Monopsony and Outside Options

Gregor Schubert, Anna Stansbury, and Bledi Taska *

March 2020

Abstract

In imperfectly competitive labor markets, wages depend in part on the availability and quality of workers' outside options. How can we measure these outside options? How much do they matter for workers' wages? We measure outside options in two groups: jobs within workers' own occupation, which we proxy for with local labor market concentration, and jobs outside the occupation, for which we construct a new index using occupational mobility data from 16 million U.S. resumes. Using a novel granular IV for changes in employer concentration, and a Bartik-type instrument for changes in the quality of local outside-occupation job options, we find that 21% of within-occupation regional wage variation in the U.S. can be explained by differences in local employer concentration, and a further 13% can be attributed to differences in outside-occupation job options. The two interact: the effect of concentration on wages is twice as high for occupations with the lowest outward mobility. Through the lens of a Nash bargaining model, our results suggest that a $1 increase in the value of outside options translates into a $0.20-$0.40 increase in wages.

*Schubert: Harvard University, gschubert@g.harvard.edu. Stansbury: Harvard University, annastansbury@g.harvard.edu. Taska: Burning Glass Technologies, btaska@burning-glass.com. Previous versions of this paper have been circulated under the title "Getting labor markets right: occupational mobility, outside options, and labor market concentration." The authors thank Gabriel Chodorow-Reich, Karen Dynan, Ed Glaeser, Xavier Gabaix, Larry Katz, Maya Sen, Betsey Stevenson, and Larry Summers for detailed comments and advice, and Justin Bloesch, John Coglianese, Oren Danieli, David Deming, Martin Feldstein, Claudia Goldin, Emma Harrington, Simon Jaeger, Max Kasy, Bill Kerr, Robin Lee, Ioana Marinescu, Jeff Miron, Nancy Rose, Isaac Sorkin, Elizabeth Weber Handwerker, Ron Yang and participants of the briq Workshop on Firms, Jobs and Inequality, the 2019 Federal Reserve System Community Development Research Conference, the IZA Summer School on Labor Economics 2019, the IZA/CAIS Workshop on Matching in Online Labor Markets, the Urban Economic Association Meeting 2019, the Wharton People and Organizations Conference, and the Harvard Labor lunch, Labor breakfast, Industrial Organization lunch, Macro lunch, and Multidisciplinary Seminar on Inequality and Social Policy for helpful comments and suggestions. This research was supported by the Washington Center for Equitable Growth (Schubert & Stansbury) and the James M. and Kathleen D. Stone PhD Scholarship in Inequality and Wealth Concentration (Stansbury).

1 Introduction

In imperfectly competitive labor markets, workers do not necessarily earn their marginal product. Instead, the worker's wage will depend to some extent on the degree to which they are able to move on to other opportunities if the firm's offer is not sufficiently generous. This is true both in neoclassical or dynamic monopsony models of the labor market, where firm monopsony power is determined by the elasticity of labor supply to the firm, and in search and matching models featuring wage bargaining, where better worker opportunities give workers a more valuable outside option, leading to higher wages (Boal and Ransom, 1997; Pissarides, 2000; Manning, 2003; Ashenfelter et al., 2013). Measuring worker outside options and estimating their effect on wages therefore plays an important role in assessing worker bargaining power and quantifying employers' monopsony power to reduce wages below workers' marginal product.

In this paper, we propose a novel strategy to measure outside options and estimate their effect on wages on a systematic basis across almost all occupations and metropolitan areas in the U.S. We first group workers' outside options into two categories: job options within workers' own occupation, and job options outside the occupation. To measure the value of options *outside* workers' own occupation, we construct a 'revealed' measure of workers' labor markets, using empirical job transition data (from a new data set of 16 million U.S. resumes) to infer the relevance of jobs in other local occupations. To measure worker's options *within* their current occupation, we use an index of local labor market concentration derived from online vacancy postings – noting that in a search model, the availability of within-occupation job options can be a function of the concentration of local employers in that occupation. We then estimate the effect of changes in these outside options on wages. Using a novel granular IV identification strategy for the effects of employer concentration on wages, and using Bartik-style shocks to wages in workers' outside-occupation job options, we show that both types of outside options matter for wages. Within a given occupation, 21% of regional wage variation can be attributed to differences in local labor market concentration across cities, and a further 16% can be attributed to differences in the quality of outside-occupation job options – as opposed to, for example, differential productivity in different cities. We also find that these different components of outside options interact: the effect of employer concentration on wages is three times higher for workers in occupations with low outward mobility than in occupations with high outward mobility.

When we interpret these estimates through the lens of a simple search model with wage bargaining, they imply that a $1 increase in the value of workers' outside options translates into a wage increase of $0.23 to $0.37. In our model, this would imply workers receive roughly 2/3 of rents or quasi rents – similar to their share of overall compensation, and a greater degree of rent-sharing within firms than has been documented by other studies in the literature (Card et al., 2016). Alternatively, this estimate may suggest that job options outside the firm are not as relevant in the wage bargain as many models would suggest (or that there are additional dimensions of outside options that are not captured by our measures of local occupational wages and employer concentration).

Overview of our analysis: We start our analysis by developing a simple labor market search model with wage bargaining, in which workers' bargained wages depend on the expected quality of their best outside job option – a function of both the job options within their current occupation (which depends on labor market concentration) and the feasible job options in other occupations. The model suggests (1) that employer concentration in workers' own occupation lowers wages, (2) that higher wages or more job options in other occupations which are feasible outside options increase wages, and (3) that the two interact with each other: more or better local outside-occupation options mitigate the effect of employer concentration within workers' own occupation.

Based on this model, we construct an index of the average value of workers' outside-occupation job options within their local area. This index is constructed as the weighted average of local wages in all occupations except the worker's own, where each weight is the product of two factors: (i) the national average empirical transition share from the worker's starting occupation to each destination occupation, and (ii) the local employment share in that destination occupation, relative to the national average. We use the occupation-to-occupation transition share as a proxy for the likelihood that a worker's best job option outside her own occupation is in each of these other occupations. We use the local relative employment share as a proxy for the local availability of job options in each destination occupation.

Why use empirical occupation-to-occupation transitions to proxy for workers' outside options? Occupational transitions are a simple, non-parametric way to identify workers' "revealed" labor markets, capturing both the feasibility and desirability of jobs in different occupations. We construct our occupation-to-occupation transition data from a new and unique data set of 16 million resumes, with over 80 million job observations over 2002–2018, collected by Burning Glass Technologies. The large sample size – an order of magnitude more than other data sources – enables us to reliably estimate average transition shares between a large proportion of all pairs of 6-digit occupations in the U.S.[1] Using this data we document that occupational mobility is high, highly heterogeneous across occupations, asymmetric, and poorly captured by aggregating up the occupational hierarchy, suggesting that existing occupational classifications are not sufficient to proxy for workers' outside options; and that occupational transitions capture underlying similarity between occupations in terms of task requirements, leadership responsibilities, and amenities.

We then use our outside-occupation option index, as well as measures of employer concentration at the occupation-city-year level, to study the effect of outside options on wages.

We construct our outside-occupation index for a panel of U.S. occupations and cities over 1999–2016. To identify causal effects, we focus on quasi-exogenous shocks to the value of workers' outside-occupation job options: we instrument for demand shocks to workers' outside-option occupations using the national leave-one-out mean wage in each of those occupations, and proxy for local exposure to these demand shocks using the local employment share in each of these occupations relative to the national average. We find a positive and significant effect of an increase in the value of outside-occupation options: within a given occupation and year, across cities, moving from the 25th to the 75th percentile value of outside-occupation options is associated with 3.0 log points higher wages. This in turn suggests that 13% of the interquartile variation in wages within a given occupation across U.S. cities can be explained by the relative availability of local

[1] Based on Standard Occupational Classification codes.

outside-occupation job options.

We construct our measure of employer concentration – the Herfindahl-Hirschman Index (HHI) of online vacancy postings, from Burning Glass Technologies – for occupation-by-city cells over 2013-2016. To identify a causal effect of concentration on wages, we develop a new instrumental variable strategy based on the "granular IV" approach developed by Gabaix and Koijen (2020). We exploit the fact that different occupation-city cells vary in their exposure to the national growth of certain employers relative to others: when a large employer grows nationwide, this will increase employer concentration by more in cities which already had a large presence of that employer. We use this distinction to construct exogenous shocks to local concentration that are plausibly orthogonal to local changes in productivity. We estimate a negative and significant effect of local employer concentration on wages: within a given occupation and year, across cities, moving from the 25th to the 75th percentile HHI is associated with 4.8 log points lower wages, suggesting that 21% of the interquartile variation in wages within a given occupation across U.S. cities could be explained by differences in local labor market concentration.

