Skill Heterogeneity and Aggregate Labor Market Dynamics*

John Grigsby

January 25, 2021

Abstract

Shifts in the composition of the workforce meaningfully affect the cyclicality of average wages. This paper develops a model to study when composition bias materially impacts aggregate wages. In the model, workers differ in their skills for a variety of occupations. Sectors employ occupations with different weights in production and are subject to idiosyncratic shocks. The model is tractable enough to be estimated using two-period panel data on workers' occupation choices and wages. The model replicates the negative comovement between employment and wages observed during the Great Recession using only shocks to sectoral labor demand. I estimate that the change in aggregate labor market dynamics has arisen because 1) skills became less transferable between the 1980s and 2000s, 2) the gap in human capital levels between least and most skilled workers has grown, and 3) the declining sectors during 2008-09 all employed a similar mix of skills, limiting the effective transferability of workers' skills. Reduced form composition-adjustment methods recover positive comovements between employment and wages in recent periods.

*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 Adrien Auclert, Stéphane Bonhomme, Ariel Burstein, Ricardo Caballero, Ali Hortaçsu, Thibaut Lamadon, Simon Mongey, Kevin Murphy, Jeremy Pearce, Gustavo Souza, Tom Winberry, Liangjie Wu, and Yulia Zhestkova, as well as seminar participants at Columbia Business School, the Federal Reserve Bank of Chicago, Harvard, Hoover Institution, Princeton, Stanford, University of Chicago, University of Michigan, Wharton and Yale. Author's contact information: john.grigsby@northwestern.edu

1 Introduction

Real wages have historically moved together with employment: both average hourly earnings and employment fell in recessions and rose together in booms. However, the movements of employment and wages have decoupled since 2000. During the Great Recession real mean wages rose despite a crash in both employment and hours, while real wages were largely flat in the subsequent recovery.[1] A sizable empirical literature highlights shifts in the composition of the workforce notably impact aggregate wage fluctuations (Solon et al., 1994; Daly et al., 2011; Devereux, 2001). Indeed, during the recent COVID-19 crisis, mean wages rose by 6% between February and May of 2020 entirely through a composition effect, as low wage workers lost their jobs at a high rate (Cajner et al., 2020). However, there is little theoretical understanding of the forces underlying composition effects. Could such shifts be large enough to generate the recent negative comovements between aggregate employment and wages? If so, what has changed to generate unusually strong composition effects on wages?

This paper develops and estimates a macroeconomic model of the labor market in order to understand what drives cyclical compositional shifts in the pool of employed workers. The model is anchored by two observations. First, the composition of industry shocks varies through time. For instance, while the 2000s saw a large construction boom and bust, there was a large cycle in the technology sector during the late 1990s. Second, individuals have heterogeneous skills that are imperfectly transferable between pursuits: economists may not easily become surgeons, for example. As a result, the aggregate impact of a collapse in demand for one set of tasks will be mediated by the skill level of workers who are employed in those tasks, as well as the ability for those workers to reallocate themselves to other productive pursuits. The aggregate response to a shock therefore depends on both the sectoral composition of that shock and the distribution of skills in the labor force.

I build a quantitative model in which multiple sectors 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 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. Sectors produce by combining occupations with differ-

[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.

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

ent weights in their production function and are subject to occupation-neutral total factor productivity (TFP) shocks, which shift their demand for labor.

A decline in a sector's TFP in this environment has three effects. The first effect is common - a decline in a sector's TFP lowers demand for the occupations heavily employed by that sector, pushing down the employment and price of those occupations. Here, there is an additional effect from labor supply spillovers: workers displaced from declining occupations 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 induces 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 declines in labor demand.

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 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 and the occupational price of labor which is absorbed into an occupation-by-time fixed effect. The approach consistently estimates these parameters of interest in two-period panel data, under some rank and exogeneity conditions.

I apply the model to study the US recession of 2008-09, which experienced increases in real wages and a crash in employment. I estimate the distribution of latent skill types and their returns to different occupations using the panel component of the March supplement of the Current Population Survey during the mid-2000s. Feeding a sequence of sectoral TFP that is taken from the data through the model generates a rise in measured aggregate wages and a sharp drop in employment during the Great Recession. Performing the same exercise for the 1990-91 recession generates positive comovements between employment and wages. Although the sole exogenous shock in the model is a shock to labor demand, the endogenous shifts in the composition of the workforce are sufficiently strong to generate the decoupling between employment and wages observed in recent periods.

To generate these negative comovements, it is necessary to have both vertical and horizontal differentiation of workers. A model, in which workers have the same average level of

human capital but differ in the occupations in which they possess it, is unable to generate strong enough selection effects to see mean wages rise in the face of negative demand shocks. On the other hand, a model, in which workers have different levels of perfectly transferable human capital (a worker fixed effect model), cannot generate increases in real wages because negative demand shocks for a subset of activities will lead workers to exert downward pressure on the price of labor elsewhere in the economy.

I estimate that the mean human capital of employed workers is generally countercyclical but has become more so since 2000, largely owing to a change in the skill distribution. 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. This is because the elasticity of non-employment to changes in the price of occupational services has grown over time. As a result, for a given set of labor demand shocks, one would expect to see larger employment fluctuations and smaller fluctuations in the price of labor in recent periods.

This shift has arisen because the distribution of skills has changed. The estimation reveals that skills have become less transferable, with the variance of skills growing within workers across occupations. In addition, the variance of skills across workers has similarly grown – the degree of absolute advantage in the economy has risen – laying the foundation for stronger selection effects today than in the past.

The composition of shocks during the Great Recession were also key to the negative comovement between employment and wages in the model. If the recession of 2009 had arisen from an aggregate shock in which all sectors declined together, then real wages would have declined approximately 6 percent. The 2009 recession was unique in that multiple sectors, all of which employ the same low-skill workers, declined at once, limiting the ability of these low-skill workers to supply their labor elsewhere in the economy. Whereas in the past, the workers expelled from a declining construction sector could find well-paid work as a miner or at a manufacturing plant, this was not the case during the Great Recession.

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. Studying the wage changes of occupation-stayers isolates shifts in the price of labor if workers' on-the-job human capital is fixed in the short run. Fixing the allocation of workers to tasks 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 principal form of selection in recent periods has been on the extensive margin of employment, rather than from employed workers moving to occupations for which they are worse-suited. Notably, despite declining cyclicality of real aggregate wages, correcting for selection following standard reduced approaches reveals no substantial change in wage cyclicality, suggesting that composition effects have become more important through time.

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)). This paper builds on this literature in two ways. First, its 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.[3] Second, the 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 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

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

[4] See Rogerson et al. (2005) for a classic survey. Chang (2011) extends these models to have sectoral

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 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.

This 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.

Many papers rationalize the large estimated elasticity of aggregate labor supply by appealing to differences between extensive and intensive margin elasticities (Rogerson and Wal-

shocks.

[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.

lenius, 2009; Chang et al., 2012). Chang and Kim (2007) show that a model with imperfect capital markets and idiosyncratic labor income risk is able to generate large cyclical movements in the labor wedge, and a low correlation between aggregate hours and productivity. In their framework, there is one-dimensional human capital that is subject to idiosyncratic shocks. Labor is supplied to a representative firm. The focus of this paper is to understand how industry shocks conspire to generate movements in the composition of the workforce, which in turn have implications for the aggregate wage.

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 results of the calibrated model and highlights the key ingredients which generate the negative comovement between employment and wages. Section 5 the changing cyclical pattern of selection 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, and provides reduced form evidence that the changing cyclicality of real wages is down to shifts in cyclical composition effects. Section 7 concludes.

2 Quantitative Model

This section builds a quantitative model with a multidimensional 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 human capital that they can supply to each occupation. Sectors hire labor 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 be considered period-by-period.

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 numeraire consumption good C. The household takes as given income from labor I, 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: $C = I + \Pi$.

2.1.2 Intermediate Goods Firms

Each sector 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 sector TFP z_s.

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

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

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, inherit risk-neutrality from the representative household, and are endowed with one unit of time which is indivisible. Workers 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. Because workers' time is indivisible, 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}. A law of one price holds for occupational

skill, and a worker of type j will earn $\gamma_{jk} w_k$ if she were to work in occupation k.

One may think of these γ_{jk} as a metaphor for the skill level of a type j worker in the various tasks employed by occupation k. 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. Each occupation provides some fixed non-pecuniary benefits ξ_k to workers.[6] Workers may additionally choose to be non-employed, in which case they receive no wages but earn an inactivity benefit, which is normalized to 0. Given this normalization, the non-pecuniary benefits ξ_k may be thought of as the negative of non-employment benefits. In addition, each worker receives an idiosyncratic preference shock ζ_{ik} for each occupation. As a result, the occupation chosen by worker i is given by

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

where $k = 0$ represents the non-employed state.[7]

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 are assumed to have marginal (cross-sectional) distribution which is type 1 extreme value with standard deviation ν. The standard deviation ν determines the weight that workers place on pecuniary versus non-pecuniary benefits of working, and therefore is a key determinant of the elasticity of labor supply. The distributional assumptions on ζ are standard in the discrete choice literature following McFadden (1974), and generate a tractable form for the occupational choice probabilities of workers:

[6] 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.

[7] 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 labor supply curves less elastic as the economy grows.

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

Aggregating workers' individual decision problems yields occupation-level labor supply 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 labor supply that clears markets, but employment as measured in the data. Because worker types differ in their labor unit endowments, the true labor supply curve in each occupation is instead given by the human-capital-weighted employment in each occupation:

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

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

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

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

$$\bar{\gamma}_k = \frac{L_k}{E_k} \tag{7}$$

Since workers are remunerated according to their human capital levels, movements in $\bar{\gamma}_k$ can shift mean earnings while leaving employment unaffected. This selection force can induce all manner of relationships between aggregate employment and measured wages, and is key to the model's ability to generate both pro- and counter-cyclical wages from exogenous labor demand shocks. Indeed, we may express the measured aggregate wage as the employment-

share-weighted average wage of each of the occupations thus:

$$\bar{\omega} = \sum_{k=1}^{K} w_k \bar{\gamma}_k \left(\frac{E_k}{E} \right) \tag{8}$$

where the symbol ω represents take-home pay. Note that take-home pay, which increases with worker skill, is distinct from the skill price w which clears markets.

The mean earnings of type j workers is given by

$$\omega_j(\mathbf{w}) = \sum_{k=1}^{K} \mathbb{P}_k(j|\mathbf{w}) \gamma_{jk} w_k \tag{9}$$

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 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. Aggregate income is given by summing over each worker type's earnings

$$I = \sum_{j=1}^{J} m_j \omega_j(\mathbf{w}). \tag{10}$$

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 of substitution (CES) production function. That is, the output of the final good is given by

$$Y = \left(\sum_{s=1}^{S} \tau_s \hat{y}_s^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \tag{11}$$

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

$$p_s = \tau_s \left(\frac{Y}{\hat{y}_s} \right)^{\frac{1}{\eta}} \tag{12}$$

Note that τ_s and z_s both affect the marginal revenue product of labor in sector s isomorphically: τ_s affects it through the price of output from sector s while z_s affects the marginal

quantity product of labor. Therefore, I normalize $\tau_s = 1$ for all sectors and interpret industry shocks as a combination of productivity and demand shocks.

2.2 Equilibrium Definition

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

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

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

The approach to characterizing equilibrium is detailed in Appendix G.

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 elucidate the general equilibrium cross-occupation spillovers that arise from workers reallocating from declining occupations to stable or growing occupations.

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 is similar across all occupations.

This structure nests three common 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 parameters determining the non-pecuniary benefits of job choice therefore permits a detailed structural estimation of flexible 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).[8] While Γ represents a useful reduced form representation of skills that grants great analytical tractability, these caveats confound attempts to decompose Γ into more fundamental components.

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 a 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 paper's 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

[8] Edmond and Mongey (2019) explore this idea further in the context of technology adoption, and show that the law of one price for skills may fail if workers are unable to unbundle their fundamental talents.

induce a very different aggregate response of employment and measured wages than a change in the price of engineering, for instance. This results from differences in the characteristics 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 generally be inframarginal to small changes in the price of engineering, and are unlikely to drop out of employment 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 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 measured relationship between aggregate employment and wages.

To build intuition for the relationship between aggregate employment and wages, consider the following partial equilibrium exercise. Suppose that there are two worker types, each accounting for half of the population, and two occupations. Using the model for labor supply, one can trace out an aggregate employment-wage schedule as the price of occupational labor services changes. Specifically, one can vary the vector of occupation prices $\mathbf{w}$ and plot the relationship between $E(\mathbf{w})$ and $\bar{\omega}(\mathbf{w})$ as implied by equations (6) and (8), respectively. I do this for three specifications of the Γ matrix:[9]

$$\Gamma^{(RA)} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \qquad \Gamma^{(AA)} = \begin{pmatrix} 1.5 & 1.5 \\ 0.5 & 0.5 \end{pmatrix} \qquad \Gamma^{(CA)} = \begin{pmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{pmatrix}. \tag{13}$$

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.

First suppose that price movements are such that $w_1 = w_2 = w$: that is, both occupations

[9]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.

have 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), or if the economy were subject to an aggregate shock. Varying the price of labor w induces workers to choose occupations according to the decision rule of equation (2). Using these choices, one can then trace out the relationship between aggregate employment (equation 6) and aggregate wages (equation 8) as w varies.

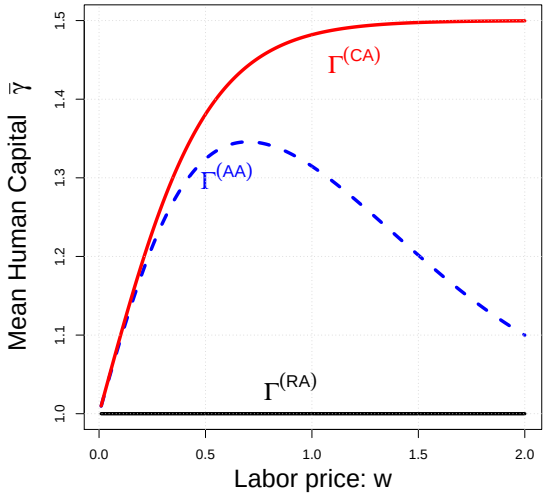
The results of this exercise for the three Γ matrices are presented in Panels A and B of Figure 1. Panel A plots the mean human capital level of employed workers $\bar{\gamma}$ as we vary the price of labor in both occupations w. Panel B plots the implied relationship between aggregate employment and wages.

