } k \neq k' \end{cases} \tag{A6}$$

That is, the probability that a worker of type j chooses the pair (k, k') is given by the probability that they first choose k, multiplied by the probability that they choose k' given that they chose k. If $k' = k$, then this is simply the probability that the worker stays in occupation k. If $k' \neq k$, then this is the product of the probability that the worker left occupation k, multiplied by the probability that they chose occupation k' given that they left occupation k.

Since the marginal distribution of the ζ_{ikt} is a mean 0 type 1 extreme value distribution with standard deviation ν (by construction), the probability that a worker chooses type j is simply given by the familiar multinomial logit form

$$Pr\{k_t = k|j\} = \mathbb{P}_k(j) = \frac{\exp(\bar{u}_{jk}/\nu)}{\sum_{\tilde{k}=0}^{K} \exp(\bar{u}_{j\tilde{k}}/\nu)} \tag{A7}$$

In addition, given the assumption that the draws of ζ_{ikt+1} are independent across k, the probability of choosing k' in period $t+1$ given that the worker left k in period t may be expressed as

$$Pr\{k_{t+1} = k'|j, k_{t+1} \neq k, k_t = k\} = \frac{\exp(\bar{u}_{jk'}/\nu)}{\sum_{\tilde{k}\neq k} \exp(\bar{u}_{j\tilde{k}}/\nu)} \tag{A8}$$

Both of these choice probabilities may be easily computed. To construct the probability that a worker leaves occupation k, note that this probability may be expressed as

$$\begin{aligned} Pr\{k_{t+1} \neq k|j, k_t = k\} &= Pr\{\bar{u}_{jk} + \zeta_{ikt} \leq \max_{\tilde{k}\neq k} \bar{u}_{j\tilde{k}} + \zeta_{i\tilde{k}t+1}\} \\ &= Pr\{\zeta_{ikt} \leq \max_{\tilde{k}\neq k} (\bar{u}_{j\tilde{k}} - \bar{u}_{jk}) + \zeta_{i\tilde{k}t+1}\} \end{aligned}$$

Let $M := \max_{\tilde{k}\neq k}(\bar{u}_{j\tilde{k}} - \bar{u}_{jk}) + \zeta_{i\tilde{k}t+1}$ denote the random variable equal to the highest utility draw for occupations $\tilde{k} \neq k$. A well-known property of the type 1 extreme value distribution is that the maximum of multiple independent type 1 extreme value distributions is itself distributed according to a type 1 extreme value. Since each of the $\zeta_{i\tilde{k}t+1}$ draws are independent, this implies that the distribution of the best outside option M is type 1 extreme value with mean $\ln\left(\sum_{\tilde{k}\neq k} \exp(\bar{u}_{j\tilde{k}} - \bar{u}_{jk})\right)$. Let the density of M for a type j worker who chose occupation k in period t be given by $\psi_{jk}(M)$. The above equation implies that

$$Pr\{k_{t+1} \neq k|j, k_t = k\} = \int G(M|j, k_t = k)\psi_{jk}(M)dM \tag{A9}$$

where $G(M|j, k_t = k)$ is the cumulative distribution function for ζ_{ikt+1} given that a worker of type j chose occupation k in period t. Observe that the condition that a worker chose occupation k in period 1 is equivalent to k being a solution to

$$k \in \underset{\tilde{k} \in \{0,\ldots,K\}}{\operatorname{argmax}} \; \bar{u}_{j\tilde{k}t} + \zeta_{i\tilde{k}t}.$$

As a result, this condition is equivalent to a restriction on ζ_{ikt}. Therefore, we may rewrite equation A9 as

$$Pr\{k_{t+1} \neq k | j, k_t = k\} = \int \left(\int G_{\zeta_{ikt+1}|\zeta_{ikt}}(M|\zeta_{ikt}) \psi_{jk}(M) dM \right) \varphi_{jk}(\zeta_{ikt}|j, k_t = k) d\zeta_{ikt} \tag{A10}$$

for ${}_{jk}(\zeta_{ikt}|j, k_t = k)$ the density of ζ_{ikt} given that a type j worker chose occupation k in period t, and $G_{\zeta_{ikt+1}|\zeta_{ikt}}(\cdot|\zeta_{ikt})$ the conditional CDF of ζ_{ikt+1} given ζ_{ikt}. Imposing the law of total probability on the expression for the joint CDF of $(\zeta_{ikt}, \zeta_{ikt+1})$ given in equation (A5) yields that this conditional CDF is given by