In addition, we show that the effect of labor market concentration on wages is more than three times higher for occupations with low outward mobility than for occupations with high outward mobility, and that that regressions of wages on employer HHI suffer from omitted variable bias if the availability and quality of outside-occupation job options is not included in the analysis (with an upward bias in the size of the coefficient of about 20-40%).[2] A simple regression of wages on within-occupation HHI alone is therefore not appropriate to identify employer concentration effects in many local labor markets—it should be considered in conjunction with the quality and availability of local outside-occupation job options.

Finally, we use our coefficient estimates from regressions of wages on HHI and outside-occupation options to back out a range of estimates for workers' bargaining power – β – that would be consistent with our Nash bargaining model of wage determination. Our coefficients suggest a β of 0.63–0.76, which would suggest that workers receive 63-76% of the rents or quasi-rents generated by a match. This is

[2]This occurs because workers with few local employers in their own occupation (and so a high HHI) also tend to have worse local outside-occupation job options as measured by our index.

substantially larger than the β implied by studies which estimate rent-sharing elasticities based on firm-level changes in productivity or profits (as reviewed in Card et al. (2016)). That is, our estimated empirical sensitivity of wages to workers' outside job options – whether within, or outside, workers' own occupations – appears to be smaller than one might expect, given these other studies. On the other hand, this is similar in size to the labor share of income, which might suggest that workers' share of rents or quasi-rents is similar to workers' share of total value added.

Relevant literature: Our paper makes a number of contributions to related literatures. First, in using empirical worker transitions to infer the extent of workers' labor markets, our work builds on other papers which use worker flows to identify the scope of workers' geographic labor markets (Manning and Petrongolo (2017) in the UK), to identify firm clusters representing labor markets (Nimczik (2018) in Austria), and to study similarity in skill requirements across occupations and industries (Shaw, 1987; Neffke et al., 2017). We introduce a new, unique data set of occupational mobility in the U.S., constructed using 16 million worker resumes from Burning Glass Technologies. To our knowledge, we are the first to use resumes to construct occupational mobility data with which to infer workers' 'revealed' labor markets and measure outside-occupation options across the entire U.S.

Second, our results add to a growing literature on the effect of workers' local outside options in wage determination. Our results and identification for outside-occupation options relate most directly to Beaudry et al. (2012), who show for the U.S. that local changes in the availability of high-wage jobs in some industries have spillover effects on wages in all other local industries. Caldwell and Danieli (2018) also construct an index of the value of workers' outside options in Germany, which is based on the diversity of jobs and locations in which similar workers are observed, and find that this index is strongly associated with wages. Macaluso (2019) also studies the quality of local occupational labor markets but focuses on how the outcomes for laid-off workers vary depending on the similarity or dissimilarity of available local jobs, measuring similarity using occupational task and skill content. Our paper adds to this literature by developing a new way of measuring outside options based on observed transitions, and is the first to estimate the effect of occupational outside options in the U.S.

Third, our estimates of the effect of labor market concentration build on several recent papers that have documented a negative empirical correlation between wages and employer concentration in U.S. local labor markets, defining markets as a single local occupation or industry (Azar et al., 2017, 2018; Rinz, 2018; Lipsius, 2018; Benmelech et al., 2018; Hershbein et al., 2019). Most recently, Jarosch et al. (2019) develop a concentration index defined on clusters of firms, identified from worker flows, and Arnold (2020) estimates the effect of increases in local labor market concentration as a result of M& A activity on wages. We make two contributions to this literature. In developing a new instrument for changes in local labor market concentration, based on local employer "granularity" we are able to obtain plausibly causal and precisely estimated effects of local labor market concentration on wages. And, in integrating this concentration literature with an analysis of the effect of outside-occupation options and a flexible definition of workers' true labor markets, we show that the effects of labor market concentration on wages depend on the degree to which workers are able to find jobs outside their occupation – and therefore, that if measures of employer concentration are to be used for antitrust screening or assessments of monopsony power, as suggested by, e.g. Marinescu and Hovenkamp (2019), they should be adjusted to take into account outside-occupation options.

Finally, in estimating the degree to which wages are sensitive to outside options, we contribute to a broader literature on imperfect competition in labor markets. This includes the theoretical and empirical literature on labor market monopsony and the elasticity of the labor supply curve to the firm (e.g. Boal and Ransom, 1997; Manning, 2003; Ashenfelter et al., 2013; Webber, 2015; Bassier et al., 2019; Sokolova and Sorensen, 2020). It also includes the large search-and-matching literature, particularly models featuring wage bargaining between workers and firms (e.g. Burdett and Mortensen, 1980; Pissarides, 2000). Within these literatures our work is closest to Jaeger et al. (2019) and Berger et al. (2019). Jaeger et al. (2019) show that wages are very insensitive to unemployment in Austria, suggesting either an implausibly high Nash bargaining parameter or that unemployment is not a particularly valuable or relevant outside option for workers. Complementing this, our paper suggests that the value of moving to another job either within or outside workers' own occupation *does* affect wages and can generate a plausible – albeit still high – Nash bargaining parameter. Berger et al. (2019) derive and estimate a model-consistent measure of oligopsonistic wage markdowns that depends on the degree of within- and across-market substitutability of labor. They find welfare losses from monopsony power of roughly 5 percent of lifetime consumption.

Overall, our paper demonstrates – in data comprising over 100,000 U.S. occupation-city cells – that workers' job options *outside* their occupation and *within* their occupation both matter for wages. The sensitivity of wages to outside options is quite substantial – with a $1 increase in the value of outside options translating into a $0.20-$0.40 increase in wages – but is smaller than might be implied by empirical evidence on rent-sharing elasticities. Moreover, we show that considering the availability of outside-occupation job options is important when studying labor market concentration and its effects.

2 Theoretical framework

Which jobs are valuable outside options for workers, and to what extent does the quality and quantity of these outside job options matter for wages? In this section, we outline a simple search model to structure our analysis. Our model shows that under plausible assumptions about the likelihood of job offers, the value of a worker's outside option in the wage bargain can be expressed as a function of employer concentration *within* her occupation and a probability-weighted average of local wages *outside* her occupation.[3] The model is based on a labor market search-and-matching framework,with undirected search by workers and Nash bargaining over wages. To simplify exposition, we omit location subscripts, but all derivations and variables in this section should be assumed to apply to a specific city k.

Job search. Each job seeker coming from occupation o in city k applies to all

[3]In appendix Section F, we also show that the same probability-weighted average wage can be justified as a measure of outside-occupation job options in a simple matching model with heterogeneity in outside options and without search frictions. The true structure of the labor market is likely to be somewhere in between. Since both extreme case models can rationalize our measure, the exact structure of the search frictions in the labor market seems likely to be less important for our measure than the assumption that there is *on average* some heterogeneity across workers in terms of the occupation that is their best outside job option.

feasible local employers $j \in N$. Each employer offers the worker a job paying $w_{j,o}$ with probability α_j.[4] Once she has received all her offers, the job seeker accepts the offer with the highest wage. If she does not receive an offer from any employers in her feasible set N, she moves to unemployment for the period, receiving payoff b. She can then search again in the next period.[5] We assume that the wage offered to job seekers is the same as the wage obtained by workers bargaining on-the-job.

Wage bargaining. Each employed worker in occupation o Nash-bargains with her employer i at the start of each period. The outcome of wage bargaining is a wage $w_{i,o}$ equal to the value of the worker's outside option $oo_{i,o}$ if she leaves her job, plus a share β –reflecting worker bargaining power – of the match surplus created by the worker in working for that firm:

$$w_{i,o} = \beta(MPL_{i,o} - oo_{i,o}) + oo_{i,o} = \beta MPL_{i,o} + (1 - \beta) oo_{i,o} \tag{1}$$

The worker's outside option value $oo_{i,o}$ is equal to her expected payoff if she leaves her current employer and applies for jobs at all other feasible firms. She does not know her outside option with certainty: instead, the expected value of her outside option is her expected wage if she leaves her current job. We assume that, in expectation, all employed workers in the same occupation at the same firm have the same outside option.

Outside option value. The probability that a worker moves to any one employer j given that she leaves her existing job is the product of the probability that she receives an offer from that employer, α_j, and the probability that the wage offered to her by that employer is the maximum of all the wages offered to her this period.

The expected value of this outside option $oo_{i,o}$ is therefore a weighted average of the wages she would be paid at those other firms (where the weight on each firm's wage is the probability that she ends up moving to that firm if she leaves her current job), and of the payoff from unemployment b (where the weight is her probability

[4] This could encompass any employer-specific characteristic which influences the propensity to make a job offer, such as the employer's current labor demand or aggregate labor market tightness.