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

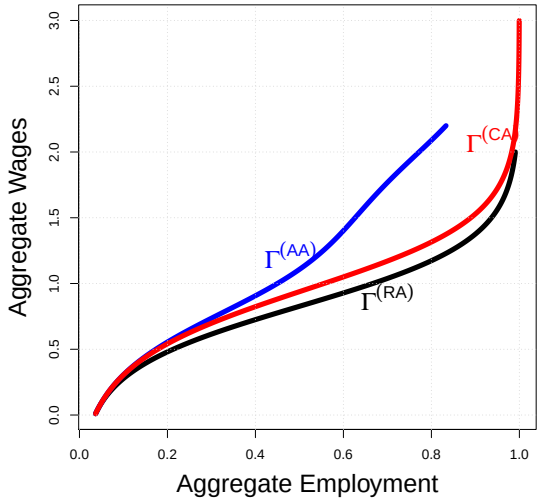
The red line shows the response under the comparative advantage skill matrix. In this case, the mean human capital level $\bar{\gamma}$ is increasing in the price per unit of labor w. To see why, consider the case in which $w = 0$. When the price of labor is 0, there is no gain for workers to sort into an occupation in which they have comparative advantage: they will earn nothing regardless of which occupation they choose. Thus there is no sorting. As the price of labor rises, so too do the gains from working in one's best occupation. Thus workers sort more and more as the price of labor increases, leading to the increasing relationship between $\bar{\gamma}$ and w. As a result, the aggregate relationship between employment and wages resembles an upward shift of the representative agent schedule as workers sort into their occupation of skill, thereby realizing higher wages for any given employment rate.

When there is absolute advantage (the blue dashed line), the aggregate wage-employment schedule becomes relatively inelastic at low levels of employment. This is because of a selection effect. Again, when the price of labor is 0, absolute advantage does not affect allocations, as both low and high type workers are equally unlikely to work. As the price of labor increases, high type workers disproportionately enter the labor force, leading to growing positive selection at low levels of w. This leads to higher wages than observed in the representative agent economy for low levels of employment. Eventually, nearly all of the high type workers are employed. When this occurs, additional increases in the price of labor w only impact 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 generates little increase in employment for a

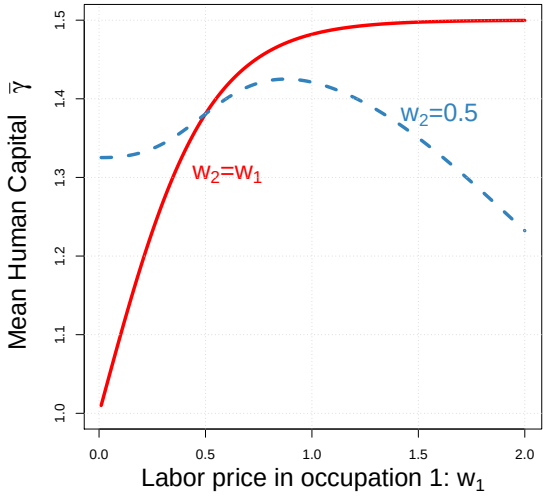
Figure 1: Aggregate Employment-Wage Schedule As Vary Γ and Occupation Prices



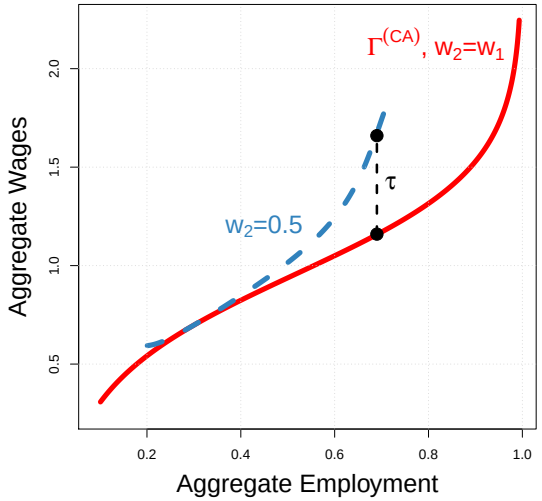
Panel A: Mean Human Capital Levels
$w_1 = w_2 = w$



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



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



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

Notes: Figure presents the behavior of the labor market induced by movements in occupational labor prices w_k in a two occupation, two-type labor supply model. Panels A and B plots the implied movements when the price of labor in occupation 1 is constrained to equal the price in occupation 2, while Panels C and D plots the implied curves when occupation 2's labor price is fixed at 0.5 and occupation 1's price is allowed to vary between 0 and 1.5. Panels A and C plot the mean human capital of employed workers $\bar{\gamma}$ against the prevailing price of labor, while Panels B and D plot the implied relationship between aggregate employment and wages. The solid black line is the representative agent curve with $\gamma_{jk} = 1$ for all j and k, while the solid blue line reports the curve when Γ has worker fixed effects. The red solid line is the curve when $w_1 = w_2$, and Γ exhibits comparative advantage. The blue dashed line is the curve when w_2 is fixed to 0.5, and Γ exhibits comparative advantage. Γ matrices defined as in equation (13).

given wage movement, yielding a steep relationship between wages and employment. 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 relationship between employment and measured wages if an increase in the price of labor induced a large enough inflow of low-type workers.

The analysis thus far assumes that the price of both occupations' services move in tandem. Now consider the opposite 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 Panels C and D 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.

The red line recreates the red curves from panels A and B, while the blue dashed line shows the curves after fixing w_2 at 0.5. Fixing the wage in occupation 2 steepens and shifts the aggregate relationship between employment and wages inward. This is because inducing type 2 workers to enter the labor force requires large movements in the price of occupation 1. For high values of w_1, the majority of type 1 workers are employed, and type 2 workers are only marginally responsive to the movements in the price of labor.

This has important implications for macroeconomic accounting frameworks. If one were to assume a representative agent labor supply curve, one might estimate that curve to be given by, for instance, the red line in Panel D. Realizing a point of data on the blue dashed line would therefore be rationalized either as evidence that workers' supply curve has shifted, or that workers are off their frictionless labor supply curve. This wedge between the realized data and the estimated labor supply curve is depicted on the figure by τ and may be interpreted 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). Therefore, compositional shifts can microfound the labor wedge or labor supply shocks which have been shown to be important to account for recent business cycle fluctuations (Brinca et al., 2016).

This exercise highlights two primary reasons why real wages may become countercyclical over time. First, it is possible that the distribution of skills changes over time. For example, if the skill distribution begins to exhibit a larger degree of absolute advantage, there may be more scope for selection in the employed pool, while if there are increases in comparative advantage, there may be increased sorting through time. Second, the distribution of labor demand shocks hitting the economy may shift the distribution of labor prices in each occupation. This would induce reallocation of workers to tasks, thereby moving the relationship between aggregate employment and wages.

2.3.3 General Equilibrium Cross-Occupation Spillovers

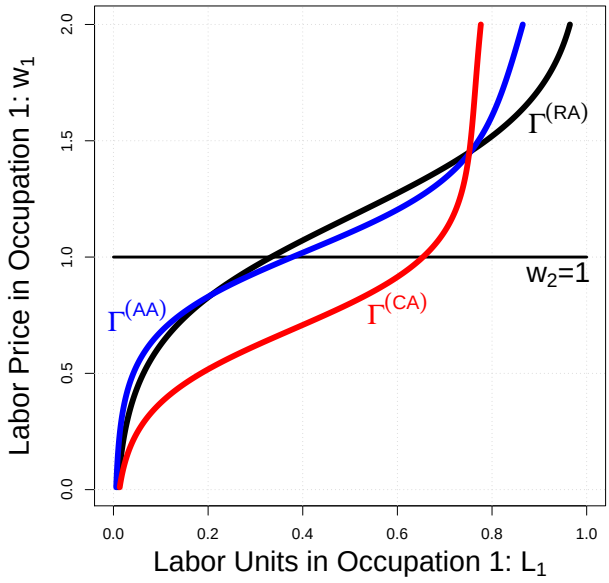
When the price of labor falls in one occupation, workers begin to seek employment elsewhere. This shifts out the labor supply curve in occupations unaffected by the initial shock. As a concrete example, if demand for mining workers falls, some workers previously employed in mining may seek employment in manufacturing or construction, thereby exerting downward pressure on the price of labor in those two occupations. Indeed, in Appendix H, I present reduced form evidence that such a force exists using an exogenous decline in the demand for mining labor between 2014 and 2016. In this subsection, I formalize and build intuition for these general equilibrium labor supply spillovers.

Consider the simple two-occupation, two-type labor supply model. Note that the labor supply curve in each occupation $L_k(\mathbf{w})$ as defined in equation (5) depends on the entire vector of labor prices $\mathbf{w} = (w_1, w_2)$. Nevertheless, we can plot the labor supply curve in occupation 1 given a fixed price of labor in occupation 2. This curve relates the units of labor supplied to occupation 1 (L_1) against the price of labor in occupation 1 (w_1) and are plotted in Panel A of Figure 2 for the three skill matrix specifications. The figure shows standard upward sloping labor supply curves for each of our skill matrices. The labor supply curve under a representative agent skill matrix strongly resembles that of the absolute advantage skill distribution. In the absolute advantage case, there are two offsetting effects. Low-skill workers do not respond to increases in the price of labor much, which, *ceteris paribus* makes the labor supply curve more inelastic. However, for a given increase in the price of labor, it is principally high skill workers who enter occupation 1. Since these workers carry more human capital with them, this will push the labor supply curve towards becoming more elastic in labor-units space. The lack of response of low skill workers is almost exactly offset by the fact that the responding high skill workers are more productive.

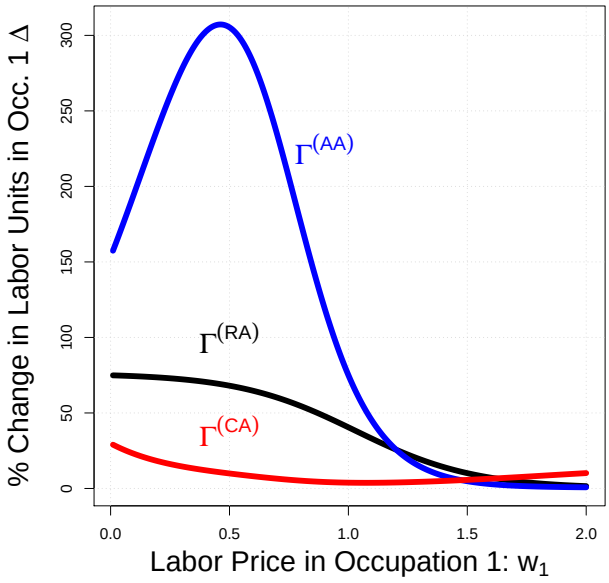
By contrast, the economy with a comparative advantage skill distribution exhibits a different labor supply curve. This is due to skill specificity. Type 1 workers are very responsive to increases in w_1. As a result, for low levels in the price of labor, the labor supply curve is very elastic under a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

Panel A is plotted assuming that w_2 were equal to 1. Suppose now that the price of labor in occupation 2 exogenously moved to $w_2 = 0.5$. This simulates a large negative demand shock to occupation 2. Because of this decline, the relative value of working in occupation 1 increases for workers. This leads to an outward shift of the labor supply curve for occupation 1. One can quantify the magnitude of this shift by measuring the horizontal movement in the

Figure 2: General Equilibrium Labor Supply Spillovers



Panel A: Labor Supply Curves in Occ. 1



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

Notes: Figure shows the behavior of occupation 1's labor supply given exogenously specific prices of labor. Panel A plots the labor supply curve in occupation 1 if the price of labor in occupation 2 is fixed at 1. Panel B plots the percentage horizontal shift in occupation 1's labor supply curve when the price of labor in occupation 2 falls to 0.5, as described in equation (14). The black line has a representative agent skill matrix, the blue line has a worker fixed effect skill matrix, and the red line has a comparative advantage skill matrix, as defined in equation (13).

labor supply curve for every given level of w_1. Specifically, one can calculate the percentage change in labor units supplied to occupation 1 induced by the change of price in occupation 2 for every given price of labor w_1:

$$\Delta(w_1) = \frac{L_1(w_1|w_2 = 1) - L_1(w_1|w_2 = 0.5)}{L_1(w_1|w_2 = 1)} \tag{14}$$

This function $\Delta(w_1)$ is plotted in Panel B of Figure 2 for our three skill matrices. The more specific are skills, the less impact will a shock to the price of labor in occupation 2 have on the labor supply curve of occupation 1. This is captured by the fact that the red curve representing the comparative advantage skill distribution is substantially below the black and blue curves, which both have perfectly transferable skills for workers between the two occupations. The blue curve representing the absolute advantage skill matrix has the largest shift for low levels of w_1. This is because the workers who move from occupation 2

to occupation 1 are principally high type workers.

To summarize, this section illustrates that skill heterogeneity exerts three novel effects on the relationship between aggregate employment and wages. First, shifts in the distribution of skills change the sorting patterns in response to a shock to the price of labor, thereby affecting the relationship between aggregate employment and measured wages. Second, changes in the relative price of labor between two occupations similarly affect sorting patterns, which implies that idiosyncratic sector shocks influence the relationship between aggregate employment and wages. Finally, the multidimensional nature of workers' skills implies that shocks to the price of labor in one occupation induce shifts in the labor supply curves of other occupations, thereby exerting pressure on the price of labor elsewhere in the economy. This force will be especially strong if skills are easily transferred between tasks. The matrix of skills Γ is a key determinant of the strength of each of these forces.

3 Model Estimation

This section describes the procedure and data used to estimate the model and to calibrate a sector-level TFP series which correct 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). Estimation proceeds following a maximum likelihood approach. I assume that individual wages in period t are observed with multiplicative measurement error ϵ_{it},[10] 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

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

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 realizations 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*

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

   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 of a simple likelihood func-

tion 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 large, representative, 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}. 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.

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.[11]

The third item of Assumption 1 requires that any two occupations belong to a connecting cycle for every type of worker. This does *not* require that every worker type must flow between every pair of occupations (k, k') bilaterally. Rather, it imposes graph connectedness in the sense of Abowd et al. (1999). This will always hold under the model, given the distributional assumptions on the idiosyncratic preference shocks. In addition, it must be that workers of different types flow in different ways, which occurs 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 through maximum likelihood. Since the identification argument follows that of Bonhomme et al. (2019) almost exactly, I relegate it to Appendix D.