$$G_{\zeta_{ikt+1}|\zeta_{ikt}}(M|\zeta_{ikt}) = \frac{\exp\left(- \left[e^{-\frac{\zeta_{ik1}}{\nu\tilde{\rho}}} + e^{-\frac{M}{\nu\tilde{\rho}}} \right]^{\tilde{\rho}} - \frac{\zeta_{ik1}}{\nu\tilde{\rho}} \right) \left[e^{-\frac{\zeta_{ik1}}{\nu\tilde{\rho}}} + e^{-\frac{M}{\nu\tilde{\rho}}} \right]^{\tilde{\rho}-1}}{\frac{1}{\nu} \exp\left(- \left[\frac{\zeta_{ik1}}{\nu} + e^{-\frac{\zeta_{ik1}}{\nu}} \right] \right)} \tag{A11}$$

Finally, to calculate the conditional density of ζ_{ikt} given a choice of k, note that

$$\begin{aligned} Pr\{\zeta_{ik1} \leq x | k_1 = k\} &= Pr\{\zeta_{ik1} \leq x | \zeta_{ik1} \geq \overbrace{\max_{k'} \bar{u}_{k'} - \bar{u}_k + \zeta_{ik2}}^{M_t}\} \\ &= \int_{-\infty}^{x} [G(x) - G(M_t)] \psi(M_t) dM_t \end{aligned}$$

for $\psi(M_t)$ the density of the maximum of non-k occupation utilities. Using Leibniz's rule for differentiating under the integral gives the expression for the density of ζ_{ikt} given a choice k:

$$\begin{aligned} \varphi(x) &= \frac{dPr\{\zeta_{ik1} \leq x | k_1 = k\}}{dx} \\ &= \int_{-\infty}^{x} g(x) \psi(M_1) dM_1 \\ &= g(x) \Psi(x) \end{aligned} \tag{A12}$$

for $g(\cdot)$ the PDF of the standard Type 1 Extreme Value distribution, and $\Psi(\cdot)$ the CDF of the maximum of non-k utilities, which we know follows a type 1 extreme value distribution with mean $\ln\left(\sum_{\tilde{k}\neq k}\exp(\bar{u}_{j\tilde{k}}-\bar{u}_{jk})\right)$. One may therefore substitute equations A12 and A11 into equation A10 in order to calculate the probability that a type j worker switches out of occupation k. Numerically integrating this expression allows for relatively efficient computation of the choice probabilities given in A6.

Thus, the log-likelihood of the data may be computed by summing over the log of the individual likelihoods expressed in A4. Doing so yields, for $\theta := (\{m_j, \xi_k, \mu_{jk}, \sigma_{jk}\}_{j,k}, \nu, \rho)$ the complete vector of labor supply parameters:

$$\mathcal{L}(\theta) = \sum_i \ln\left(\sum_{j=1}^{J} \underbrace{m_j Pr\{k_1 = k_1(i)|j;\theta\} \cdot Pr\{k_2 = k_2(i)|k_1 = k_1(i), j;\theta\}\psi(\omega_{i1}|\theta)\psi(\omega_{i2}|\theta)}_{l_{ij}}\right) \tag{A13}$$

It is relatively straightforward to find local maxima of this log-likelihood function. This is because the analytical derivatives are mostly computable. The derivative of the log-likelihood function with respect to a parameter $\tilde{\theta}$ may be expressed as:

$$\sum_i \sum_{j=1}^{J} \frac{l_{ij}}{l_i}\left[\frac{\frac{\partial m_j}{\partial\tilde{\theta}}}{m_j} + \frac{\frac{\partial P_{k_{i1}}(j)}{\partial\tilde{\theta}}}{P_{k_{i1}}(j)} + \frac{\frac{\partial Pr\{k_2=k_2(i)|k_1=k_1(i),j;\theta\}}{\partial\tilde{\theta}}}{Pr\{k_2 = k_2(i)|k_1 = k_1(i), j;\theta\}} + \frac{\frac{\partial\psi(\ln\omega_{i1}|\theta)}{\partial\tilde{\theta}}}{\psi(\ln\omega_{i1}|\theta)} + \frac{\frac{\partial\phi(\ln\omega_{i1}|\theta)}{\partial\tilde{\theta}}}{\phi(\ln\omega_{i2}|\theta)}\right]$$

Analytical derivatives are computationally tractable for every piece of this gradient, with the exception of the probability of switching occupations between period 1 and 2. The functional form of these gradients is available upon request. For this piece, I employ finite-difference approximations to the gradient. Given these gradient functions, I use the KNITRO's Interior/Direct algorithm with 20 starting parameter vectors.