[5] Implicitly, each period a fraction of workers are exogenously displaced and search for a new job. These correspond to the transitions we observe in our data. While employed workers can also choose to leave their job, in equilibrium they do not because their employer offers them a wage which is weakly greater than the expected value of their outside options.

of becoming unemployed if she leaves her current job):[6]

$$oo_{i,o} = \sum_{j=1}^{N} \alpha_j \cdot w_{j,p} + \prod_{j=1}^{N} (1 - \alpha_j) \cdot b_t \tag{2}$$

Note that $\prod_{j=1}^{N}(1-\alpha_j)$ is the probability that the worker becomes unemployed if she leaves her current job (she receives no job offers). For simplicity, in this paper we do not consider job options outside the worker's city, and will assume that the outside option value of unemployment is negligible, i.e. $\prod_{j=1}^{N}(1-\alpha_j) \cdot b \approx 0$. [7]

Outside- vs. within-occupation job options: Since we focus in this paper on occupational labor markets and outside-occupation job options, we segment the set of feasible employers for a worker currently employed by i, into two categories: the outside options represented by employers in the worker's own occupation o, which we denote $oo_{i,o}^{own}$, and the outside options represented by employers in other occupations p, which we denote $oo_{i,o}^{occs}$. We therefore split equation (2) above into the probability that the worker's best job offer is in occupation p, denoted $Prob_{o \to p}$, as well as the probability that the best job offer *within* occupation p is in some firm j, and the wage offered by firm j. We further assume that the probability that a worker in city k would move to a job in firm j in occupation p, conditional on moving to *some* job in occupation p, is proportional to firm j's local employment

[6] b represents the value of any unemployment benefit the worker receives *plus* the monetary equivalent of any utility the worker receives from being unemployed.

[7] While these components of outside options certainly matter, the data suggests that for most workers they will be less important than the jobs in other occupations in workers' current city. Occupational mobility is substantially higher than geographic mobility: only around 3% of U.S. workers move between metropolitan areas each year. And the outside option value of unemployment is likely to be small for most workers, both since unemployment rates are generally in the single digits, and since unemployment benefits are low in the U.S. Jaeger et al. (2019) find that the outside option value of non-employment is negligible for most workers in Austria, which has both higher unemployment and more generous unemployment benefits than the U.S. (Note also that any time-varying differences in the value of unemployment will be captured by city-year fixed effects in our main empirical specifications.) In appendix Section F, we sketch how our theory might be extended to consider job options outside workers' city.

share defined as $\sigma_{j,p}$, following Burdett and Mortensen (1980). This implies:

$$
\begin{aligned}
oo_{i,o} =& oo_{i,o}^{own} + oo_{i,o}^{occs} \\
&= Prob_{o\to o} \cdot \sum_{j\neq i}^{N_o} Pr(\text{job at firm } j \text{ is best offer in occ } o) \cdot w_{j,o} \\
&+ \sum_{p\neq o}^{N_{\text{occs}}} Prob_{o\to p} \cdot \sum_{m=1}^{N_p} Pr(\text{job at firm } m \text{ is best offer in occ } p) \cdot w_{m,p} \qquad (3) \\
&= Prob_{o\to o} \cdot \sum_{j\neq i}^{N_o} \sigma_{j,o} \cdot w_{j,o} + \sum_{p\neq o}^{N_{\text{occs}}} Prob_{o\to p} \cdot \sum_{m=1}^{N_p} \sigma_{m,p} \cdot w_{m,p}, \qquad (4)
\end{aligned}
$$

where N_x is the set of employers in occupation x, and N_{occs} is the set of occupations.

The second term in this expression – $oo_{i,o}^{occs}$ – states that the ex-ante value of the component of workers' outside options based on jobs in *other* occupations is simply the weighted average of wages in other occupations, weighted by the probability that the best job offer the worker receives will be in occupation p.[8]

The first term – $oo_{i,o}^{own}$, the component of workers' outside options based on jobs in their *own* occupation – is a function of the probability that a worker's best offer will be in their own occupation ($Prob_{o\to o}$), and the wages offered by each local firm in that occupation ($w_{j,o}$). But note that this means that wages at any firm are a function of wages at other firms in the local area – which, in turn, are a function of wages at other firms, including the original firm.[9] To solve this "reflection problem", we first put our expression for the value of workers' outside options back into the Nash bargaining equation:

[8]Our model assumes rational expectations for both workers and employers. Greater uncertainty around occupational transition probabilities should not affect the form of the outside-occupation option index, but will increase noise in the empirical measure. Systematically biased expectations by workers and/or employers will affect the level of the outside option index but not our regression results. The only case in which uncertainty might bias our regression results is if workers and firms in certain occupations are *both* more likely to systematically under-estimate workers' ability to leave the occupation, *and* more likely to have higher wages than other occupations (or vice versa).

[9]This is also true of wages in other occupations, which are functions of wages in the initial occupation. We abstract away from this as the "reflection problem" from one occupation to another is second order given the magnitude of most occupation-to-occupation transition paths.

$$
\begin{aligned}
w_{i,o} &= \beta MPL_{i,o} + (1-\beta)\left(oo_{i,o}^{own} + oo_{i,o}^{occs}\right) \\
&= \beta MPL_{i,o} + (1-\beta)\left(Prob_{o\to o} \cdot \sum_{j\neq i}^{N_o} \sigma_{j,o} \cdot w_{j,o}\right) + (1-\beta) oo_{i,o}^{occs} \qquad (5)
\end{aligned}
$$

Making the simplifying assumption that the MPL is the same in all firms in the same occupation and city, i.e. $MPL_{j,o,k} = MPL_{o,k} \ \ , \forall i, j$, we can then iteratively substitute the wage equation (5) into itself to take into account the reflexive nature of the effect of outside options on wages. This results in the following simple expression for equilibrium average wages in a local occupation

$$
\ln \bar{w}_o = \ln\left(\beta MPL_o + (1-\beta) oo_o^{occs}\right) + \ln \psi_o. \qquad (6)
$$

where ψ_o is a power function of mobility-weighted higher-order concentration indices and is defined as

$$
\begin{aligned}
\psi_o &= \sum_{r=0}^{\infty} (1-\beta)^r Prob_{o\to o}^r \Omega_{o,r} \\
\Omega_{o,r} &= \sum_i \sigma_i \sum_{j\neq i} \sigma_j \sum_{l\neq j} \sigma_l \cdots \sum_{z\neq y} \sigma_z (1-\sigma_z),
\end{aligned}
$$

where the r index indicates the order of aggregation, i.e. the degree to which higher-order indirect connections are accounted for, and all the sums are taken over N_o – all employers in occupation o.[10] It is helpful to note that $\Omega_{o,0} = \sum_i \sigma_i = 1$ and $\Omega_{o,1} = \sum_i \sigma_i(1-\sigma_i) = 1 - HHI_o$. Moreover, note that ψ_o is a sequence of increasingly smaller terms which go to zero very quickly if either the probability of finding a job in another occupation $(1 - Prob_{o\to o})$ or worker bargaining power β are sufficiently far from zero. Therefore, approximating the concentration term ψ_o as a function of only concentration indices up to order $r = 1$, and adding city indices k, we can write local occupational average wages, as a function of local

[10]Full derivations shown in appendix E.

within-occupation concentration $HHI_{o,k}$ and outside-occupation options $oo^{occs}_{o,k}$:

$$\ln \overline{w}_{o,k} \approx \ln \left(\beta MPL_{o,k} + (1-\beta) oo^{occs}_{o,k}\right) + ln\left(1 + (1-\beta) Prob_{o\to o}(1 - HHI_{o,k})\right) \tag{7}$$

This suggests several testable implications. First, better outside-occupation options should be associated with higher average wages. Second, the empirical relationship between wages and employer concentration (HHI) should be negative, and should be stronger for occupations where workers have little ability to get jobs in other occupations (large $Prob_{o\to o}$). Third, the empirical relationship between wages and HHI may be biased if concentration within an occupation is correlated with the availability of outside-occupation job options (as is empirically the case).

3 Measuring Outside Options

Our theory suggests that we need two empirical objects to estimate the effect of within- and outside-occupation options on wages: an index of the value of outside-occupation options, and a measure of within-occupation local employer concentration.