To construct the likelihood of the data, consider the likelihood of observing a single worker i for two periods, labeled 1 and 2. This worker chooses occupation k in period 1 and k' in

[11] Bonhomme et al. (2019) show that first-order Markov processes for wages may be accommodated with four-period panel data. The identification problem in two-period panels is that if wages are persistently high for a given individual, one is unable to identify whether that is because they have high γ for their job, or because idiosyncratic wage draws are highly persistent.

period 2, realizing wages ω_{i1} and ω_{i2} in periods 1, and 2, respectively. Let the parameters of the model be given by θ, which includes $\gamma_{jk}w_k$,[12] ξ_k and the parameters governing the idiosyncratic taste shocks ζ_{ikt} and measurement error θ_ϵ. Let $\psi(\omega|k,j,\theta)$ be the density of idiosyncratic wages implied for a type j worker in occupation k. Unemployed workers' wage density has mass 1 and does not affect the likelihood function. The likelihood of observing this worker may be written as

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

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

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

$$\mathcal{L}(\theta) = \sum_i \sum_{k=0}^{K} \sum_{k'=0}^{K} \mathbf{1}\{k_1(i)=k\}\mathbf{1}\{k_2(i)=k'\} \ln l_i(k,k',\omega_{i1},\omega_{i2}|\theta) \tag{15}$$

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

Assumption 2. ***Distributional Assumptions***

1. *The measurement error in wages ϵ_{it} is log-normally distributed with unit mean 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*

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

[12] Note that w_k is common to all workers, and thus is absorbed into an occupation-by-time fixed effect in estimation.

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. Workers switch between occupations either because they draw new taste shocks or because the price per unit of their labor w moves in different occupations. The assumption that taste shocks are independent through time is strong, as it generates close to random mobility. Stickiness in occupation choices therefore principally loads into small variance in ζ_{ikt}, which is in turn related to the micro elasticity of labor supply. Nevertheless, as discussed in section 5, the estimated value of this variance closely aligns with those found in the literature estimating micro labor supply elasticities.

Finally, the third item of Assumption 2 normalizes the γ_{jk} to have unit mean within an occupation. This normalization disentangles the variation in mean occupation wages that arises from the price of occupation services w_k and the workers' ability γ_{jk}.[13]

Intuitively, identification is achieved through occupation switchers. Workers switch occupations either because the market price of labor moves or because of their idiosyncratic taste shocks. When a worker switches occupations, her type j is fixed across that move. Since the ϵ_{it} are i.i.d. across the job switches, her wage change is principally determined by differences in $\gamma_{jk}w_{kt}$. However, the w_{kt} is the market price of human capital which affects all workers equally. As a result, the w_{kt} act similarly to an occupation-by-time fixed effect, for which the marginal distribution of wages in occupation k in period t is highly informative. After controlling for changes in the price of labor with these fixed effects, the last determinant of the worker's wage change are her relative skills in source relative to destination occupation. The distribution of wage changes for workers switching from occupation k to k' therefore informs the distribution of relative skills in the economy. In addition, the frequency of moves from occupation k to k' further pin down the relationship between γ_{jk} and $\gamma_{jk'}$. Finally, the normalization that the mean skill level in each occupation equals 1 converts the distribution of relative skills into a distribution of skill levels.

The parameters governing the non-pecuniary benefits are principally affected by occupation choices and flows. The likelihood that a worker chooses low expected utility jobs is determined by the variance ν of the idiosyncratic taste shocks. 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

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

because engineering is an unpleasant occupation, but rather because the annualized cost of acquiring and 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 data from the Current Population Survey (CPS).[14] I then split occupations into terciles of educational attainment to rank occupations according to their general skill requirement.

Next, I cluster occupations within each education tercile according to the skill content required by the occupation. To do this, I employ data from O*NET, which surveys thousands of occupation holders about the level of skill and knowledge required to perform their job. Skills include both hard skills, such as mathematics and science, and soft skills, such as critical thinking and social perceptiveness. Knowledge categories include specific occupational knowledge such as Personnel and Human Resources and Foreign Languages. A sample questionnaire from O*NET is reproduced in Appendix 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.

A brief overview of the clustered occupations is provided in Table 1, with a fuller picture provided in Appendix E. 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

[14] Throughout, I harmonize occupation codes to follow the 2010 Census occupation coding provided by IPUMS, and use the crosswalk to detailed SOC codes from census. More details are in Appendix E.

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.

occupations. For instance, cluster 12 pairs nurses together with surgeons. It is natural that these occupations might be clustered together within a broader medical clustering. However, surgeons are generally thought to be higher skill workers than are nurses. This would be captured by the γ_{jk} - the worker types with high γ_{jk} for medical occupations may be thought of as surgeons, while those with lower γ_{jk} may be nurses.

With the occupation clusters in hand, I turn to the estimation of the Γ matrix. I assume that the number of types J is equal to 8.[15] I use the March Supplement of the CPS going back to 1984 focusing on workers aged between 21 and 60 years old. I link individuals over time following Flood and Pacas (2008) to generate two-year panels; more details are provided in Appendix E. Worker earnings ω_{it} are measured by the total labor income of workers over the prior year, deflated by the CPI-U.[16] I drop workers who report earning less

[15]The macroeconomic model simulation results do not change if I choose $J = 10$, or $K = 20$, but the standard errors on the estimates get large. Choosing $J = 5$ does not permit sufficient heterogeneity to generate strong selection patterns. In contrast, Bonhomme et al. (2019), on which the estimation is based, allowed for the equivalent of $K = 10$ and $J = 6$.

[16]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

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 separately before the recession of 1990-91 (1984-1989).

3.3 Estimated Skill Distributions

Table 2 reports the transpose of the estimated matrix Γ, along with the mass of each type of worker m_j for the period 2002-2006. Each column reports the γ_{jk} vector for a given worker type j, while each row reports the γ_{jk} entry for a given occupation k. Worker types are ordered according to the mean of their γ_{jk} vector, reported in the row labeled $\mathbb{E}_k[\gamma_{jk}]$. In addition, the final column reports the non-pecuniary benefit of each occupation ξ_k, while the final two rows report the variance and geometric range of each column vector. The corresponding table for the 1984-1989 period is reported in Appendix A.

The table shows, for instance, that a type 1 worker supplies 0.81 units of human capital to routine occupations (cashiers, security guards etc.), but only 0.05 units of human capital to skilled business services occupations (such as financial analysts or management consultants). In contrast, type 6 workers supply 2.57 units of human capital to business services occupations, but only 0.38 units of human capital to routine occupations. Recall that the γ_{jk} are normalized to have unit mean (weighted by worker type shares) within each occupation. As a result, these γ_{jk} may be interpreted as the amount of human capital a type j worker has in occupation k relative to a mean worker in the economy.

The estimation is an excellent fit in-sample. For brevity, the details of the model fit are provided in Appendix A and only briefly summarized here. The correlation between the estimated mean and variance of occupational wage distributions with those of the data is between 0.99 and 1. Similarly, the employment shares implied by the model match the data almost exactly. In addition, the model accurately predicts the share of flows of from occupation k that go to any other occupation k'. At the (k, k') level, the correlation of occupation flows predicted by the model to those in the data is 0.84. However, the model overpredicts the share of people who switch occupation. This is due to the i.i.d. assumption on the idiosyncratic preference shocks ζ_{ikt}.The model's performance out of sample will be explored more fully in Section 4.

differences from these observables as reflecting differences in workers' human capital.

Table 2: Estimated Γ, m_j and ξ_k, 2002-2006 CPS

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

Notes: Table reports the estimated matrix of skills Γ, mass of worker types m_j for the period 2002-2006. A cell (k, j) in the matrix reports the estimated units of human capital that a worker of type j supplies to occupation k on average. The final column reports the net non-pecuniary benefits of each occupation ξ_k. The final 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.

3.4 Externally Calibrated Parameters

Table 3 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 sectors S to match the number of 3-digit NAICS sectors. I assume that the number of worker types J is 8,[17] and that the elasticity of substitution η between intermediate sectors in the production of the final good is 4, following Broda and Weinstein (2006).[18]

[17]Reducing the number of types to 5 or fewer removes the model's ability to match the aggregate data as there is little scope for composition effects with few types. Increasing the number of types renders estimation more noisy, but does not drastically change the macro implications of the model.

[18]Broda and Weinstein (2006) estimate the mean elasticity of substitution across 3-digit SITC products,

Table 3: Calibration Overview

Parameter	Description	Source
Structural Estimation		
γ_{jk}	Effective Labor supply of type j	Maximum Likelihood
σ_{jk}	Variance of idiosyncratic Wage Draw	Maximum Likelihood
m_j	Share of workers who are type j	Maximum Likelihood
ξ_k	Compensating Differential of Occ k	Maximum Likelihood
ν	S.D. of T1EV shocks	Maximum Likelihood
External Calibration		
S	Number of Sectors	57 (# 3-Digit NAICS)
J	Number of types	8
K	Number of occupations	15
η	Elast. of Subs. Between Sectors	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 Sector s	BEA Labor Share
α_{sk}	Share of Occupation k in Sector s	OES Share in Wage Bill
z_{st}	TFP series for sector s	Adjusted VA/Worker

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 sector s, x_s, will be equal to labor's share of value added in sector s, while α_{sk} will be the share of sector s's wage bill that is accounted for by occupation k.[19] Hence x_s is chosen to match the BEA's estimate of the labor share of production in each sector, while the α_{sk} is chosen to match the share of the wage bill in each of the 15 occupation clusters in the BLS' Occupation Employment Statistics data series. These quantities are assumed to be fixed to the average share in each sector over the period 2002-2006.

rather than sectors. The true elasticity of substitution across 3-digit sectors may therefore be somewhat lower than 4. Reducing the elasticity of substitution across sectors would have the effect of reducing the dispersion of labor demand shocks for each occupation, as a shock to a particular sector would be partially capitalized into the price of that sector'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.

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

3.4.1 Estimating Sector-Level TFP Series

I estimate TFP using a Solow residual approach. The traditional method for calculating sector 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 a sector'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 sector 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 sector s. Because workers are indifferent over sectors conditional on their occupation, the total labor units employed in occupation k in sector s are

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

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

$$\ln z_{st} = \ln \text{Value Added}_{st} - x_s \sum_{k=1}^{K} \alpha_{sk} \ln(\bar{\gamma}_{kt} E^s_{kt}) - (1 - x_s) \ln(\text{Non-Labor Input}). \quad (16)$$

The employment in each sector in each occupation, E^s_{kt}, is observed in the data. The challenge arises because $\bar{\gamma}_{kt}$ is not observed. 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 sector-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 sector.[20] These data go back to 1987. To calculate the employment of each occupation in

[20]Note that these value added data include sectoral prices. Thus the z_s arising from this calculation may be interpreted as reflecting shocks both to productivity and to demand for sector s as described at the end

each sector, I combine data from the Quarterly Census of Employment and Wages (QCEW) with data from the CPS. The QCEW provides the total employment and wages by sector and locale using administrative data derived from tax records. Using the CPS, I calculate the share of employment in each 3-digit NAICS sector that is accounted for by each of the 15 occupation clusters. Combining these gives an estimate of the total number of employees in each sector-occupation pair. Finally, I use equation (16) to estimate sectoral TFP series.

This adjustment is meaningful. Appendix Table A5 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 smallest sectors in both 1990 and 2008, measured by value added. 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 Application: Labor Market Dynamics During the Great Recession

I estimate the labor supply parameters in the period 2002-2006 and feed through a sequence of realizations of selection-adjusted TFP levels from 2008-2012. Figure 3 plots the aggregate labor market dynamics implied by the model. The figure plots the level of mean average earnings of employed workers and the measure of workers employed, relative to the pre-recession peak of 2008. The blue solid line plots the evolution of real wages, while the green dashed line plots the evolution of employment. 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.

Table 4 reports numbers associated with these patterns. Each row of the table represents the movement of aggregate labor market variables between 2008 and 2009, either in the data

of Section 2.1.4.

Figure 3: Aggregate Wage and Employment Responses in Calibrated Model, 2008-2012



Notes: Figure plots the model-implied aggregate behavior of wages and employment in response to the calibrated sectoral TFP series around the Great Recession. Wages and employment normalized to be 1 in 2008; therefore figure plots employment and wage behavior relative to their levels in 2008. Labor supply parameters estimated using data from 2002-2006.

or a particular calibration of the model. Column 1 shows the implied change in real wages, column 2 shows the change in employment, while column 3 shows the ratio of the change in employment to the change in wages. This ratio is the elasticity of labor supply that would be inferred by a frictionless representative agent model with labor demand shocks.

The table shows that real wages rose by 2.7% between 2008 and 2009, while employment fell by 8.8% in the data.[21] The calibrated model reveals a wage increase of 2.5% and employment decline of 10.4% over the same period. Therefore, the selection forces endogenously generated by the model are sufficiently strong to generate the negative correlation between employment and wages observed in the data. That is, even when the only exogenous shock is to labor demand, the endogenous response of heterogeneous labor supply is sufficient to generate negative comovements between measured employment and wages in the aggregate.

In the Appendix, I additionally detail some of the disaggregated moments that the model produces. The model's occupation-level predicted changes in employment and wages between 2008 and 2009 have a correlation with the data changes in employment and wages of approximately 0.47 and 0.48, respectively.

The final four rows of the table illustrate the necessity of each ingredient of the model

[21]The data numbers consider year over year changes from March 2008 to March 2009, and reflect employment changes, rather than aggregate hours changes. Wages in the data are defined to be average hourly earnings in the BLS' Current Employment Statistics, deflated by the CPI-U.