Appendix E Model Appendix

This section contains details of the economic model. First, I clarify the characterization of equilibrium. Next I discuss the numerical method to solve the model. Consider the problem of the sector s firm. The first order conditions for optimality for this firm is given by

$$l_{sk} = \frac{p_s x_s z_s \alpha_{sk}\left(\prod_{k'=1}^{K} l_{sk'}^{\alpha_{sk'}}\right)^{x_s}}{w_k} \tag{A14}$$

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

$$l_{sk'} = l_{sk} \left(\frac{\alpha_{sk'}}{\alpha_{sk}} \frac{w_k}{w_{k'}} \right) \tag{A15}$$

Substitute this into equation (A14) to arrive at

$$l_{sk}^{1-x_s} = p_s x_s z_s \left(\frac{\alpha_{sk}}{w_k} \right)^{1-x_s} \left(\prod_{k'=1}^{K} \left(\frac{\alpha_{sk'}}{w_{k'}} \right)^{\alpha_{sk'}} \right)^{x_s} \tag{A16}$$

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

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

$$l_{sk}^{1-x_s} = \frac{(Y)^{\frac{1}{\eta}} x_s z_s \left(\frac{\alpha_{sk}}{w_k} \right)^{1-x_s}}{M_s^{x_s} y_s^{\frac{1}{\eta}}} \tag{A17}$$

Plugging equation (A15) into the production function for sector s reveals that

$$y_s = z_s \left(\frac{M_s \alpha_{sk}}{w_k} \right)^{-x_s} l_{sk}^{x_s} \tag{A18}$$

which we may then substitute into the amended first order condition A17

$$l_{sk}^{\eta - x_s(\eta-1)} = Y x_s^{\eta} z_s^{\eta-1} M_s^{-x_s(\eta-1)} \left(\frac{\alpha_{sk}}{w_k} \right)^{\eta - x_s(\eta-1)} \tag{A19}$$

Note that 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 $\nu_s = \eta - x_s(\eta - 1)$

$$l_{s'k} = l_{sk}^{\frac{\nu_s}{\nu_{s'}}} \underbrace{\left(\frac{\alpha_{s'k}}{w_k}\right)\left(\frac{w_k}{\alpha_{sk}}\right)^{\frac{\nu_s}{\nu_{s'}}} \left[\left(\frac{x_{s'}}{x_s}\right)^{\eta}\left(\frac{z_{s'}}{z_s}\right)^{\eta-1}\left(\frac{M_s^{x_s}}{M_{s'}^{x'_s}}\right)^{\eta-1}\right]^{\frac{1}{\nu_{s'}}}}_{:=\psi_{s',s}} \tag{A20}$$

As a result, A18 implies that the equilibrium output in sector s' is given by

$$y_{s'} = z_{s'} \left(\frac{M_{s'}\alpha_{s'k}}{w_k}\right)^{-x_{s'}} \psi_{s',s}^{x_{s'}} l_{sk}^{\frac{x_{s'}\nu_s}{\nu_{s'}}} \tag{A21}$$

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

$$Y(l_{sk}) = \left(\sum_{s=1}^{S}\left[z_{s'}\left(\frac{M_s\alpha_{s'k}}{w_k}\right)^{-x_{s'}} \psi_{s',s}^{x_{s'}} l_{sk}^{\frac{x_{s'}\nu_s}{\nu_{s'}}}\right]^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}} \tag{A22}$$

Finally, we may plug this into equation (A19) 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 (A15) and (A20) 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:

$$L_k^D(\mathbf{w}|\mathbf{z}) = \sum_{s=1}^{S} l_{sk}(\mathbf{w}|\mathbf{z}) \tag{A23}$$