Outside-occupation options. To construct our index of the value of outside-occupation options, we need an empirical analog for the probability that a job in occupation p and city k will be the best job offer for a worker in occupation o and city k ($Prob_{o\to p}$). We use the product of two variables: (1) the national average empirical transition share between occupation o and occupation p for workers that leave their job, $\pi_{o\to p}$, and (2) the relative employment share of occupation p in city k compared to the national average, $\frac{s_{p,k}}{s_p}$. The national occupation-to-occupation transition share is a proxy for the likelihood that, nationwide, the average worker's best job option outside her own occupation would be in each of these other occupations; the local relative employment share adjusts this for the local availability of job options in each destination occupation:

$$\begin{aligned} Prob_{o\to p} &= \frac{\text{workers moving from occ } o \text{ to occ } p}{\text{workers leaving job in occ } o} \cdot \frac{\text{emp. share in occ } p \text{ in city } k}{\text{national emp. share in occ } p} \\ &= \pi_{o\to p} \cdot \frac{s_{p,k}}{s_p} \end{aligned} \tag{8}$$

This implies that we can write the average expected value of outside-occupation job options for workers (the second term in equation (4)) in occupation o, city k, and year t as the weighted average of local wages in other occupations $\bar{w}_{p,k,t}$, with the weights we just defined:

$$oo^{occs}_{o,k,t} = \underbrace{\sum_{p\neq o}^{N_{occs}} \pi_{o\to p} \cdot \frac{s_{p,k,t}}{s_{p,t}} \cdot \bar{w}_{p,k,t}}_{\text{outside-occupation job options}} \tag{9}$$

To our knowledge, we are the first to use empirical transitions observed in worker career paths to construct a measure of local outside options in this way.

We construct this outside-occupation option index at the annual level for as many SOC 6-digit occupations and US cities over the years 1999–2016 as our data allows (and show summary stats in Table 1).[11] We use data from the BLS Occupational Employment Statistics (OES) to construct the relative employment shares $\frac{s_{p,k,t}}{s_{p,t}}$ and average wages $\bar{w}_{p,k,t}$ by SOC 6-digit occupation, city, and year.[12] Since there is no existing data which allows us to calculate occupation-to-occupation transitions at a granular SOC 6-digit level for a large number of of occupations, we construct the occupation-to-occupation transition shares $\pi_{o\to p}$ from a new and unique data set of 16 million U.S. resumes over 2002–2018, collected by Burning Glass Technologies. We describe this data in more detail below.

Concentration indices. To calculate Herfindahl-Hirschman Indices (HHIs) of

[11] We use "cities" to refer to the CBSAs (metropolitan and micropolitan statistical areas) and NECTAs (New England city and town areas) for which data is available in the BLS OES.

[12] Creating a consistent panel of occupations over time requires us to crosswalk SOC classifications when the OES updates occupation codes. Details on the crosswalk used can be found in appendix Section B. Note also that, of the possible 786,335 occupation-city cells, wage data in the BLS OES only exists for approximately 115,000 each year. The missing occupations and cities are primarily the smaller ones.

employer concentration, we use Burning Glass Technologies' database of online vacancy postings (following Azar et al., 2018; Hershbein et al., 2019), calculating the vacancy HHI at the level of a SOC 6-digit occupation by city by year.

Using occupational transitions to identify outside-occupation options

When constructing our outside-occupation option index, we use empirical occupation-to-occupation transitions to infer the likelihood that a job in a given occupation is a valuable outside option. Why?

The outside option value of a job in another occupation can be thought of as having two dimensions. One dimension is feasibility: the likelihood that the worker could easily become a typically productive worker in the new occupation. The other dimension is desirability: the degree to which the worker would like to do a job in the new occupation, compared to a job in their current occupation. Occupational transitions capture some combination of *both* feasibility and desirability: by definition the occupational transitions that occur were feasible, and in most cases since occupational transitions involve some element of choice, the new occupation is likely on average similarly or more desirable than the old occupation (incorporating the value of amenities such as work-life balance, status, or career concerns).

We believe that using occupational transitions is a better way to identify workers' outside options than the two other most commonly-used methods of estimating occupational similarity: estimating the similarity between occupations in the skills or tasks that they normally require, or inferring the extent of labor markets from the similarity in the demographic characteristics or qualifications of the workers who are found in certain occupations.[13] Task/skill similarity measures do not capture unobserved constraints that prevent moves in practice (e.g. regulation), and do not capture the desirability of moves between different occupations; in addition, they assume symmetry in the likelihood of moves between pairs of occupations. Characteristic similarity measures suffer from similar issues. In contrast, a transition-based measure will reveal whether or not a move is difficult, even if the precise barrier is

[13]For example, Macaluso (2019) measures occupational skill similarity using the vector difference of occupational skill content, and Gathmann and Schönberg (2010) use the angular separation of occupations' task content vectors.

unknown to the researcher. We note, however, that a transition-based measure relies on the assumption that the transitions which are observed are representative of the options faced by *all* workers, including those not switching occupations (see online appendix Section C for a more detailed discussion).

Resume data from Burning Glass Technologies

Our data on occupational and job transitions is constructed from a new proprietary data set of 16 million unique resumes with more than 80 million job observations over 2002–2018, provided by labor market analytics company Burning Glass Technologies ("BGT"). Resumes were sourced from a variety of BGT partners, including recruitment and staffing agencies, workforce agencies, and job boards. Since we have all data that people have listed on their resumes, we are able to observe individual workers' job histories and education up until the point where they submit their resume, effectively making it a longitudinal data set.

For the purpose of calculating occupational transition probabilities, the biggest advantage of the BGT data set is its size: the largest publicly-available data set on which occupational transitions can be calculated, the CPS, has at least an order of magnitude fewer occupational transition observations (aggregated over the same time period 2002–2018). This matters a great deal when calculating pairwise occupational transition shares between the 840 SOC 6-digit occupations: with 705,600 possible transition cells, data sets with even a few million observations will not be big enough to capture a great many of the transition paths (and even our data is not big enough to capture all of them with precision).

We would ideally use the BGT data to estimate annual transition probabilities between full-time jobs in particular occupations. Unfortunately, resumes often do not list the months in which jobs started and ended, and do not always indicate if jobs were part-time. We therefore approximate the share of workers moving from occupation o to occupation p with the share of all workers observed in occupation o at any point in year t who are observed in occupation p at any point in year $t + 1$.[14] Similarly, we approximate the share of workers in occupation o who take a new job

[14]We drop jobs which are listed as lasting for 6 months or less to exclude temporary work, summer jobs, and internships.

as the share of all workers observed in a given job in occupation o at any point in year t who are observed in a different job at any point in year $t+1$.[15]

Our measure of occupational transitions

Our main measure of occupational transitions $\pi_{o\to p}$ is the probability of a worker moving from occupation o to occupation p conditional on leaving her job, defined formally as follows:

$$\begin{aligned}\pi_{o\to p} &= Pr(\text{move from occ } o \text{ to occ } p | \text{leave job}) \\ &= \frac{\text{Share from occ } o \text{ moving into occ } p}{\text{Share from occ } o \text{ moving into a new job}} \\ &= \frac{\text{\# working in occ } o \text{ in year } t \text{ who are observed in occ } p \text{ in year } t+1}{\text{\# working in occ } o \text{ in year } t \text{ who are observed in a new job in year } t+1}\end{aligned} \tag{10}$$

We estimate these occupation transition probabilities at the national level for a large proportion of the possible pairs of SOC 6-digit occupations: we exclude the occupations for which we have fewer than 500 observations in the BGT data (roughly the bottom 10% of occupations), resulting in 786 origin SOC 6-digit occupations in our data and 285,494 non-empty occupation-to-occupation transition cells out of a total 705,600 possible transition cells (840 x 840). We average the observed annual occupation-to-occupation transitions over all observations over starting years 2002–2015,[16] to capture as much as possible the underlying degree of occupational similarity rather than transitory fluctuations in mobility.

We calculate one further statistic from this data: the 'occupation leave share', which we use as an approximation to the share of people leaving their jobs who also

[15]Note that this measure will also capture mobility between occupations in the form of working in two different occupations at the same time, as well as job mobility that consists of taking a new job while continuing to work in an old job. Implicitly, we are assuming that taking up a secondary job in an occupation indicates its viability as an outside option to the same degree as moving primary occupations. Under this assumption, our measure is more appropriate than one that focuses only on transitions that involve abandoning a previous job or occupation entirely.