Table 4: Wage and Employment Changes During Great Recession

Specification	Wage Change (1)	Employment Change (2)	Implied Elasticity (3)
Data	+2.7%	-8.8%	-3.8
Model: Calibrated	+2.5%	-10.4%	-4.6
Model: $\gamma_{jk} = 1 \, \forall j, k$	-3.5%	-1.0%	0.27
Model: Only Comparative Advantage	-1.6%	-2.1%	1.3
Model: Only Absolute Advantage	-1.3%	-7.9%	5.9
Model: No Home Sector	-2.7%	0.0%	0.0

Notes: Table reports the wage (column 1) and employment change between 2008 and 2009 in the data and a variety of model calibrations. Column 3 reports the ratio of employment changes to wage changes over this time period. Wages in the data correspond to average hourly earnings in the Current Employment Statistics, deflated by the CPI-U. The "Model: Calibrated" uses the skill distribution estimated in the CPS from 2002-2006. The model with only comparative advantage divides each worker type's skill vector by its mean so that all workers have the same average human capital. The model with only absolute advantage sets each worker type's skill vector to be a constant equal to its estimated mean. The final row reports estimates from a model in which there is no home sector.

to generate the strong selection patterns. Each row selectively removes one element of the model, and re-estimates the equilibrium response to the change in sectoral TFP between 2008 and 2009. The third row considers the case with no labor supply heterogeneity: every worker has one unit of human capital that they can supply to any occupation. In this model, real wages decline by 3.5% while employment falls by just 1%. Without skill heterogeneity, there is no scope for selection to buttress measured wages. As a result, the economy behaves as though a representative agent faced a decline in demand for their labor: both prices and quantities fall. The labor demand shock trades along the representative agent's relatively inelastic labor supply curve. Indeed, the implied elasticity of labor supply in this model is 0.27, roughly in the range of micro labor supply elasticities found in the literature surveyed by Chetty et al. (2011a). This highlights that shifting composition can rationalize the disconnect between estimated micro and macro labor supply elasticities.[22]

Both horizontal and vertical differentiation between workers are important. The fourth row of Table 4 considers the case in which there is no absolute advantage in the economy, but comparative advantage remains. To construct this counterfactual, I suppose that each worker type has the same mean γ_j, but maintains the estimated pattern of comparative advantage. That is, I construct a counterfactual Γ matrix by dividing each column of Table

[22]A common alternative employed to rationalize this disconnect is to assume that there is a difference between extensive and intensive margin elasticities of labor supply (e.g. Rogerson and Wallenius (2009) and Chetty et al. (2011b)).

2 by its mean. When faced with the decline in TFP between 2008 and 2009, this Γ matrix yields wage and employment declines of 1.6% and 2.1%, respectively. In the absence of absolute advantage, composition effects have limited strength, and therefore employment and wages continue to move together.

A worker fixed effect model also does not generate negative comovements between employment and wages. To construct this counterfactual, presented in the fifth row of the table, I assume that all worker types' vector of skills is a constant equal to the mean of their estimated γ_{jk} vector. Thus type 1 workers have 0.55 units of human capital, while type 8 workers have 7.54 units of human capital, but they may supply those units equally well across all occupations. In this model, employment falls by 7.9% while real wages fall by 1.3%. Here, there remains a great deal of selection: low skill workers leave the employed pool more than high skill workers in response to the negative demand shocks. This puts upward pressure on measured wages. However, these workers have labor supply which is relevant to all possible pursuits. Therefore, the negative demand shock to routine, construction and manufacturing jobs observed during the Great Recession exert a great deal of downward pressure on the price of, for instance, medical labor. This downward price pressure more than overcomes the selection force generated in the pure absolute advantage model, thereby preserving a positive covariance between employment and wages in the aggregate.

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

The final row of Table 4 considers a model in which workers do not have a home option and all workers are forced to work. In this case, the labor demand shock of course has no impact on employment levels. Therefore, there is no selection effect along the extensive margin. In addition, because all workers must work, labor is supplied inelastically. Removing the home option therefore generates wage declines of 2.7%.

This section shows that a model with workers of heterogeneous and imperfectly transferable skills, a non-employment option and firms employing heterogeneous occupations is able to replicate the aggregate employment and wage dynamics during the Great Recession.

Note that a similar pattern has played out during the recent COVID-19 pandemic. The lowest paid workers left the employed pool at a faster rate than those at the top, exerting strong upward pressure on average wages. These workers were concentrated in low-skill service occupations, and thus did not exert substantial downward pressure on the price of labor in pursuits requiring more specialized skills, such as healthcare provision. Indeed,

Cajner et al. (2020) show that while the selection-adjusted nominal wage fell by around 1% in the entertainment sector between March and June, the selection-adjusted wages of the healthcare sector continued to rise throughout 2020. Although the lack of high-quality industry production accounts for 2020 precludes equilibrium analysis of the coronavirus crisis, the behavior of sectoral wages during 2020 suggests that labor supply spillovers across sectors were limited despite large selection effects in the aggregate wage.

5 What Generated the Negative Wage-Employment Comovement During the Great Recession?

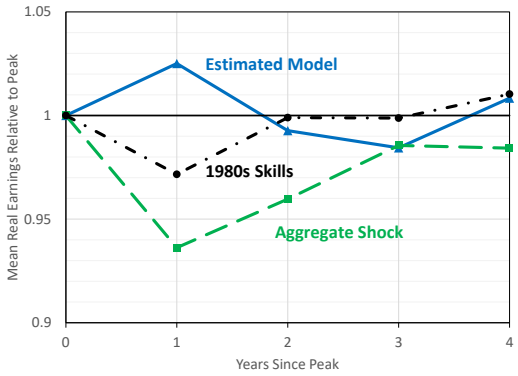
As the partial equilibrium exercises of Section 2.3.2 make clear, changes in the behavior of aggregate employment and wages can arise from two sources. First, the distribution of skills may have changed, thereby changing the patterns of selection and the degree of cross-occupation labor supply spillovers. Second, the distribution of shocks could conspire to change the relative prices of different occupations, thereby changing the allocation of workers to tasks. In this section, I explore the reasons why the behavior of aggregate employment and wages were different during the Great Recession than in prior recessions.

5.1 Changes in Labor Supply

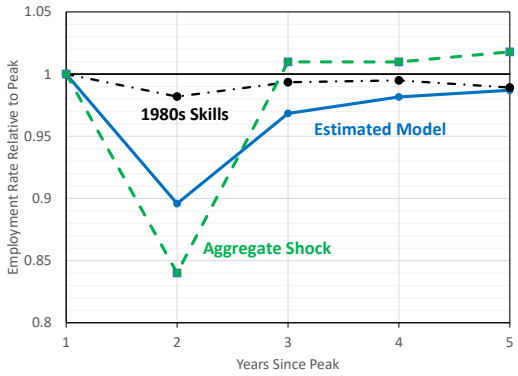
I begin by considering the importance of changes in the skill distribution. Figure 4 displays the predicted change in aggregate wages (Panel A) and employment (Panel B) relative to 2008 under the estimated model (blue solid line), and in two counterfactual economies. The black dash-dot line reports the evolution of employment and wages in a model in which the skill distribution of 1984-89 were subjected to the TFP shocks of the Great Recession. That is, it studies the impact of the Great Recession's labor demand shocks were they to occur 20 years earlier. This counterfactual exercise shows that real wages would have declined by approximately 3% with employment falling approximately 2% were the Great Recession to occur with the skill distribution of the 1980s. This stands in stark contrast to the estimated model which predicts rising wages. Thus, changes in the nature of labor supply were important to generate the wage and employment patterns observed recently. Below, I study the ways that labor supply has changed over this period.

To begin, consider the effect of unilateral increases in the price of each occupation w_k relative to the estimated equilibrium labor prices as of 2007. Increasing these prices induces flows out of non-employment. Using these flows, one can calculate the implied labor supply elasticity of non-employment to the price of each occupation as $d\ln(1-E)/d\ln w_k$. Figure

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



Panel A: Wages



Panel B: Employment

Notes: Figure reports the model-implied behavior of aggregate wages (Panel A) and employment (Panel B) under counterfactual skill distributions and sectoral shocks. The blue solid line reports the behavior of the estimated model around the Great Recession. The black dash-dot line shows the labor market evolution under a counterfactual in which the Great Recession sectoral shocks occurred with the skill distribution estimated during the 1984-89 period. The gray dashed line shows the movement of employment and wages in the case where the skill distribution is as estimated in 2002-06 and all sectors saw the same movement in TFP.

5 plots these implied elasticities for each occupation. The black bars plot the elasticities for the 1984-1989 period, while the gray bars plot the elasticities for the 2002-2006 period.[23]

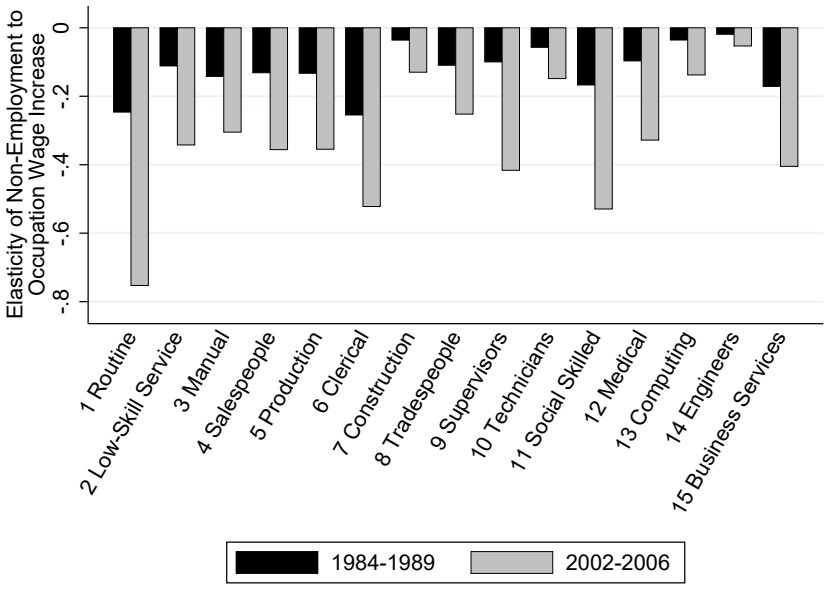
The figure shows substantial variation in the elasticity of non-employment to changes in occupation prices. There is no single "aggregate labor supply curve." Rather, the movements of employment and wages arise by aggregating movements along each of these primitive occupation-specific labor supply curves. 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 wage-employment schedules.

Low-wage occupations generally have higher non-employment elasticities than do high wage occupations, such as engineering. In both periods, the occupation cluster with the highest non-employment elasticity is routine occupations. This is intuitive, and reflects that many low wage workers move from routine occupations to non-employment and vice versa. The estimation reveals that workers most on the margin of non-employment have low average γs and 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 gen-

[23] Appendix Figure A4 additionally plots estimated own-price and cross-price elasticities of labor supply.

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 relative to the 2007 equilibrium price in each occupation.

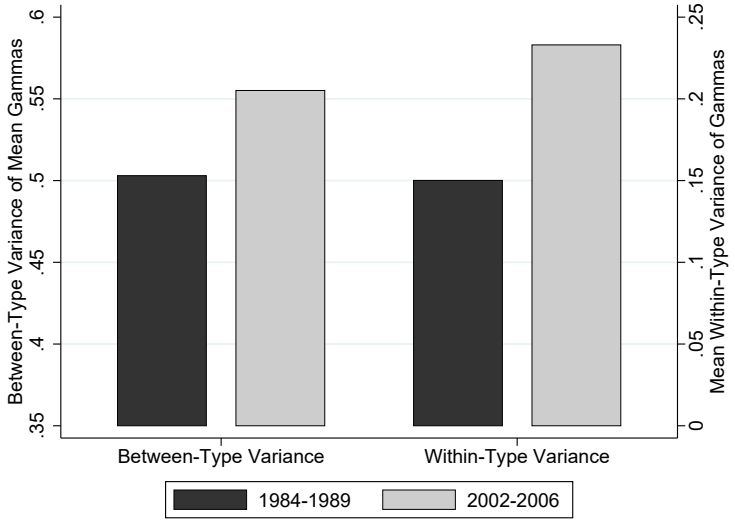
erally risen through time. Whereas the mean elasticity of non-employment to changes in the price of occupation-specific labor was -0.12 in 1984, that fell to -0.33 in 2002-2006. As a result, for any given change in the price of labor (or set of labor demand shocks), one might expect to see larger fluctuations in employment in the mid-2000s relative to the late 1980s.[24]

This change in the elasticity of labor supply primarily results from two forces. First, 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 labor prices in their occupation choice. This estimated decline stems in part from the decline in the average wage change of occupation switchers in the data, from 6.4% in the 1984-89 period to 4.5% from 2002-06.

Second, there have been changes in the distribution of absolute and comparative advantage, driven by changes in the Γ matrix. A natural measure of a worker's absolute advantage is the mean of her γ vector. 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

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

Figure 6: Absolute and Comparative Advantage: 1984-1989 and 2002-2006



Notes: Figure plots the estimated within and between type variance of skills in the economy, captured by the Γ matrix of Table A2 and 2. 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 17.

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 variance in skills may be decomposed into a within-type and an across-type variance. The across-type variance is informative about the difference mean skill for various workers. If this variance is high, then some workers have a substantially higher mean level of skill than others. 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

$$Var^{BTWN} := \sum_{j=1}^{J} m_j (\mathbb{E}_k[\gamma_{jk}] - 1)^2; \qquad Var^{WTHN} := \sum_{j=1}^{J} m_j Var_k(\gamma_{jk}), \tag{17}$$

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

Figure 6 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} has increased from 0.50 to 0.56, an increase of 10.4% in the 20 years leading up to the Great Recession. Since the mean of γ_{jk} is one within each occupation, these variances may be interpreted as the coefficient of variation of skills, squared.

There is an even larger increase in within-type variance, while the mean variance of the γ_{jk} vectors was 0.15 in the late 1980s, it was 0.23 prior to the 2008 recession, an increase of 55.2%. This suggests both that skills have become more specific and that the gap between the highest- and lowest-skill workers has grown.