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

$$L_k(\mathbf{w}) = \sum_{j=1}^{J} m_j\gamma_{jk} \underbrace{\left(\frac{\exp(\bar{u}_{jk}/\nu)}{\sum_{k'=0}^{K}\exp(\bar{u}_{jk'}/\nu)}\right)}_{\mathbb{P}_k(j)} \tag{A24}$$

One may solve for equilibrium by equating labor demand for occupation k, given by equation A23, with the labor supply for this occupation, given by equation A24. Note that since workers do not have preferences over which industry 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 industry. 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

$$C = \underbrace{\sum_{j=1}^{J} m_j \sum_{k=1}^{K} \gamma_{jk} w_k \mathbb{P}_k(j)}_{I} + \underbrace{\sum_{n=1}^{N} (1 - x_s) p_s y_s}_{\Pi} \tag{A25}$$

will equal aggregate output given by equation A22.

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 industries at a fine level of aggregation without adding substantial computational burden.

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 (A24) and labor demand (A23) given 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. I supply 30 starting values to the optimizer.

Appendix F Reduced Form Evidence for Labor Supply Spillovers

In this section, I test the model's implication that a negative shock to a sector s will induce positive labor supply spillovers to sectors with skills related to s. To do so, I exploit the sudden precipitous decline in labor demand in the Mining and Utilities sectors between 2014 and 2016. Towards the end of 2014, the Chinese government, fearing the formation of a credit bubble, implemented contractionary monetary policies. Concurrently, the booming American macroeconomy prompted the Federal Reserve to raise interest rates slowly, strengthening the dollar in the process. This further put pressure on many emerging economies, whose firms had many debt obligations denominated in dollars. The result of the Chinese expansion and strengthening dollar was a steep decline in emerging markets' demand for commodities, leading to a sharp drop in prices. Crude oil fell from $106 per barrel at the end of 2014, to

Figure A4: Shock to Mining Employment

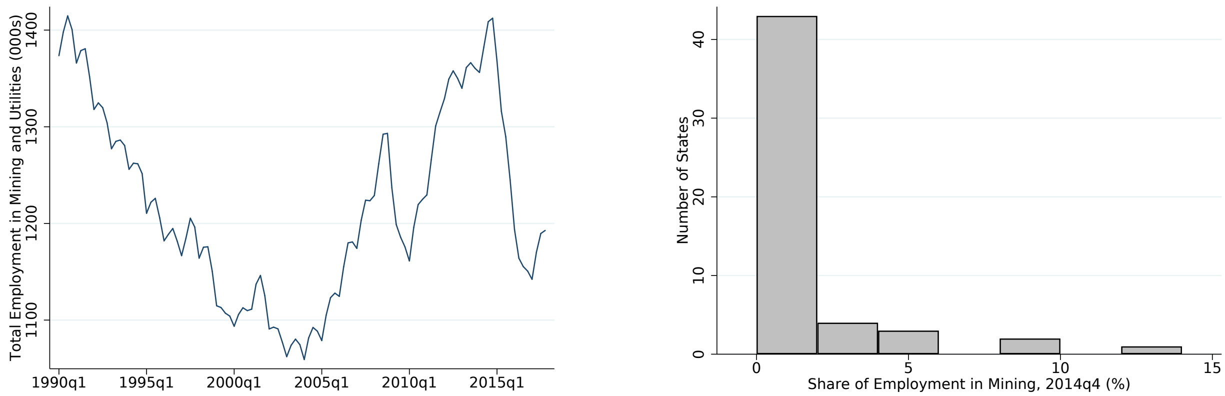

Panel A: Aggregate Mining Employment

Panel B: Histogram of State-Level Share of 2014 Employment in Mining

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

just over \$30 per barrel in early 2016, while prices for aluminum, copper, tin, and other hard commodities similarly fell. The end result was a decline in mining employment of over 30% in the span of just 2 years. The time series of aggregate mining and utilities employment is shown in Panel A of Figure A4. That the decline in employment was restricted to mining and utilities merits emphasis - this period was one of rapid expansion of employment in the US, with both employment and mean wages rising on aggregate.