[16]Where 2015 refers to the *starting* year, i.e. takes individuals in occupation o in 2015 and observes their occupation in 2016. We exclude observations from starting years 2016 or 2017 to avoid bias from the resume collection process: if we observe someone applying for a job in 2017 who has also changed job in 2016, they are not likely to be representative of the average worker.

leave their SOC 6-digit occupation:

$$
\begin{aligned}
\text{leave share}_o &= \frac{\text{Share from occ } o \text{ leaving occ } o}{\text{Share from occ } o \text{ moving into a new job}} \\
&= \frac{\text{\# working in occ } o \text{ in year } t \text{ \& no longer in occ } o \text{ in year } t+1}{\text{\# working in occ } o \text{ in year } t \text{ \& in a new job in year } t+1}
\end{aligned}
\tag{11}
$$

The BGT resume data set is largely representative of the U.S. labor force in its distribution by gender and location. However, it over-represents younger workers and white-collar occupations. Since we are estimating occupational transition probabilities within each occupation, the over-representation by occupation is not a substantial concern as long as we still have sufficient data for most occupations to have some degree of representativeness *within* each occupation. We correct for the over-representation by age by re-weighting the observed occupational transitions by the U.S. age shares by occupation, provided by the BLS for 2012-2017.[17]

Descriptive evidence on occupational mobility

Is it sensible to use occupational transitions to infer outside-occupation options? Using the BGT data, we document a number of facts which suggest it is.

First, occupational mobility is high, highly heterogeneous, and poorly captured by aggregating up the SOC hierarchy, suggesting that SOC occupational categories are not appropriate measures of workers' true labor markets – and that a flexible approach to labor market definition is needed. In our data, the average probability of a worker leaving her 6-digit occupation given that she leaves her job - the "occupation leave share" defined in equation (11) above - is 23%.[18] The full distribution,

[17]See the Data Appendix for further discussion on representativeness and sample selection.

[18]This is constructed from the average share of workers leaving their occupation (11%) and the average share of workers leaving their job (46%) in any given year. Occupational mobility of 11% in our data is somewhat lower than the occupational mobility estimate from Kambourov and Manovskii (2008), who find occupational mobility of 0.20 at the Census 3-digit level in the CPS for the late 1990s, but is in a similar range to Xu (2018) who finds occupational mobility of 0.08 in 2014. The fact that our measure is relatively low compared to Kambourov and Manovskii (2008) is interesting, since sample selection bias would be expected to *overstate* occupational mobility in our data set if the people applying for jobs (whose resumes we observe) are more mobile than average. Note however that our measure of outward occupational mobility is not strictly comparable to the concept

shown in Table 1, indicates that there is large variation in the average share of workers leaving their occupation when they leave their job, with the 25th percentile of occupations at 19% and the 75th percentile at 28%.[19] Aggregating up the SOC classification hierarchy - which groups occupations with ostensibly similar occupations - is not sufficient to capture the full extent of workers' true labor markets. For the median occupation, 87% of moves to a different 6-digit occupation are also to a different 2-digit occupation, but with substantial variation (see Table 1).[20] For example, only 39% of systems software developers move 2-digit occupations when they move across 6-digit occupations, compared to 95% of flight attendants.[21]

Second, the occupational transition matrix is highly asymmetric, suggesting that the relevance of other occupations is not symmetric across occupation pairs (which is the case in many task- and skill-based measures of occupational similarity). The asymmetry in occupational transitions partly reflects career progression, and partly reflects the fact that some occupations are fall-back job options for many different other occupations. For instance, workers in an occupation with specialized skills may move to one which requires generalist skills (e.g. retail salespersons) but the reverse flow is less feasible.

Third, as would be expected, there are few observed transitions between most

of annual occupational mobility used by Kambourov and Manovskii (2008) and Xu (2018) because of the nature of our resume data, as discussed above. This is consistent with the average length of a job in our data being 2 years. Job mobility of 46% in our data implies the average duration of a job is 2 years. Leaving your job does not necessarily entail leaving your firm: the CPS reports that median employee tenure in 2018 was 4.2 years, so an average job duration of 2 years in our data is consistent with workers working on average 2 consecutive jobs at the same employer.

[19] Almost all of the occupations with low leave shares are highly specialized, including various medical, legal and educational occupations (see Table 2). In contrast, many of the occupations with high leave shares require more generalizable skills, including restaurant hosts/hostesses, cashiers, tellers, counter attendants, and food preparation workers. The difference in mobility can be substantial: over 30% of telemarketers or hosts and hostesses leave their occupation when they leave their job – compared to around 10% for pharmacists, lawyers or licensed practical and vocational nurses. This suggests that the SOC 6-digit occupation is a substantially better measure of the true labor market for some occupations than for others.

[20] Note that management roles are often considered a separate 2-digit occupational group from non-management roles in the same field or specialty. If we exclude transitions to and from management occupations entirely, at the median 67% of SOC 6-digit occupational transitions still cross SOC 2-digit boundaries (see Table 1).

[21] 2-digit SOC occupation groups fail to capture both lateral moves across occupations groups *and* promotions into managerial roles. See online appendix Figures A9 and A10 for examples.

pairs of occupations: the transition matrix is sparse (as shown in online appendix Figure A7), suggesting that workers' relevant labor markets can often be constructed out of relatively small clusters of related occupations (as we do in this paper). There are only a few 'thick' occupational transition paths where the transition probability is greater than negligible.[22]

Fourth, empirical occupational transitions reflect systematic similarities between occupations in terms of their task requirements, wages, amenities, and leadership responsibilities. This suggests that the empirical occupational transitions we observe do indeed reflect the underlying feasibility and/or desirability of an occupation as an outside option. To determine the effect of differences in characteristics between occupations on the likelihood of observing worker transitions, we estimate regressions of the form

$$\pi_{o \to p} = \alpha_o + \beta f(X_{occ}) + \gamma f(\Delta w_{occ}) + \epsilon_{op}, \tag{12}$$

where function $f(\cdot)$ represents the absolute or relative difference in characteristic X_{occ} or wage w_{occ} when moving from occupation o to p, and α_o is an origin occupation fixed effect. (We only summarize the main results in this section, with details provided in online appendix Section D.) We consider a number of occupational characteristics X_{occ}, derived from the O*Net database: the vector difference in the importance scores for all "Skill" task content items (see Macaluso (2019)); task composites capturing the distinction between cognitive vs. manual, routine vs. non-routine task contents, and social skills, based on Autor et al. (2003) and Deming (2017); characteristics that proxy for flexibility on the job (Goldin, 2014), such as time pressure and the need for establishing and maintaining interpersonal relationships; and characteristics measuring leadership responsibilities. We also control for, and measure the effect of, wage differences.

As would be expected if occupational transitions capture structural similarities between occupations, in every pairwise regression of occupational mobility on the absolute difference in characteristics, the coefficients are significantly negative or

[22]For example, conditional on moving into a new occupation, there are only 189 pairs of 6-digit occupations which have a transition probability of 10% or greater (out of 284,797 observed non-zero occupational transition cells).

statistically insignificant, as shown in Figure 1.[23] Similarly, as would be expected if occupational transitions capture career progressions, we find that workers are more likely to move towards jobs with higher wages and that they transition on average into jobs that require more leadership responsibility.[24][25]

4 Empirical Approach

Having constructed empirical indexes of outside-occupation options and local labor market concentration, we can now estimate the effect of changes in these outside options on worker wages.

Regression specification. To be able to estimate the parameters of interest, we totally differentiate the log wage expression in equation (7) derived from our model around a city-occupation with median characteristics (see Online Appendix section E). We obtain the following baseline specification:

$$\ln \bar{w}_{o,k,t} = \alpha + \alpha_{o,t} + \alpha_{k,t} + \gamma_1 \ln oo^{occs}_{o,k,t} + \gamma_2 \ln HHI_{o,k,t} + \xi_{o,k,t} \tag{13}$$

where $\alpha_{o,t}$ and $\alpha_{k,t}$ are occupation-year and city-year fixed effects, and $\xi_{o,k,t}$ is the local occupation productivity residual that is orthogonal to these fixed effects.

Note that (based on the derivation in Online Appendix E) the outside option and concentration parameters we estimate in regression 13 correspond to

$$\gamma_1 = \frac{(1-\beta)\bar{oo}^{occs}_o}{\beta M\bar{P}L_o + (1-\beta)\bar{oo}^{occs}_o} \quad , \quad \gamma_2 = -\frac{(1-\beta)\pi_{o\to o} H\bar{H}I_o}{1+(1-\beta)\pi_{o\to o}(1-H\bar{H}I_o)} \tag{14}$$

where characteristics are computed for a reference city-occupation cell. We use

[23] In a similar vein, Macaluso (2019) finds that mobility between SOC 2-digit occupation groups in the US is highly correlated with task similarity.

[24] This is in regressions of occupation-to-occupation transitions on the *relative* (target minus origin) difference in occupational characteristics. Nimczik (2018) similarly finds that most job transitions in Austria move up the career ladder.

[25] Our measured transitions also seems to reflect major labor market trends documented elsewhere (Acemoglu and Autor, 2011; Goldin, 2014): occupational transitions have on average been towards occupations that require more analytical and social skills, and out of occupations with more routine task requirements; and workers have on average been moving into occupations that require more contact with others (and thus have less time flexibility).

these expressions below to give a structural interpretation to our causal estimates.