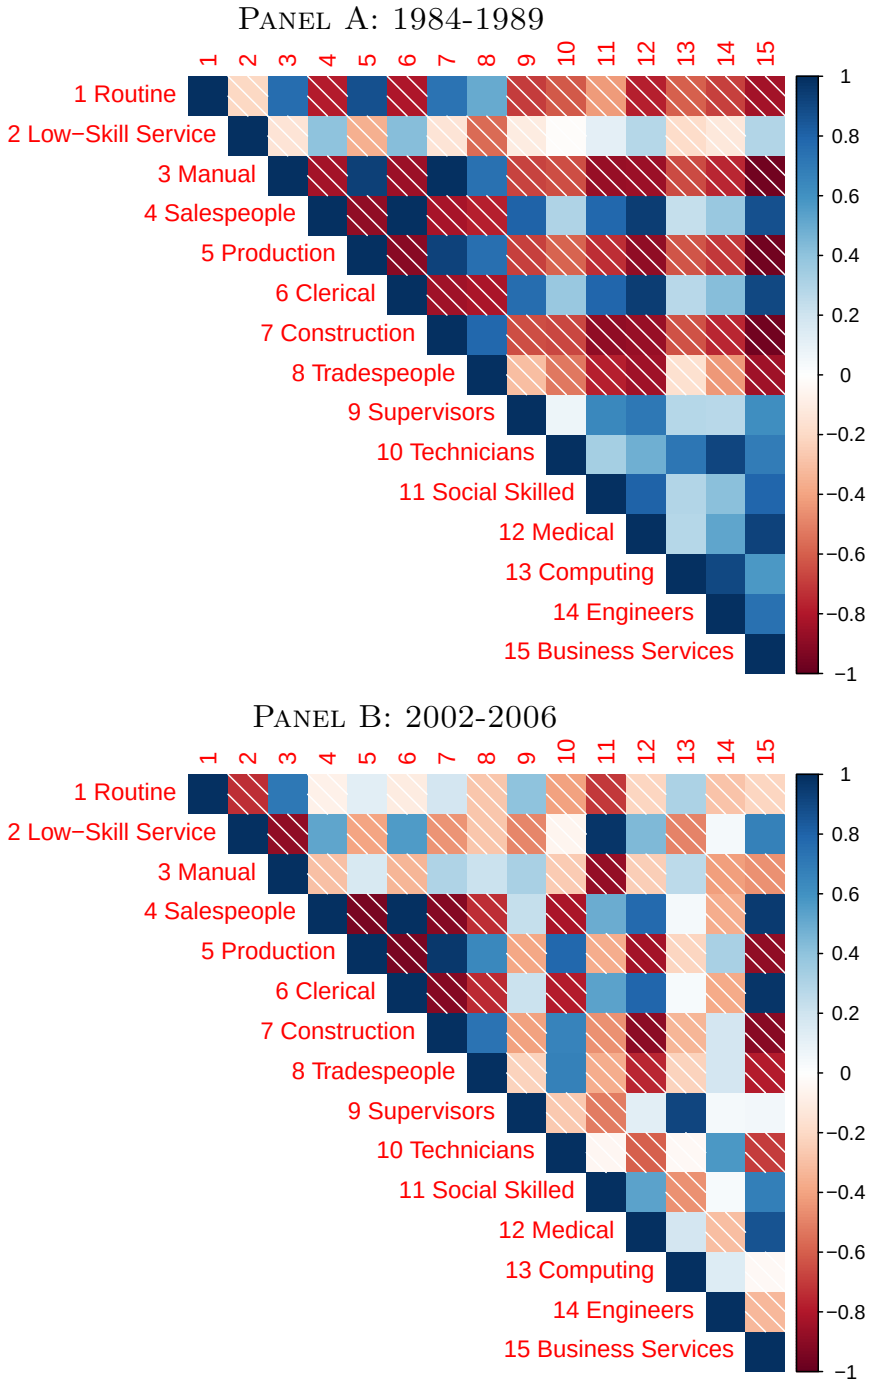
Heuristically, this result arises from two moments in the data. The increase in within-type variance owes to an increase in the variance of log earnings changes on occupation switches from 0.37 in the late 1980s to 0.52 in the mid 2000s. If this variance is high, the more one infers that individual workers' skills are better tailored to particular applications. Meanwhile, the increase in cross-type variance 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. Indeed, a regression of log earnings on occupation cluster fixed effects has an R^2 of 0.19 in the 2002-06 period, but just 0.16 in the 1984-89 period, suggesting increased earnings variance within occupations.

The degree to which skills are transferable affects aggregate wage dynamics by dictating the size of labor supply spillovers as workers reallocate from declining occupations to growing occupations.[25] The degree of skill transferability between any two occupations may be captured as the correlation of the row vectors of the Γ matrix. For instance, if the correlation between the Manual and Production occupations' γ vectors is high, it suggests that workers who have high skills in Manual occupations tend to also have skills in Production occupations. Put differently, workers who are good at manual labor, such as stock and material movers, may easily transfer their skills to production occupations to be serviceable welders or machinists.

Figure 7 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

[25] In Appendix H, I provide reduced form evidence that isolated labor demand shocks generate labor supply 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 7: Correlation of Occupation Skills, 1984-1989 and 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. Darker colors indicate that the magnitude of the correlation is closer to 1.

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 than they are for surgeons. Framed in this way, it is unsurprising that job-switchers out of medical professions tend to have similar skills to teachers and salespeople.

The comparison between 1984-1989 and 2002-2006 is also instructive. In the late 1980s, skills were highly transferable across high skill occupations, as represented by the large amount of blue squares in the bottom right corner of the correlogram. 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.

The results presented in this section suggest that a representative agent framework with interchangeable skills has become a worse approximation of reality. The gap between the most and least skilled workers in the economy has risen. Comparative advantage has similarly risen: workers have become more specialized over the last twenty years. Finally, the transferability of skills has declined amongst high-skill occupations, occupations employing manual labor, and occupations employing social skills. There may be many reasons for these changes, such as changes in education policy or a change in the task composition of occupations. Understanding the source of these changes is outside of the scope of this paper, but is fertile ground for future research. These changes have conspired to increase the primitive occupation-specific labor supply elasticities over time, thereby leading to larger employment declines and smaller wage declines in the face of labor demand shocks.

5.2 Labor Demand Shock Changes

In this section, I consider the importance of the set of labor demand shocks observed during the Great Recession in determining the observed decoupling of employment and wages.

The majority of industries receiving large negative shocks to selection-adjusted TFP primarily employ manual laborers – with the exception of the insurance industry, the largest declines in labor demand were concentrated in manufacturing, transportation, construction, and mining. This is highlighted by Figure 8, which shows the percentage change in selection-adjusted TFP between 2008 and 2009 (Panel A) and during the 1990-91 recession (Panel B) by the share of an industry's wage bill that accrues to manual occupations. Manual occupations are defined to be routine ($k = 1$), manual ($k = 3$), production ($k = 5$), construction ($k = 7$), and tradespeople ($k = 8$) occupations. Each dot is a different 3-digit NAICS sector, and its size is proportional to the value added share of that sector in the immediate pre-recession year. The figure shows that nearly all sectors mostly employing manual laborers saw large declines in labor demand during the Great Recession, but this was not true during the 1990-91 recession.

The green dashed lines in Figure 4 report a counterfactual evolution of aggregate employment and wages were there no sectoral heterogeneity in TFP shocks around the Great Recession. Specifically, it assumes that all sectors had declines in TFP of 5.9%: the average decline of TFP observed in the data, weighted by sectoral value added. Under this counterfactual set of labor demand shocks, the model implies that real wages would have declined by approximately 6%. The result of the concentration of shocks amongst sectors that principally employ manual occupations is that workers who have skills in manual occu-

Figure 8: Industry TFP shocks by Share of Wage Bill in Manual Occupations



Panel A: 2008-2009



Panel B: 1990-91

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

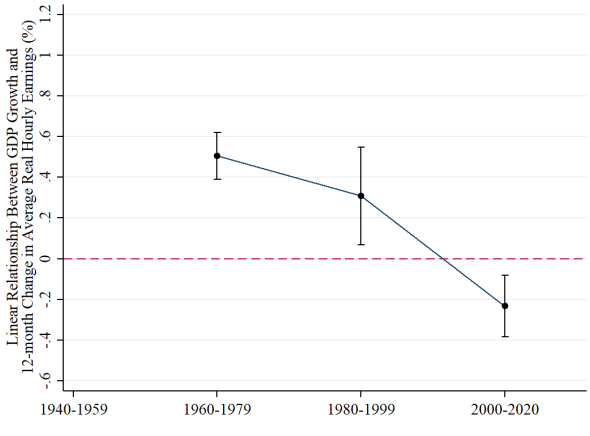
pations received a large negative demand shock for their skills, and had little scope to apply their human capital elsewhere. As a result, they left the employed pool, but exerted limited downward pressure on the price of labor in other occupations. Since these workers tend to be low skill (i.e. have low γ_{jk}s), this generated a large composition effect with limited spillovers to the rest of the economy. Both the unique nature of labor demand shocks and the shifting structure of labor supply contributed to the negative comovement between employment and wages observed during the Great Recession.

6 Reduced Form Composition-Correction

This section provides reduced form evidence that shifting composition effects account for nearly all of the change in the cyclicality of real wages. I first estimate the relationship between year-over-year GDP growth rates and year-over-year changes in aggregate wages or employment, allowing their cyclicality to differ every 20 years. Specifically, I estimate time series regressions of the form

$$\Delta \ln y_t = \alpha + \sum_p \beta_p \Delta \ln GDP_t \mathbf{1}\{t \in p\} + \epsilon_t \tag{18}$$

Figure 9: The Cyclicality of Labor Market Variables over Time



Panel A: Real Average Hourly Earnings



Panel B: Employment-to Population Ratio

Notes: Figure presents estimates of labor market cyclicality in the U.S. for twenty-year time periods going back to the 1940s. It shows the estimated linear relationship between year-over-year real GDP per capita growth rates and growth rates in real wages (Panel A) and employment (Panel B) by estimating β_p from Equation (A3) using OLS. Bars represent 90% confidence intervals, constructed with Newey-West standard errors allowing for up to 10 years of autocorrelation. Data sources: BLS, BEA, and FRED.

with OLS, where Δ is the year-over-year change operator, y_t is either the real average hourly earnings of production workers or the employment-to-population ratio, p indexes 20-year periods (e.g. 1940-1959, 1960-1979, etc.), GDP_t is quarterly real GDP per capita, and ϵ_t is an error term. The coefficient β_p represents the cyclicality of the variable y in period p.

The estimated cyclicalities of aggregate wages β_p are presented in Panel A of Figure 9, along with 90% confidence intervals calculated using Newey-West standard errors allowing for 10-year autocorrelation. The figure shows a stark decline in real aggregate wage cyclicality. While the elasticity of real aggregate wage growth to GDP growth was 0.50 (standard error: 0.07) in the 1960s and 70s, this fell to 0.31 (SE: 0.15) from 1980-1999. Strikingly, aggregate wages have become *countercyclical* since 2000, with an elasticity of wage growth to GDP growth of -0.23 (SE: 0.09). This trend looks set to continue through 2020, given the large increase in aggregate wages observed during the COVID-19 pandemic (Cajner et al., 2020).

Panel B of Figure 9 shows that the employment-to-population ratio has become more cyclical through time. While in the 1940s and 50s, the elasticity of employment growth to GDP growth was 0.46 (SE: 0.03), this had risen to 0.90 (SE: 0.22) by 2000-2020.[26] Restricting attention to prime age males strengthens this trend.

[26] All variables downloaded from FRED, maintained by the St. Louis Fed. Employment-to-population ratios are derived from the CPS, while average hourly earnings of production and non-supervisory employees come from the Current Employment Statistics, both produced by the BLS. GDP per capita taken from the BEA and deflated by the GDP price deflator. Wages deflated by core CPI.

How much of the decline in wage cyclicality can be attributable to changes in the nature of compositional shifts in the workforce over the business cycle? To address this question, I begin by following the existing selection-correction method of Solon et al. (1994), who assume the following statistical model for individual wages

$$\ln \omega_{it} = \alpha_i + \beta_1(U_t - \delta_1 - \delta_2 \cdot t - \delta_3 \cdot t^2) + \beta_2 \cdot t + \beta_3 \cdot t^2 + \beta_4 X_{it} + \beta_5 X_{it}^2 + \epsilon_{it} \tag{19}$$

where U_t is a contemporaneous cyclical indicator, 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 (19) in first differences, one controls for characteristics of a worker which are fixed over time, such as the worker's education, race, sex, and permanent unobserved ability. Therefore estimating

$$\Delta \ln \omega_{it} = \eta_0 + \beta_1 \Delta U_t + \eta_1 \cdot t + \eta_2 \cdot X_{it} + \nu_{it}, \tag{20}$$

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.

I estimate the same specification using microdata from the March Supplement of the Current Population Survey (CPS). As above, I allow the relationship between the aggregate cycle and selection-corrected wages to vary across 20-year periods. Specifically, I estimate

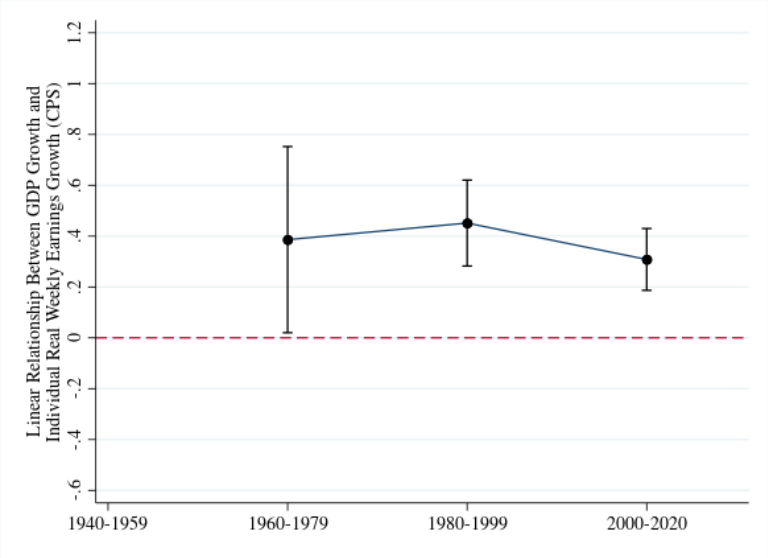
$$\Delta \ln \omega_{it} = \tilde{\eta}_0 + \tilde{\beta}_{1p} \Delta \ln GDP_t + \tilde{\eta}_1 \cdot t + \tilde{\eta}_2 X_{it} + \Delta \epsilon_{it}. \tag{21}$$

Figure 10 shows that the cyclical behavior of wages is much more stable after controlling for composition bias. It presents estimates of $\tilde{\beta}_{1p}$ and 90% confidence bands constructed using standard errors clustered at the year level. In the 1960s and 70s, the elasticity of individual real wage growth to GDP per capita growth was 0.39 (SE: 0.22) after controlling for selection.[27] This elasticity fell only slightly to 0.31 (SE: 0.07) by the 2000-2020 period. The relative stability of this coefficient through time, coupled with the sharp decline in aggregate wage cyclicality, provides reduced form evidence that cyclical compositional shifts in the workforce have become especially important over the last twenty years.