This mining shock had heterogeneous impacts on local communities. For some states, such as West Virginia, Texas and North Dakota, mining constituted a significant share of employment, while for others, such as Massachusetts and Florida, mining is a relatively small share of employment. As a result, this aggregate mining shock likely produces a larger labor supply shock in states like Texas than it did in Florida, providing a laboratory to study the impact of a sectoral decline on related sectors. Let λ_r^{MINING} be the share of region r's employment that is in mining as of the fourth quarter of 2014, and let $\Delta \ln E_{MINING,-r}$ denote the percent change in mining employment in all states other than r between the fourth quarter of 20014 and the fourth quarter of 2016. We then let the predicted employment loss from mining in a region r be given by $\sigma_r = |\lambda_r \Delta \ln E_{MINING,-r}|$ - that is, the interaction of the national employment change in mining with the pre-existing share of employment in state r. If this negative labor demand shock to mining constitutes a labor supply shock to sectors with related skills, then we would expect that the share of non-mining employment to rise in sectors more related to mining, while the wages of those sectors would fall relative to unrelated sectors. These patterns should be more concentrated in states with a higher pre-existing mining share of employment.

To test these hypotheses, I construct a measure of the skill distance between sectors using the commonly-employed O*NET survey data. To do this, suppose there is a cost $c(h', h)$ of acquiring skill level h' given that a worker is already at skill level h. Construct the distance between k and k' as $d(k, k') = G(\sum_m c(h_m(k'), h_m(k)))$ for $h_m(k)$ the level of skill m required by k, and G some function. I choose $c(h, h') = \max\{0, h' - h\}^2$, and $G(x) = \sqrt{x}$ as a baseline case, which implies that $d(k, k')$ is a directed Euclidean distance.

Now one must define how related two sectors' skills are to one another. To do so, I turn to data provided by O*NET. Given the responses to this survey, one can construct vectors of skills required for each occupation, and therefore calculate the distance between each occupation as defined above. It should be noted that these survey measures do not provide a *cardinal* measure of skill relatedness, and may be subject to multiple problems with measurement error.

Finally, I aggregate to industry-level skill vectors by combining the O*NET occupation-level data. Specifically, let χ_{sk} be the share employees in sector s who are employed in occupation k (from CPS; in future can use OES), and let $h_m(k)$ be the level of skill m required for occupation k according to O*NET. Then define the level of skill m required by industry s to be the weighted average of $h_m(k)$, where the weights are the shares of s's employment in occupation $k : \chi_{sk}$. That is,

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

One can interpret this measure to be the expected skill vector a worker would require in industry s if one were to randomly samply workers in that industry. Given these skill vectors, we can then construct the distance between two sectors using the same function $d(s, s')$ as before.

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

Table A8: Response of Sectors to Mining Shocks

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

Notes: Table reports coefficients estimated from equation A26. Sectors with related skills are defined to be those sectors in the bottom quartile of skill distance with Mining sectors. Only tradable sectors in the top and bottom quartile of skill distance included. Standard errors clustered at 3-digit NAICS industry code level reported in parentheses.

the decline in mining. I estimate the following regression at the region-industry level

$$\Delta \ln y_{sr} = \alpha \cdot \mathbf{1}\{n \text{ is Related}\} + \eta \cdot \sigma_r + \beta \mathbf{1}\{s \text{ is Related}\} \cdot \sigma_r + \epsilon_{sr} \tag{A26}$$

where ΔZ is an operator which takes the difference in the variable Z between the fourth quarters of 2016 and 2014. I do this for two dependent variables y: real average weekly wages from the QCEW, and the share of non-mining employment in region r that is in sector s. The hypothesis is that $\beta > 0$ for employment, and $\beta < 0$ for wages.

The results are presented in table A8. Columns 2 and 4 control for state-industry-specific trends (i.e. long run growth between 1990 and 2014), while columns 1 and 3 do not. The table shows that sectors with skills related to mining experienced larger declines in wages and increases in employment, relative to sectors with unrelated skills, in states which had large pre-existing mining shares, suggesting that the decline in mining from 2014-2016 did indeed lead to a disproportionate positive labor supply shock for related sectors relative to unrelated sectors. A one standard deviation increase in the size of the regional exposure to the mining decline induces an increase of 4 percentage points in the share of workers employed in sectors with skills related to mining, relative to sectors with unrelated skills. This is coupled with a relative decline in log wages of approximately 4% in these sectors. These patterns are consistent with positive labor supply spillovers from the mining sectors to other sectors most related to mining.