As in the construction of our outside option index, we use BLS OES data for average occupational wages by city and year for the dependent variable $\bar{w}_{o,k,t}$. Our full data set for the regressions comprises 1.94 million occupation-city-year observations: 394 cities and 753 6-digit SOC occupations over 17 years (1999–2016).[26] We also run versions of most regressions separately by quartile of the occupation leave share, measured among the included occupations.[27]

Identification: outside-occupation options

Endogeneity issues may be expected to bias the coefficients on our outside-occupation option measure upwards: in a year where a worker's city experiences a positive demand shock for an occupation similar to her own, there may also be a positive local demand shock for her own occupation (driven, for example, by a common product market shock or a regulatory change). In addition, there is a reverse causality or reflection problem: if occupation p is an outside option for workers in occupation o, and occupation o is an outside option for workers in occupation p, then a wage increase in o will increase wages in p and vice versa.

To identify causal effects, we therefore need exogenous shocks to the wages in workers' outside-occupation options which do not affect and are not affected by the local wages in their own occupation.[28] We instrument for local wages in each outside option occupation with plausibly exogenous national demand shocks to that occupation. Specifically, to instrument for wages in each outside-option occupation p in a worker's own city k, we use the leave-one-out national mean wage for occupation p, excluding the wage for occupation p in city k.

In addition, to avoid endogeneity concerns over the local employment shares,

[26] Although not all SOC occupations have data for all cities or all years. We have only 753 out of a possible 841 SOC occupations because (a) some occupations are not present in the BLS OES data, and (b) we exclude all occupations for which we have fewer than 500 person-year observations in the BGT occupational transitions data. Our regression results are robust to including these occupations.

[27] Note that, in the regressions, each quartile will not have the same number of observations, as the number of cities and years in which each occupation is observed differs between occupations.

[28] A range of papers identify the effects of plausibly exogenous empirical shocks on outside options in specific settings. Caldwell and Harmon (2018), for example, shows in Danish data that wage increases at firms at which a worker's former co-workers now work affect the worker's own wage, likely through information about outside options.

we instrument for the local relative employment share in each occupation using the initial employment share in that occupation in 1999, the first year in our panel.[29] Our instrument for the oo^{occs} index, $oo^{occs,inst}$, therefore becomes the weighted average of national leave-one out mean wages in occupation p, $\bar{w}_{p,\not{k},t}$, where the weights are the product of the year 1999 relative employment share in each of those occupations in the worker's own city, $\frac{s_{p,k,1999}}{s_{p,1999}}$, and the national occupation transition shares from the worker's occupation o to each of the other occupations, $\pi_{o\rightarrow p}$:

$$oo^{occs,inst}_{o,k,t} = \sum_{p}^{N_{occs}} \left(\pi_{o\rightarrow p} \cdot \frac{s_{p,k,1999}}{s_{p,1999}} \cdot \bar{w}_{p,\not{k},t} \right) \tag{15}$$

Identifying variation within the same occupation across different cities comes from differences in each city's initial exposure to outside option occupations.[30] Identifying variation over time within the same occupation-city cell comes from national (leave-one-out) changes over time in wages of local outside-option occupations. That is, in a year when there is a national wage shock to one of occupation o's outside option occupations p, cities which had a higher proportion of their jobs in occupation p in 1999 should see bigger increases in the wage of occupation o (because they were more exposed to the shock to their outside options).

In appendix Section G, we provide more details on the conditions required for our instrument to identify the effect of outside-occupation options on wages, using the approach in Borusyak et al. (2018). Intuitively, the assumptions require that the national leave-one-out mean wage $\bar{w}_{p,\not{k},t}$ in outside option occupation p is correlated with the local wage of occupation o in location k (relevance condition), but does not affect the local wage in initial occupation o through a direct channel other than increasing the quality of local outside options $oo^{occs,inst}_{o,k,t}$. Note that this only needs to hold *conditional* on controlling for fixed effects which, in our main specification, include the national wage trend in occupation o itself, and wage trends common to all occupations in city k. The inclusion of these fixed effects means that differences

[29]Or we use the first year the occupation-city cell is in the data, if it is not present in 1999.

[30]This refers to the relative employment share of each occupation p in city k compared to the national average, in either 1999, or the first year that occupation-city cell is present in the OES.

in city-level trends or national productivity of different occupations do not represent an issue for our identification strategy.[31] As an example, note that a correlation at the national level in the wages of a pair of occupations (e.g. of Compliance Officers and Financial Analysts), perhaps due to common industry shocks (e.g. to the banking industry) does *not* invalidate this identification strategy, because we are holding national wage trends constant for each occupation and are identifying outside option effects from the differences between cities *within* occupations.

Identification: concentration index

Similarly, endogeneity issues may be expected to bias the coefficients on our HHI concentration measure. The direction of the bias is ambiguous: an increase in concentration of employment could reflect the expansion of a highly productive large firm, which would lead to an increase in monopsony power of that firm (expected to reduce wages) but also an increase in productivity (expected to increase wages). On the other hand an increase in concentration of employment could reflect a lack of local dynamism, with few new firms being created, which would lead to an increase in monopsony power (expected to reduce wages) but also a general decrease in productivity (which would also reduce wages).

We therefore instrument for local labor market concentration, using the "granular" instrumental variable approach (GIV) of Gabaix and Koijen (2020), which uses plausibly exogenous idiosyncratic firm-level variation to instrument for changes in market-level aggregates.

Note that, mechanically, the growth in the local employer concentration of occupation o will be a function of the growth in occupational employment in all em-

[31]This instrumental variable strategy is closely related to that of Beaudry et al. (2012), who argue that they are able to avoid endogeneity and reflection problems in their index of cities' industrial composition by using national industry wage premia to substitute for city-level industry wages.

ployer firms j (leaving aside firm entry):

$$\begin{aligned}\Delta HHI_{o,k,t} &= \sum_j \sigma^2_{j,o,k,t} - \sum_j \sigma^2_{j,o,k,t-1} \\ &= \sum_j \left(\sigma^2_{j,o,k,t} - \sigma^2_{j,o,k,t-1}\right) \\ &= \sum_j \sigma^2_{j,o,k,t-1}\left(\frac{(1+g_{j,o,k,t})^2}{(1+g_{o,k,t})^2} - 1\right)\end{aligned}$$

The increase in concentration is a function both of initial concentration and of the growth rates of firm-level vacancies $g_{j,o,k,t}$ relative to overall vacancy growth $g_{o,k,t}$.

This means we can use plausibly exogenous *firm-level* changes in vacancies as instruments for the overall change in local labor market concentration. For large firms operating in more than one local labor market, the leave-one-out firm-level national growth in all vacancies, $\tilde{g}_{j,t}$ is likely to be uncorrelated with determinants of occupation-specific productivity growth in any given local labor market k. This makes $\tilde{g}_{j,t}$ a valid instrument for $g_{j,o,k,t}$.

Our instrumental variable for the change in concentration is therefore:

$$\Delta HHI^{inst}_{o,k,t} = \sum_j \sigma^2_{j,o,k,t-1}\left(\frac{(1+\tilde{g}_{j,t})^2}{(1+g_{o,k,t})^2} - 1\right)$$

One concern with this instrumental variable is that growth in large national firms may also differentially affect total labor demand in a given occupation in a given city. Exploiting the fact that the effect of a large firm's growth on local labor market concentration is quadratic, whereas the effect of a large firm's growth on local labor demand is linear, we control for the (linear) growth rate of local vacancies.[32] With this control, we should be estimating the effect of a change in local labor market concentration due to changes in large firms' employment, holding constant any direct effects on local labor demand (which would be captured by the growth in

[32] In robustness checks, we include two additional controls: (1) a control for the linear *predicted* growth rate of local vacancies, as predicted by national firm growth, of the form $\tilde{g}_{o,k,t} = \sum^{N_o} \sigma_{j,o,k,t-1}\tilde{g}_{j,t}$, and (2) a control for the equal-weighted predicted growth rate of local vacancies, as detailed in Gabaix and Koijen (2020). These additional controls have almost no effect on the estimated coefficients.

total vacancies).[33]

We believe this GIV approach is a novel instrument for local labor market concentration. Our instrument exploits the fact that (a) increases in local labor market concentration are often driven by specific, already-large, firms growing, (b) these large firms usually operate across many labor markets, (c) that different local labor markets are differentially exposed to the growth of these large national firms based on the initial share of employment with that firm, and (d) that the employment growth of these large firms nationally is likely to be orthogonal to productivity changes or other conditions in a specific local labor market. Consider a hypothetical example of baristas in Portland, OR, and in Hartford, CT. Imagine there are two large coffee chains in each city: Starbucks, and Dunkin'. In Portland, Starbucks is relatively more dominant than Dunkin', and the reverse in Hartford, meaning that Starbucks' local employment share is much larger in Portland than in Hartford. In years where Starbucks grows substantially faster than Dunkin' nationwide, employer concentration of baristas will grow by more in Portland than in Hartford.