Note that the classic selection-correction equation (19) implicitly assumes that workers are vertically differentiated: some workers are high type, while others are low type. Horizontally differentiated workers may generate composition effects both by changing their

[27] Note that the estimates are noisy due to relatively small sample sizes in the early iterations of the CPS.

Figure 10: The Cyclicality of Real Weekly Earnings, Controlling for Composition Bias



Notes: Figure presents estimates of wage cyclicality in the U.S. for twenty-year time periods after controlling for composition bias. It shows the estimated linear relationship between year-over-year real GDP per capita growth rates and individual growth rates in real weekly earnings in the Current Population Survey by estimating $\tilde{beta}_1$ from Equation (21) using OLS. Bars represent 90% confidence intervals, constructed with standard errors clustered at the year level.

employment status and by switching between jobs for which they are well- or poorly-suited. Appendix B develops a more nuanced method to additionally correct for such "allocation effects."[28] Building off the model of Section 2, I show that considering the wage changes of occupation-stayers isolates movements in the price of labor under the assumption that workers' skills are fixed over short time periods. Doing so recovers pro-cyclical wages over the last 20 years. However, the estimated cyclicality under this augmented model is not notably different from that estimated from a Solon et al. (1994) correction, suggesting that the principal form of composition bias over this period is over the extensive margin of employment, rather than the allocation of workers across different productive tasks.

7 Conclusion

Composition effects confound attempts to measure the cyclicality of wages. This paper argues that such composition effects arise when shocks are concentrated in sectors which predominantly employ workers with a similar set of non-transferable skills. I propose a model in

[28]Note that the BLS's Employment Cost Index (ECI), which builds a wage index holding fixed the share of industry-occupation cells in employment, suffers from a similar bias by not accounting for which workers fill those jobs. Nevertheless, the ECI also shows mild declines in real wages during the Great Recession, further suggesting the importance of composition effects.

which workers differ in their skills for various occupations, and differentiated sectors combine each occupation with different weights. When labor demand declines in a particular sector, its workers reallocate to other activities. If those workers have highly transferable skills, they may 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, exerting a compositional force on the measured mean wage. These compositional forces would manifest as either labor wedge or aggregate labor supply shock in representative agent economies.

The model-relevant distribution of skills may be estimated using 2-period panel data on workers' occupation choices and earnings. Identification is achieved through the earnings changes of occupation-switchers, which reflect their relative skills in their source and destination occupation.

I apply the model to study the negative comovements of real wages and employment during the Great Recession. I estimate the distribution of skills using data from 2002-2006 from the linked CPS March Supplement. Feeding a sequence of realistic shocks to sectoral labor demand through the model replicates the large decline in employment from 2008 to 2009. More impressively, the model replicates the roughly 2% rise in real wages over this period. Although the only shock is to labor demand, the endogenous compositional shifts are sufficiently strong to cause employment and average wages to move in opposite directions.

This unusual negative comovement arose both because the distribution of skills in the economy have grown more diverse - the gap between low and high skill workers have grown, and workers have become more specialized - and because all of the sectors that traditionally employ the same low-skill manual laborers saw large declines in labor demand. The model implies that if the Great Recession had been an aggregate shock or had it occurred with the skill distribution of the late 1980s, employment and wages would have declined together.

Recognizing that the impact of sectoral shocks on aggregate employment and wages depends 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. 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.

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.

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

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

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.

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.

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).

Cajner, Tomaz, Leland D. Crane, Ryan A. Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz, "The U.S. Labor Market during the Beginning of the Pandemic Recession," *Becker-Friedman Institute Working Paper # 2020-58*, 2020.

Chang, Briana, "A Search Theory of Sectoral Reallocation," 2011.

Chang, Yongsung and Sun-Bin Kim, "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations," *American Economic Review*, 2007, *97* (5), 1939–1956.

_ , _ , Kyooho Kwon, and Richard Rogerson, "Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins," 2012.

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

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review Papers and Proceedings*, 2011, *101* (2), 471–75.

_ , John N Friedman, Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses and Micro v. Macro Labor Supply Elasticities," *Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

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.

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., 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.

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

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.

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

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.

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.

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*.

Kambourov, Gueorgui and Iourii Manovskii, "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 2009, *76*, 731–759.

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

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.

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

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).

Rogerson, Richard and Johanna Wallenius, "Micro and Macro Elasticities in a Life Cycle Model with Taxes," *Journal of Economic Theory*, 2009, *144*, 2277–2292.

_ , 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.

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

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

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.

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

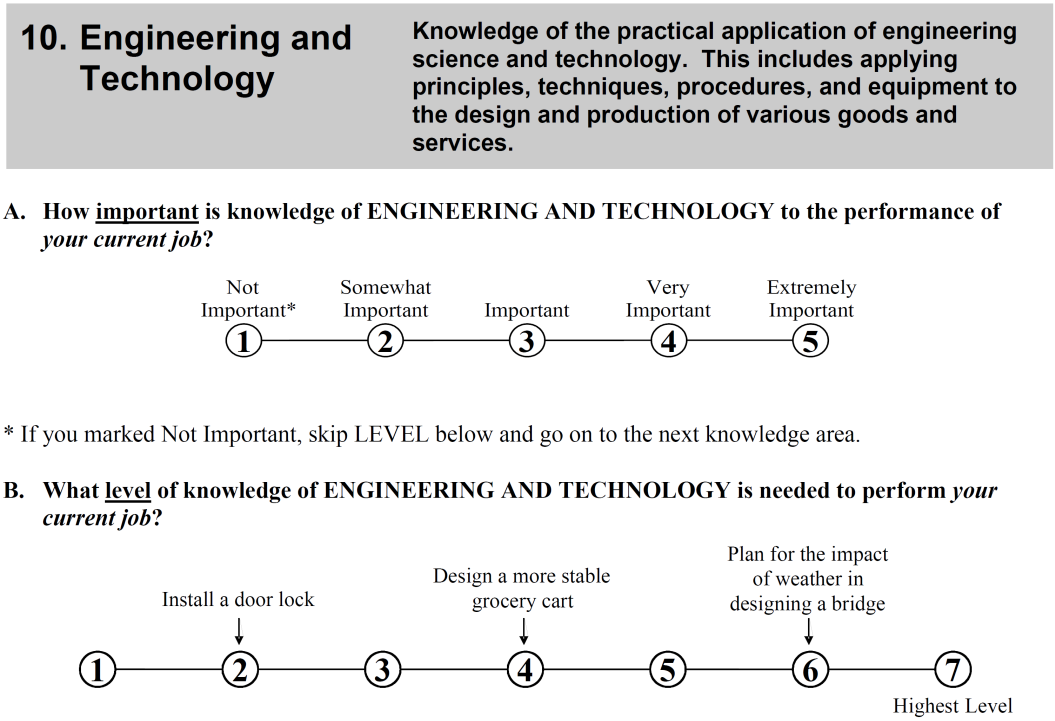


Table A2: Estimated Γ, m_j and ξ_k, 1984-1989 CPS

	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.

Table A3: 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	

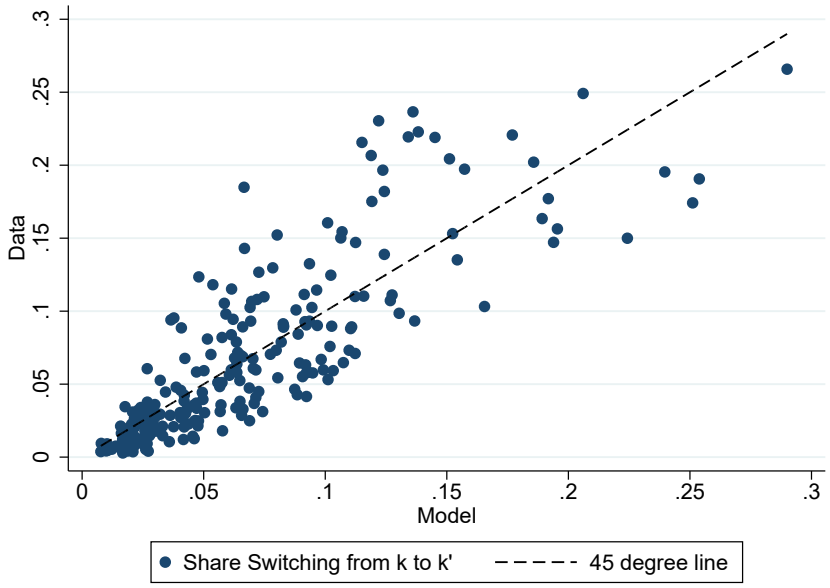
Notes: Table reports the in-sample fit of the estimated model for the period 2002-2006. Columns 1 and 2 report employment shares in each of the 15 occupations and the non-employment rate implied by the model and in the data, respectively. Columns 3 and 4 similarly report the mean log wage, while columns 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.

Table A4: 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	

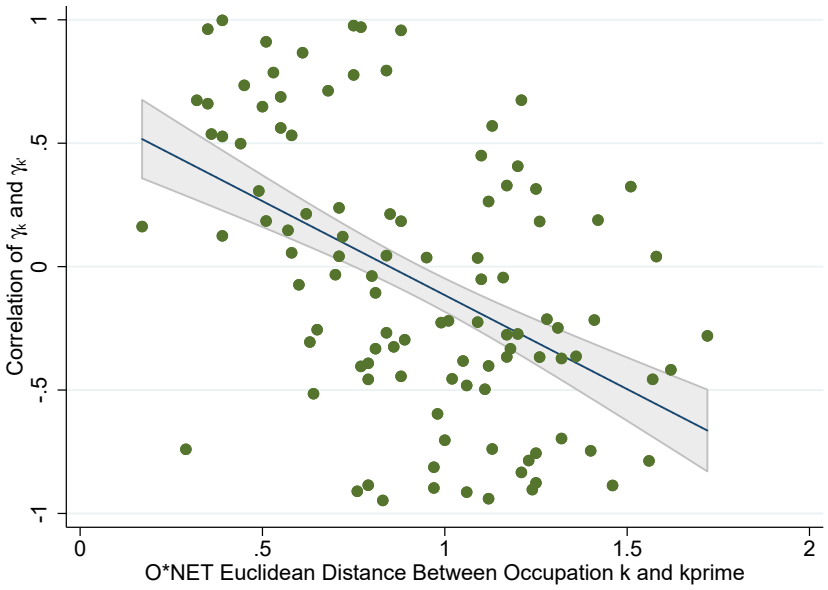
Notes: Table reports the in-sample fit of the estimated model for the period 1984-1989. Columns 1 and 2 report employment shares in each of the 15 occupations and the non-employment rate implied by the model and in the data, respectively. Columns 3 and 4 similarly report the mean log wage, while columns 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.

Figure A2: In-Sample Fit: Occupation Switching Patterns, 2002-2006



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

Figure A3: Correlation of Skill Relatedness in Γ with Euclidean Skill Distance in O*NET



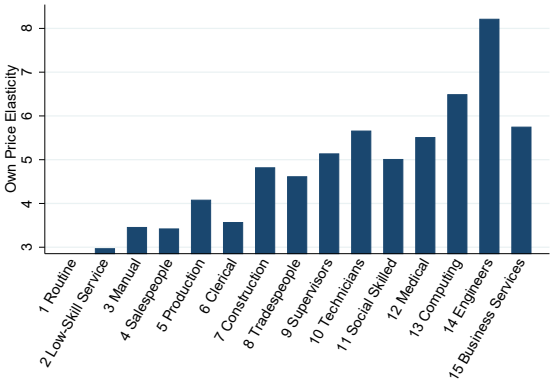
Notes: Figure compares the structurally-estimated skill transferability from the model to a common measure of skill relatedness from O*NET. Each dot corresponds to a pair of occupation clusters (k, k'). Occupations clustered by O*NET skill and knowledge vectors within terciles of the share with at least 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 7. Line of best fit reported, with shaded area representing 95% confidence interval using White heteroskedasticity robust standard errors.

Table A5: TFP Series: Annual Percentage Changes in the Raw BLS Multifactor Productivity Series Versus Series Adjusted for Human Capital Selection

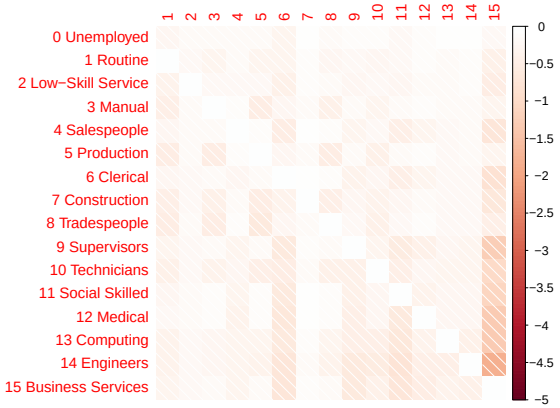
NAICS		1990-1991		2008-2009	
Code	Sector Title	BLS Raw	Adjusted	BLS Raw	Adjusted
211	Oil and gas extraction	0.9	-0.3	22.6	-3.8
212	Mining, except oil and gas	-0.0	-2.2	-5.9	-5.2
221	Utilities	-1.4	-0.7	3.8	-0.5
230-238	Construction	-0.5	-2.9	0.0	-6.0
311-312	Food and beverage and tobacco products	-0.8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

Notes: Data processing and other information services includes NAICS codes 516, 218, and 519. Aggregate TFP constructed as the mean of sector TFP series, weighted by value-added in each sector. 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 16.