For comparison, note that some recent empirical work on labor market concentration has instrumented for changes in concentration in a given occupation and city with changes in concentration in the same occupation in other cities (e.g. Azar et al. (2018)). However, that identification strategy is subject to the concerns that national occupation trends in concentration are correlated with unobservable national trends in occupational productivity, which could be confounding any measured wage effects. In contrast, our strategy allows us to control for national occupational trends because we are exploiting across-city variation in concentration due to differences in exposure to firm-level growth. Another recent approach is that of Arnold (2020) who uses cross-sectional variation in local exposure to M&A activity to generate plausibly exogenous variation in local labor market concentration. Our approach is similar in using differential local exposure to large national firms' activity – but, by using the GIV approach, we are not restricted only to studying M&A activity by specific firms and in specific geographies.

[33] See online appendix Section G for more details on conditions for identification. Controlling for national trend exposure directly to prevent it from confounding a nonlinear instrumental variables approach is similar to the "double Bartik" approach in Chodorow-Reich and Wieland (forthcoming).

5 Empirical results

Effect of outside-occupation options on wages

First, we estimate the effect of outside-occupation options on wages over 1999–2016 using the regression shown in equation (13) (without including an HHI measure of within-occupation labor market concentration, because our HHI data only covers 2013–2016). Panel A of Table 4 shows the simple OLS estimates and Panel B shows the instrumental variable estimates.

Column (1) shows that, only including year fixed effects, we find a strong positive relationship between outside-occupation options and wages. Column (2) has occupation-year and city fixed effects and column (3) has occupation-year and city-year fixed effects: they show that occupation-city-year cells which have higher oo^{occs} compared to the national average for their occupation have significantly higher wages. Column (4) has occupation-by-city and occupation-by-year fixed effects, and so identifies only off annual variation in outside-occupation options compared to their mean for each occupation-by-city and occupation-by-year unit.[34]

The IV coefficient magnitudes in columns (2) through (4) suggest that a 10 log point higher outside-occupation option index is associated with 0.3-0.8 log points higher wages in the workers' own occupation. This implies that, *within* a given occupation and year, moving from the 25th to the 75th percentile value of outside-occupation options[35] leads to 1.4-3.5 log points higher wages in the workers' own occupation. These results would also imply that nationwide demand shocks to relevant outside option occupations can cause meaningful indirect spillover effects to local occupational wages.[36]

[34]While our sample size of $\sim$ 1.9M data points makes weak instrument concerns unlikely to be an issue, we can formally verify this by computing the first-stage F-statistic for the instrument. As the bottom of Panel B in Table 7 shows, the instrument for other occupation outside options is highly statistically significant, with F-statistics between 348 and 4921 depending on the exact specification. Moreover, the coefficient ranges between 0.81 and 1.01, in line with an expected positive relationship that is somewhat attenuated by measurement error.

[35]This represents the average interquartile range of the logged outside-occupation option index *within* each occupation and year across different cities, which has a value of 0.46.

[36]Note that our outside-occupation option index include transitions to *all* SOC 6-digit occupations for which we have sufficient data, including management occupations. Transitions into management occupations may represent promotion rather than lateral moves, so may be less likely to represent relevant outside options. Our inclusion of management occupations likely therefore attenuates our

These estimates can be applied to explain patterns of pay variation for specific occupations in specific cities. Consider Baltimore, MD, and Houston, TX. Both metropolitan areas are large (1-3 million people), with a similar average wage in 2016 (\$25-\$26 per hour). Statisticians in Baltimore earn 26% more than statisticians in Houston. These differences could clearly partly be explained by differences in the productivity of the occupations in the different cities; however, our estimates of the value of outside-occupation job options would suggest that between 7 and 17% of this gap could be attributed to differential availability of outside-occupation job options in the two cities.

In the Appendix, we explore a number of variations on our baseline analysis of the effect of outside-occupation job options on wages. First, we re-run our baseline regressions with the outside-occupation option index computed from occupational mobility at the SOC 2-digit and SOC 3-digit levels (appendix Table A4), to address concerns about possible mismeasurement in occupational mobility at the level of detailed 6-digit SOC codes. Second, we re-run our baseline regressions separately for three different time periods, corresponding roughly to the boom (2002-2007), the recession (2007-2012), and the recovery (2012-2016), to address concerns that our results may be driven by the idiosyncrasies of any of these periods (appendix Table A5). Third, we re-run our baseline regressions with weights corresponding to local occupational employment, to address concerns that the results may be driven by the dynamics of smaller occupations (appendix Table A6). In all three cases, the results remain similar to our baseline specification, with a large and significant effect of our instrumented outside-occupation option index on wages.

Channels: wage bargaining and mobility

Our theoretical model focuses on a bargaining channel: higher wages in an outside-occupation job option lead to higher wages in the worker's own occupation because of more bargaining leverage. However, there is another channel by which outside-occupation job options can affect wages: mobility. As the wages in an outside option occupation p rise, some workers from initial occupation o will move to occupation p. The supply of workers in occupation o falls, and so the wage rises.

results.

Note that *both* of these mechanisms imply that the alternative occupation p is a relevant outside option for workers in occupation o: the difference is simply that in the bargaining case, workers don't have to exercise their option to move.

Our results in Table 5 provide evidence consistent with the presence of such a mobility channel. Columns (1) and (2) show that, in years where the wages of outside-option occupations rise in a given city, the employment share of the initial occupation o falls (perhaps because workers move on net to the improving outside occupations). However, columns (3) through (6) show that the bargaining channel matters even when taking into account mobility. The coefficients for the effect of the outside-occupation options on the wage are still large, positive and significant - even when controlling for the employment share of the original occupation.

Labor market concentration and outside-occupation options

Recent research on wages and monopsony power has often adopted the "market definition approach" common in antitrust policy, which considers the relevant market of substitutable jobs and excludes all other jobs from the analysis. Most papers define the relevant labor market as an occupation or industry within a geographic area.[37] Our theoretical model suggests that the market definition approach, by failing to consider workers' job options outside their occupation (or industry, or firm cluster), likely obscures substantial heterogeneity in the wage-concentration relationship and may bias the size of the estimated coefficient.

To test our model predictions, we first regress log wages on log vacancy HHIs at the occupation-city level over 2013–2016, with city-by-year and occupation-by-

[37]For instance, a large, negative and significant relationship between wages and employer concentration is shown by Azar et al. (2017) in online vacancy data within an SOC occupation group (6-digit or 4-digit), commuting zone and quarter; by Benmelech et al. (2018) using employment HHIs at a 3- or 4-digit SIC code level for county-industry-year cells over three decades, using establishment-level data from the Census of Manufacturing; and by Rinz (2018) and Lipsius (2018) using HHIs by industry and geography (Commuting Zone, and MSA, respectively) using Longitudinal Business Database data for the entire US. Considering a broader set of affected outcomes, Hershbein et al. (2019) show that employment HHIs at the industry-CZ and vacancy HHIs at the occupation-CZ level are negatively related to wages, and that firms in concentrated labor markets demand higher skills in their job postings. Jarosch et al. (2019) define a new labor market concentration index and estimate it for labor markets defined as as clusters of firms, rather than by industry/occupation/geography; they use worker flows to identify the clusters, as in Nimczik (2018).

year fixed effects,[38] segmenting our data into four quartiles of occupations' average "leave share" (our proxy for outward mobility) and with and without controlling for our instrumented outside-occupation job options.[39] Results are shown in Panels A and B of Table 6. For the full sample, without controlling for outside-occupation options, there is a negative and significant relationship between average wages and vacancy concentration (as found in previous research).[40]

When we control for outside-occupation options, however, the coefficient on the HHI falls statistically significantly by about a third, consistent with omitted variable bias (in our data, the vacancy HHI is strongly negatively correlated with workers' outside-occupation options, as workers in occupation-city labor markets with worse options *within* their occupation also have worse options *outside* their occupation – illustrated in appendix Figure A12).

Our results using the GIV instrument for concentration are shown in Panels C and D of Table 6. The coefficients on the instrumented HHI are statistically significant and more than twice as large as in the OLS specifications. Comparing panels C and D, we can see that controlling for the instrumented outside options index reduced the estimated causal effects of concentration by about 16%.