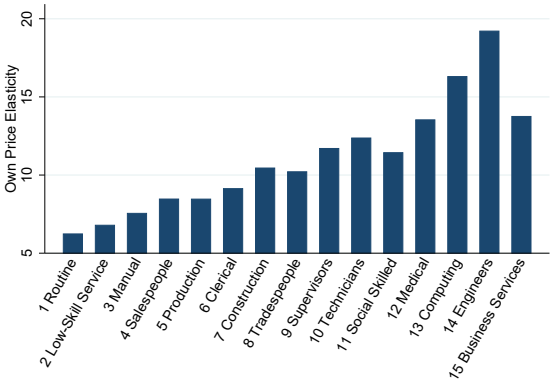
Figure A4: Estimated Own- and Cross-Price Elasticities of Labor Supply by Occupation



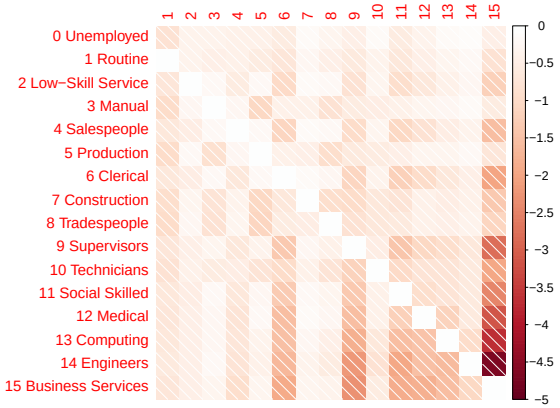
Panel A: Own Price Elasticities: 1984-1989



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 Model-implied selection correction

Classic methods to composition-adjust aggregate wages implicitly assume that workers are vertically differentiated: some workers are high type, while others are low type (Solon et al., 1994). My framework suggests a 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 for which they have different human capital levels. Fluctuations in workers' wages may partially 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 working in North Dakota's oil sector 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

$$\bar{\omega}_t = \sum_{k=1}^{K} \left(\frac{E_{kt}}{E_t} \right) \bar{\gamma}_{kt} w_{kt}. \tag{A1}$$

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 to estimate the distribution of γ_{jk} each year using the estimation approach of Section 3. I present the results of this approach in Appendix C. However, the data requirements for this method are large. This approach is thus limited 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 alternative 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_{ik} in occupation k. 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:

$$\Delta \ln \omega_{it} = (\ln \gamma_{ik} - \ln \gamma_{ik'}) + (\ln w_{kt} - \ln w_{k't-1}). \tag{A2}$$

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

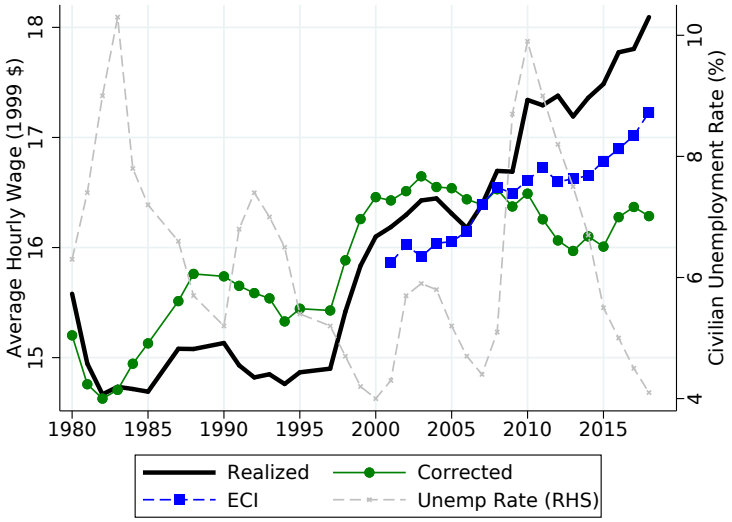
$$\Delta \ln \bar{\gamma}_{kt} = \Delta \ln \bar{\omega}_{kt} - \mathbb{E}[\Delta \ln w_{ikt} | k_t(i) = k_{t-1}(i) = k].$$

This approach relies on two 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, changes in workers' log wages must solely reflect changes in their marginal product or the price of labor.[29] 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. Indeed, 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 is able to account for cyclical changes in the allocation of workers to jobs in return.

To partially account for these concerns, 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

[29] If workers' wages reflect a *constant* markdown on their marginal product this decomposition will remain valid, as these markdowns will be differenced out.

Figure A5: 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 reports the movement of the BLS' Employment Cost Index (ECI), which fixes the employment shares of broad occupation-by-sector cells at their 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.

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 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 A5 presents various aggregate wage series implied by the CPS. The solid black line corresponds to the realized real average hourly earnings. 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 the behavior of the Employment Cost Index (ECI). The ECI is a data product produced by the BLS whose goal is to measure the cost of hiring a worker after controlling for selection. This wage series holds fixed the share of employment at the occupation-by-3 digit NAICS level. It uses 9 broad occupational groups, and coverage began in 2001. The data quality is very high: to produce the ECI, the BLS surveys a large number of establishments' administrative records. Thus the ECI effectively controls for the selection in *which jobs* grow or shrink through a cycle. Indeed, one can see that the wage series implied by the ECI does indeed fall slightly during the Great Recession, though this decline is muted. However, it does not account for *who works* in a given job. The allocation of workers to jobs may have large influence on the wages even if the share of jobs in a given occupation-sector cell is held fixed.

The green line with circle markers plots a wage series employing the novel selection correction described above. This isolates movements in aggregate wages arising solely from the wage movements of occupation stayers, controlling for worker age and occupation effects. The selection-corrected series diverges sharply from the uncorrected series after 2007. Whereas realized wages rose during 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 with regressions of the form

$$\Delta \ln \bar{\omega}_t = \beta_0 + \beta_1 \Delta U_t + \epsilon_t \tag{A3}$$

which is the aggregate version of equation (19), for the period 1980-2018. Table A6 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 procyclicality 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

Table A6: Cyclicality of Selection-Corrected Wage Series, 1980-2018

	Realized	Fixed Emp Shares	Selection Corrected	Solon et al.	SBP: Stayers
	(1)	(2)	(3)	(4)	(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 (A3). 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 (20), 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.

selection-corrected wages fall by a statistically significant 33 basis points. 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 (20) 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 sectors 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.

This need not have been the case. As highlighted above, the sectors principally affected during the Great Recession employed many low-skill manual laborers who were marginally attached to the workforce. As a result, the primary margin of selection was into and out of employment, rather than across different occupations within employment. If, for instance, there were a set of negative shocks to the data processing sector, we may have observed many data scientists reallocating themselves to become software engineers, even if they were not as well-suited to that task. In this case, the occupation-stayers correction would account for the changes in allocations *within* employment, whereas the SBP correction would not.

It is worth nothing that the procyclicality measured in both the aggregate wage series

and from the equation (20) regressions is lower than that reported in much of the literature. This is primarily due to the difference in periods. Re-estimating aggregate equation (A3) 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.

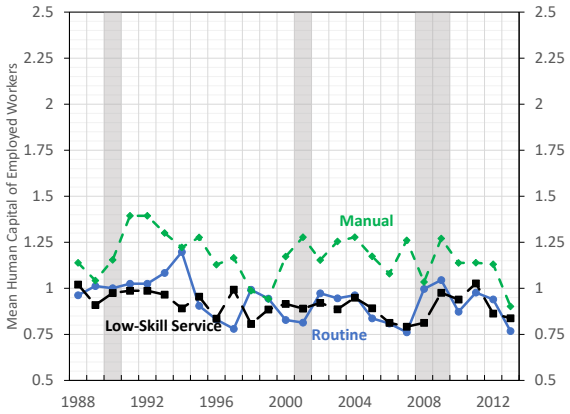
This is not a failure of the approach, however. As the model makes clear, one should not expect to have constant real wage cyclicality even after controlling for selection because the nature of labor supply spillovers change the extent to which the prices of occupational services vary over the cycle. 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 partly 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.

Appendix C Estimated Human Capital Selection in the Employed Pool of Workers

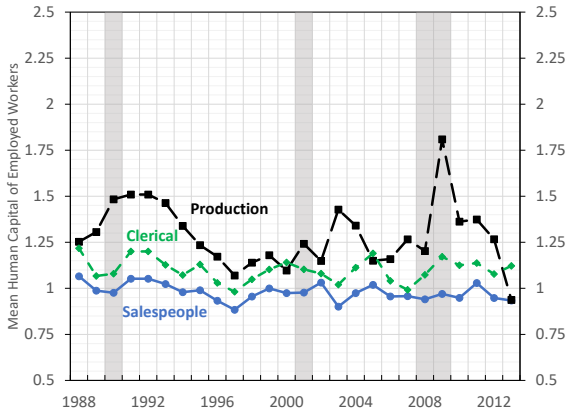
Figure A6 plots the time series of estimated mean human capital level of employed workers $\bar{\gamma}_{kt}$ for each of the 15 occupation clusters, as well as the aggregate mean human capital level of employed workers. To calculate these mean human capital levels, I re-estimate the maximum likelihood 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.

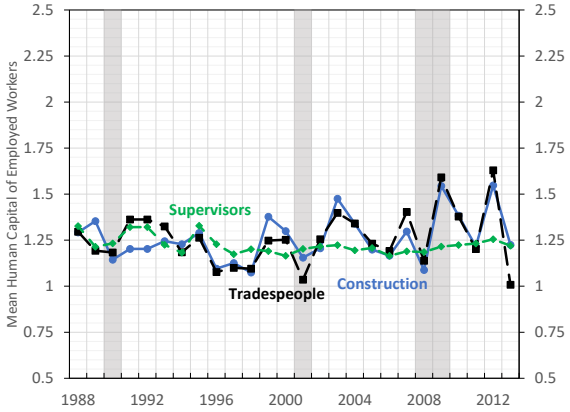
Figure A6: Time Series of Estimated Mean Human Capital of Employed Workers $\bar{\gamma}_{kt}$



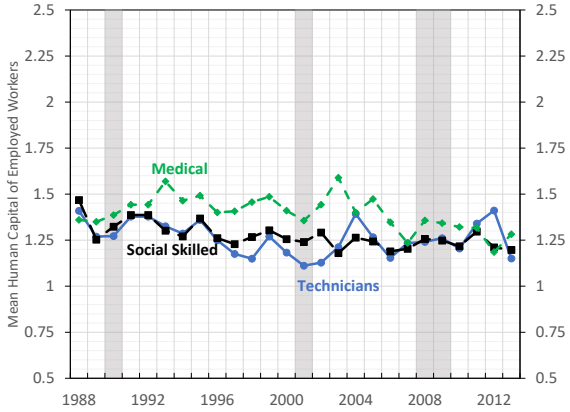
Panel A: Occupations 1 to 3



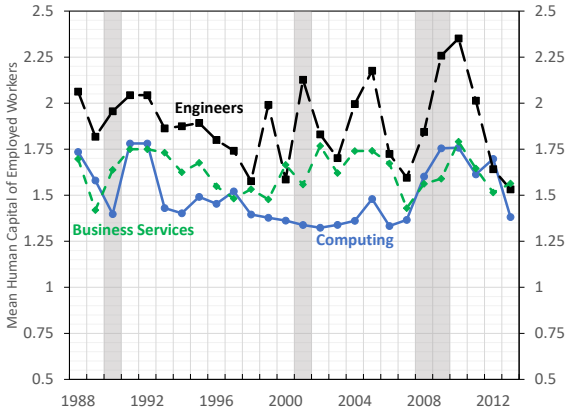
Panel B: Occupations 4 to 6



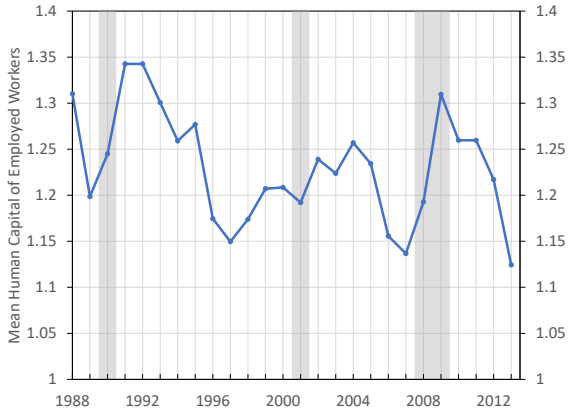
Panel C: Occupations 7 to 9



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.

In aggregate, the mean human capital of employed workers rose by 10% from 2008-2009, but only 4% in 1991. Given that wage rose by about 2% during the Great Recession and declined by about 2.3% in the 1990 recession, this change in the cyclicality of mean human capital can account for all 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 in 1990-91.

Appendix D 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:

$$Pr\left\{\omega_{i1} \leq \tilde{\omega}_1, \omega_{i2} \leq \tilde{\omega}_2 | k_1(i) = k_r, k_2(i) = \tilde{k}_{r'}, m_{i1} = 1\right\} = \sum_{j=1}^{J} p_{k_r, \tilde{k}_{r'}}(j) \Omega(\tilde{\omega}_1 | k_r, j) \Omega(\tilde{\omega}_2 | \tilde{k}_{r'}, j) \tag{A4}$$

where

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

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 A4 in matrix form, we have:

$$A(k_r, \tilde{k}_{r'}) = \Omega(k_r) D(k_r, \tilde{k}_{r'}) \Omega(\tilde{k}_{r'})^T \tag{A5}$$

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 A5, 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

$$q_k(j) = \frac{m_j\mathbb{P}_k(j)}{\sum_{j'=1}^{J} m_{j'}\mathbb{P}_k(j')} \tag{A6}$$

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 A6 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 E 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). 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.[30] 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 quar-

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

ters, retired workers, those serving in the armed forces, or employed workers with missing wage information.

I harmonize all sector codes to the 2010 NAICS coding using the crosswalks of provided by the Census bureau, and available at https://www.census.gov/topics/employment/sector-occupation/guidance/code-lists.html. 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-sector cells, which is an important input into the estimation of sector-level total factor productivity series. Using the CPS, I calculate the share of employees in each 3-digit NAICS code who belong to each of the K occupation clusters. I then interact this share with the sector-level employment provided by the Bureau of Economic Analysis (BEA) to construct an estimate for the total employment in each occupation-sector cell for every year.

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

Finally, Tables A7-A9 report 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

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

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

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

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

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

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

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

in the OES. The mean annual income in each SOC code according to the BLS is also listed.

Appendix F 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

$$l_i = \sum_{j=1}^{J} Pr\{j(i) = j\} \cdot Pr\{k_1(i) = k_1, k_2(i) = k_2 | j(i) = j\} \cdot \tag{A7}$$
$$\psi(\omega_{i1}|k_1(i) = k, j(i) = j) \cdot \psi(\omega_{i2}|k_2(i) = k, j(i) = j)$$

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 distributional assumptions on measurement error in wages, this distribution $\psi(\cdot)$ is log-normal with a different mean $\gamma_{jk}w_k$ 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

$$u_{ikt} = \underbrace{\gamma_{jk}w_{kt} + \xi_k}_{\bar{u}_{j(i)kt}} + \zeta_{ikt}$$

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, one may 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

$$G(\{\zeta_{ikt}, \zeta_{ikt+1}\}_k) = \exp\left(-\sum_{k=0}^{K}\left[e^{-\frac{\zeta_{ikt}}{\tilde{\rho}}} + e^{-\frac{\zeta_{ikt+1}}{\tilde{\rho}}}\right]^{\tilde{\rho}}\right) \tag{A8}$$

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 both period t and $t+1$. What's more, the expected utility of doing so may evolve between those choices if the price of labor moves. 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

$$\mathbb{P}_{kk'}(j) = \begin{cases} Pr\{k_t = k|j\} \cdot Pr\{k_{t+1} = k|k_t = k, j\} & \text{if } k = k' \\ Pr\{k_t = k|j\} \cdot Pr\{k_{t+1} \neq k|k_t = k, j\} \cdot Pr\{k_{t+1} = k'|j, k_{t+1} \neq k, k_t = k\} & \text{if } k \neq k' \end{cases} \tag{A9}$$

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

$$Pr\{k_t = k|j\} = \mathbb{P}_k(j) = \frac{\exp(\bar{u}_{jk}/\nu)}{\sum_{\tilde{k}=0}^{K} \exp(\bar{u}_{j\tilde{k}}/\nu)} \tag{A10}$$

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

$$Pr\{k_{t+1} = k'|j, k_{t+1} \neq k, k_t = k\} = \frac{\exp(\bar{u}_{jk'}/\nu)}{\sum_{\tilde{k}\neq k} \exp(\bar{u}_{j\tilde{k}}/\nu)} \tag{A11}$$

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

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

Let $M_{t+1} := \max_{\tilde{k}\neq k}(\bar{u}_{j\tilde{k}t+1} - \bar{u}_{jkt+1}) + \zeta_{i\tilde{k}t+1}$ 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

$$Pr\{k_{t+1} \neq k|j, k_t = k\} = \int G(M_{t+1}|j, k_t = k)\psi_{jk}(M_{t+1})dM \tag{A12}$$

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

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

As a result, this condition is equivalent to a restriction on ζ_{ikt}. Therefore, we may rewrite equation A12 as

$$Pr\{k_{t+1} \neq k|j, k_t = k\} = \int \left(\int G_{\zeta_{ikt+1}|\zeta_{ikt}}(M_{t+1}|\zeta_{ikt})\psi_{jk}(M_{t+1})dM_{t+1} \right) \varphi_{jk}(\zeta_{ikt}|j, k_t = k)d\zeta_{ikt} \tag{A13}$$

for $\varphi_{jk}(\zeta_{ikt}|j, k_t = k)$ 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 (A8) yields that this conditional CDF is given by

$$G_{\zeta_{ikt+1}|\zeta_{ikt}}(M|\zeta_{ikt}) = \frac{\exp\left(-\left[e^{-\frac{\zeta_{ikt}}{\nu\tilde{\rho}}} + e^{-\frac{M}{\nu\tilde{\rho}}} \right]^{\tilde{\rho}} - \frac{\zeta_{ikt}}{\nu\tilde{\rho}} \right) \left[e^{-\frac{\zeta_{ikt}}{\nu\tilde{\rho}}} + e^{-\frac{M}{\nu\tilde{\rho}}} \right]^{\tilde{\rho}-1}}{\frac{1}{\nu}\exp\left(-\left[\frac{\zeta_{ikt}}{\nu} + e^{-\frac{\zeta_{ikt}}{\nu}} \right] \right)} \tag{A14}$$

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

$$\begin{aligned} Pr\{\zeta_{ikt} \leq x|j, k_t = k\} &= Pr\{\zeta_{ikt} \leq x|\zeta_{ikt} \geq \overbrace{\max_{k'} \bar{u}_{jk't} - \bar{u}_{jkt} + \zeta_{ik't}}^{M_t}\} \\ &= \int_{-\infty}^{x} [G(x) - G(M_t)]\psi(M_t)dM_t \end{aligned}$$

Using Leibniz's rule for differentiating under the integral gives the expression for the density of ζ_{ikt} given a choice k:

$$\begin{aligned} \varphi(x) &= \frac{dPr\{\zeta_{ikt} \leq x|k_t = k\}}{dx} \\ &= \int_{-\infty}^{x} g(x)\psi(M_t)dM_t \\ &= g(x)\Psi(x) \end{aligned} \tag{A15}$$

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 $\ln\left(\sum_{\tilde{k}\neq k}\exp(\bar{u}_{j\tilde{k}t}-\bar{u}_{jkt})\right)$. One may therefore substitute equations A15 and A14 into equation A13 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 A9.

Thus, the log-likelihood of the data may be computed by summing over the log of the individual likelihoods expressed in A7. Doing so yields, for $\theta := (\{m_j, \xi_k, \gamma_{jk}w_{kt}, \sigma_{jk}\}_{j,k}, \nu, \rho)$ the complete vector of labor supply parameters:

$$\mathcal{L}(\theta) = \sum_i \ln\left(\sum_{j=1}^{J} \underbrace{m_j Pr\{k_1 = k_1(i)|j;\theta\}\cdot Pr\{k_2 = k_2(i)|k_1 = k_1(i), j;\theta\}\psi(\omega_{i1}|\theta)\psi(\omega_{i2}|\theta)}_{l_{ij}}\right) \tag{A16}$$

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 G 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

$$l_{sk} = \frac{p_s x_s z_s \alpha_{sk}\left(\prod_{k'=1}^{K} l_{sk'}^{\alpha_{sk'}}\right)^{x_s}}{w_k} \tag{A17}$$

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

$$l_{sk'} = l_{sk}\left(\frac{\alpha_{sk'}}{\alpha_{sk}}\frac{w_k}{w_{k'}}\right) \tag{A18}$$

Substitute this into equation (A17) to arrive at

$$l_{sk}^{1-x_s} = p_s x_s z_s \left(\frac{\alpha_{sk}}{w_k}\right)^{1-x_s}\left(\prod_{k'=1}^{K}\left(\frac{\alpha_{sk'}}{w_{k'}}\right)^{\alpha_{sk'}}\right)^{x_s} \tag{A19}$$

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

$$l_{sk}^{1-x_s} = \frac{(Y)^{\frac{1}{\eta}} x_s z_s \left(\frac{\alpha_{sk}}{w_k}\right)^{1-x_s}}{M_s^{x_s} y_s^{\frac{1}{\eta}}} \tag{A20}$$

Plugging equation (A18) into the production function for sector s reveals that

$$y_s = z_s\left(\frac{M_s\alpha_{sk}}{w_k}\right)^{-x_s} l_{sk}^{x_s} \tag{A21}$$

which we may then substitute into the amended first order condition A20

$$l_{sk}^{\eta-x_s(\eta-1)} = Y x_s^{\eta} z_s^{\eta-1} M_s^{-x_s(\eta-1)}\left(\frac{\alpha_{sk}}{w_k}\right)^{\eta-x_s(\eta-1)} \tag{A22}$$

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

$$l_{s'k} = l_{sk}^{\frac{\nu_s}{\nu_{s'}}} \underbrace{\left(\frac{\alpha_{s'k}}{w_k}\right)\left(\frac{w_k}{\alpha_{sk}}\right)^{\frac{\nu_s}{\nu_{s'}}}\left[\left(\frac{x_{s'}}{x_s}\right)^{\eta}\left(\frac{z_{s'}}{z_s}\right)^{\eta-1}\left(\frac{M_s^{x_s}}{M_{s'}^{x_s'}}\right)^{\eta-1}\right]^{\frac{1}{\nu_{s'}}}}_{:=\psi_{s',s}} \tag{A23}$$

As a result, A21 implies that the equilibrium output in sector s' is given by

$$y_{s'} = z_{s'}\left(\frac{M_{s'}\alpha_{s'k}}{w_k}\right)^{-x_{s'}} \psi_{s',s}^{x_{s'}} l_{sk}^{\frac{x_{s'}\nu_s}{\nu_{s'}}} \tag{A24}$$

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

$$Y(l_{sk}) = \left(\sum_{s'=1}^{S}\left[z_{s'}\left(\frac{M_{s'}\alpha_{s'k}}{w_k}\right)^{-x_{s'}} \psi_{s',s}^{x_{s'}} l_{sk}^{\frac{x_{s'}\nu_s}{\nu_{s'}}}\right]^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}} \tag{A25}$$

Finally, we may plug this into equation (A22) 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 (A18) and (A23) 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:

$$L_k^D(\mathbf{w}|\mathbf{z}) = \sum_{s=1}^{S} l_{sk}(\mathbf{w}|\mathbf{z}) \tag{A26}$$

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

$$L_k(\mathbf{w}) = \sum_{j=1}^{J} m_j \gamma_{jk} \underbrace{\left(\frac{\exp(\bar{u}_{jk}/\nu)}{\sum_{k'=0}^{K}\exp(\bar{u}_{jk'}/\nu)}\right)}_{\mathbb{P}_k(j)} \tag{A27}$$

One may solve for equilibrium by equating labor demand for occupation k, given by equation A26, with the labor supply for this occupation, given by equation A27. Note that since workers do not have preferences over which sector to work for, and because workers

are perfect substitutes within an occupation conditional on their units of effective labor, the law of one price will hold within each occupation. These occupation prices will determine the quantities of effective labor in each occupation employed by each sector. Furthermore, Walras' Law implies that equating the labor demand and labor supply in each occupation will imply that final goods clearing is also satisfied. That is, total income, given by

$$C = \underbrace{\sum_{j=1}^{J} m_j \sum_{k=1}^{K} \gamma_{jk} w_k \mathbb{P}_k(j)}_{I} + \underbrace{\sum_{n=1}^{N} (1 - x_s) p_s y_s}_{\Pi} \tag{A28}$$

will equal aggregate output given by equation A25.

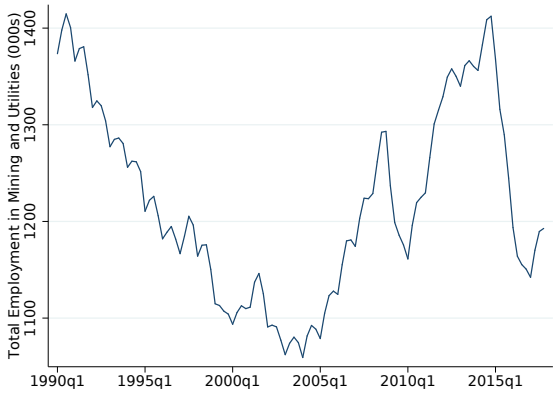
Note that the structure of the model implies that one need only solve for the K occupation prices in order to characterize the equilibrium. For this reason, one can consider sectors at a fine level of aggregation without adding substantial computational burden.

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 (A27) and labor demand (A26) subject to a choice of wage vector $\mathbf{w}$. I additionally include the squared difference between aggregate output and consumption as an equilibrium condition, as doing so improves performance of the optimizer. The ISRES routine is a semi-global optimization method put forward by Runarsson and Yao (2005). Arnoud et al. (2019) finds that ISRES performs well in many economic applications. I supply 30 starting values to the optimizer.

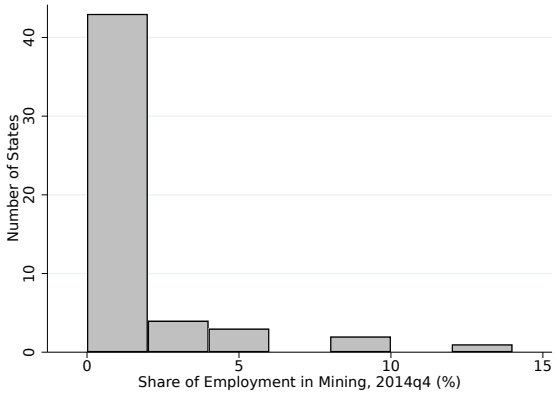
Appendix H Reduced Form Evidence for Labor Supply Spillovers

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 A7: Shock to Mining Employment



Panel A: Aggregate Mining Employment



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

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

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 A7. 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 generates a larger labor demand shock in states like Texas than it did in Florida, providing a laboratory to study the impact of a sectoral decline on related sectors. Let λ_r^{MINING} be the share of region r's employment that is in mining as of the fourth quarter of 2014, and let $\Delta \ln E_{MINING,-r}$ denote the percent change in mining employment in all states other than r between the fourth quarter of 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. Indeed, this is one of the primary motivations for the model presented in the main text. The goal here is to provide model-free reduced form evidence of cross-sector labor supply spillovers that is mediated through skill transferability.

Finally, I aggregate to sector-level skill vectors by combining the O*NET occupation-level data. Specifically, let χ_{sk} be the share of 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 sector s to be the weighted average of $h_m(k)$, where the weights are the shares of s's employment in occupation $k : \chi_{sk}$. That is,

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

One can interpret this measure to be the expected skill vector a worker would require in sector s if one were to randomly samply workers in that sector. Given these skill vectors, one 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 sector-state level for every quarter back to 1975. I restrict attention to the set of tradable 3-digit NAICS sectors which have skills which are highly related or unrelated to mining. Sectors with highly related skills are in the bottom quartile of skill distance to mining – $d(s, Mining)$ is small – while those with unrelated skills are in the top quartile of skill distance. Restricting attention to tradable sectors isolates local labor supply effects by abstracting from movements in local labor demand resulting from the decline in

Table A10: 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 A29. Sectors with related skills are defined to be those sectors in the bottom quartile of skill distance with Mining sectors. Only tradable sectors in the top and bottom quartile of skill distance included. Standard errors clustered at 3-digit NAICS sector code level reported in parentheses.

mining. I estimate the following regression at the region-sector level

$$\Delta \ln y_{sr} = \alpha \cdot \mathbf{1}\{s \text{ is Related}\} + \eta \cdot \sigma_r + \beta \mathbf{1}\{s \text{ is Related}\} \cdot \sigma_r + \epsilon_{sr} \tag{A29}$$

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 A10. Columns 2 and 4 control for state-sector-specific trends (i.e. long run growth between 1990 and 2014), while columns 1 and 3 do not. The table shows that sectors with skills related to mining experienced larger declines in wages and increases in employment, relative to sectors with unrelated skills, in states which had large pre-existing mining shares, 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 is associated with an increase of 4 percentage points in the share of workers employed in sectors with skills related to mining, relative to sectors with unrelated skills. This is coupled with a relative decline in log wages of approximately 4% in these sectors. These patterns are consistent with positive labor supply spillovers from the mining sectors to other sectors most related to mining.