When segmenting the data by quartile of leave shares, we again find substantial heterogeneity: The concentration effect more than triples in magnitude from -0.016 for the most outwardly mobile quartile of occupations to -0.055 for the least mobile – as one would expect if these workers are less able to switch occupations.[41]

This heterogeneity analysis suggests that, considered alone, the HHI defined at the level of a 6-digit SOC occupation is not an appropriate measure of labor

[38]We calculate the HHI for each occupation-city unit for each year. The timespan of our HHI data is too short to analyze meaningful changes over time in the HHI within a city-occupation unit.

[39]The results using the simple outside-occupation options index without instrumenting look very similar in size and significance but may suffer from omitted variable bias in the coefficient on $log(oo^{occs}_{o,k,t})$ and are therefore omitted here.

[40]The estimated average relationship of labor market concentration and log wages is similar to that found in other studies: our baseline coefficient estimate with city-year and occupation-year fixed effects of -0.014, which is the same as Hershbein et al. (2019). The estimates in other papers are not directly comparable: Rinz (2018) uses employer concentration measured at the 3- or 4-digit industry level, Azar et al. (2017) use a measure of wages based on posted wages, and Azar et al. (2018) estimate wage-concentration relationships for national industry averages.

[41]These results are visualized in appendix Figure A11.

market concentration. To illustrate this point, compare the examples of pharmacists and amusement and recreation attendants in Cape Coral-Fort Myers (Florida). In this metro area, both occupations have an HHI of around 0.25, which by product market guidelines would be considered to be highly concentrated (Marinescu and Hovenkamp, 2019). Pharmacists, however, have an occupation leave share of only 9%, whereas amusement and recreation attendants have an occupation leave share of 27%. The within-occupation HHI suggests that both groups face similar labor market concentration, but in reality, the amusement and recreation attendants in in Cape Coral-Fort Myers are likely to face a much less concentrated labor market because they have more job options outside their occupation.

To understand the size of these concentration effect estimates, note that our results suggest that, going from the 25th to the 75th percentile of concentration across all occupations (from an HHI of 0.047 to an HHI of 0.28) would be associated with a 4.8 log points lower wage.[42] The magnitude of our baseline HHI effect size is larger than but relatively comparable to Arnold (2020).[43] For the quartile of occupations with the lowest outward mobility, an increase in the HHI of this size would be associated with a 9.8 log points lower wage, compared to 2.9 log points for the most mobile occupations.[44]

To put these estimated effect sizes – for both outside-occupation options and the HHI – in context, we can compare them to the average interquartile range of log wages within a given occupation and across cities, which was 23 log points in 2016. Our baseline coefficient estimates (Table 6D) would imply that 21% of the regional variation in wages within a given occupation can be attributed to differences in local labor market concentration within that occupation, and a further 13% of the regional variation in wages within a given occupation can be attributed to differences in

[42]The difference between the 25th (HHI = 0.047) and 95th (HHI = 1) percentile of concentration of all occupation-city cells would be associated with a 8.3 log points lower wage. We can also consider variation within occupations: in 2016, the average interquartile range of log HHIs within an occupation and across cities was 1.38, and a concentration change of that size would be associated with a wage decline of 3.7 log points.

[43]Arnold (2020) finds that for the average worker, concentration reduces wages by around 5% relative to a zero-concentration counterfactual. Our coefficient estimate suggests that moving from the level of concentration faced by the 1st percentile of workers to the average worker reduces wages by 8.2 log points.

[44]We show a number of robustness checks in appendix Table **??**.

the quality of local outside-occupation job options – as opposed to, for example, differential productivity in different cities.[45]

Interpreting the coefficients through the lens of the model

We can use the Nash bargaining structure in our model to estimate an implied value for the degree of worker bargaining power β which would be consistent with the regression coefficients from our empirical analysis (which estimate the responsiveness of wages to outside options).

Our baseline estimated coefficients $\hat{\gamma}_1$ and $\hat{\gamma}_2$ – the elasticity of wages with regard to outside-occupation options and concentration (see panel D in Table 6) – can be expressed as a function of the Nash bargaining parameter β and of the characteristics of a reference city-occupation cell (as shown in equation 14).

We start by rearranging the expression for γ_2 to obtain an expression for the implied Nash bargaining parameter β, of the form

$$\hat{\beta} = \frac{\hat{\gamma}_2(1+\pi_{o\to o}) + \pi_{o\to o}\overline{HHI}_o(1-\hat{\gamma}_2)}{\hat{\gamma}_2\pi_{o\to o} + \pi_{o\to o}\overline{HHI}_o(1-\hat{\gamma}_2)}.$$

Assuming that the reference city has characteristics of the median unit across occupation-city labor markets in 2016, shown in Table 1, we can plug in the values of $\pi_{o\to o}$ (0.76) and $\overline{HHI}_{o,t}$ (0.12), alongside our estimated $\hat{\gamma}_2$ (-0.027). This would be consistent with a $\hat{\beta}$ of 0.63.

Similarly, we can rearrange the structural equation for the outside-occupation options elasticity $\hat{\gamma}_1$ to obtain

$$\hat{\beta} = \frac{(1-\hat{\gamma}_1)\overline{oo}_o^{occs}}{\hat{\gamma}_1\overline{MPL}_o + (1-\hat{\gamma}_1)\overline{oo}_o^{occs}}.$$

Then, we can solve for an estimate $\hat{\beta}$ from this expression using our estimate of the elasticity of the wage with respect to outside-occupation options, $\hat{\gamma}_1$ (0.065), together with the year 2016 median outside-occupation options $\overline{oo}_o^{occs}$ (4.7) and an additional assumption about the (unobserved) median marginal product $\overline{MPL}_o$.

[45] Figures A13 and A14 in the appendix visualize our regression results for our baseline regression of log wages on log outside-occupation options and log HHI.

If we assume no markup over wages (so wages equal marginal productivity) and simply plug in the median occupation-city wage in 2016 ($20.82), we would get $\hat{\beta} = 0.76$. But in a search and matching model like ours, the average wage will be lower than the average marginal product of labor, so this is an overestimate for β. But even if we assume an average markdown of the wage from the marginal product of 40%,[46] the estimated $\hat{\beta}$ only falls to 0.66.

The instrumental variable approach means that our causal effect estimates are weighted local average treatment effects where the weights reflect the effect size for the city-occupation cells in which the endogenous variable is most strongly affected by the instrument. As a result, the effective weights on different sample units varies between the outside option and concentration effect estimates. The fact that our estimated coefficients imply a relatively small range of plausible β estimates of 0.63–0.76 is therefore reassuring.

It is interesting to note that these estimates for β are very similar to the aggregate labor share. That is, if taken literally these estimates would suggest that labor's share of the rents or quasi-rents created by their employment relationship is similar to labor's share of overall output.

On the other hand, these estimates for β are substantially larger than the typical rent-sharing elasticity estimated across a range of studies (as reviewed in Card et al. (2016)) is in the range of 0.1-0.15. That is, the empirical sensitivity of wages to workers' outside job options – whether within, or outside, workers' own occupations – appears to be smaller than one might expect, given the relatively low empirical sensitivity of wages to within-firm changes in productivity or profits.

It is interesting to view our result in conjunction with Jaeger et al. (2019), who find that wages in Austria are almost totally insensitive to changes in unemployment benefits. While our measured sensitivity of wages to outside options may be smaller than the estimates from rent-sharing papers would predict, it is still substantially *greater* than the wage sensitivity that Jaeger et al. (2019) find for unemployment benefits: they estimate that a $1 increase in unemployment benefits is associated with a $0.01 increase in wages, whereas our β estimates would suggest

[46]That is, $\overline{MPL}_o = \frac{1}{1-0.4}\overline{w}_o$. Note that is on the higher end of plausible markdown estimates for the U.S., see, e.g. Sokolova and Sorensen (2020)

that a \$1 increase in the value of outside job options would be associated with a \$0.24-\$0.37 increase in wages. But the results in both papers suggest that workers' wages are substantially *less* sensitive to the value of outside labor market conditions than would be implied by rent-sharing elasticities of 0.1-0.15.

6 Conclusion

In this paper, we (1) document new facts about occupational mobility, showing that existing occupational definitions are not appropriate measures of workers' labor markets; (2) propose a simple, tractable, and theoretically-consistent way to take into account workers' outside-occupation job options using a 'probabilistic' concept of a labor market; (3) show that outside-occupation job options are important in wage determination; (4) demonstrate that wage-concentration regressions which fail to take into account outside-occupation job options are biased and obscure the fact that the effects of concentration is much higher in occupations with lower outward mobility; and (5) use a new instrumental variable approach to estimate causal effects of concentration on wages, taking into account outside-occupation options.

More generally, our paper suggests that labor market analysts can easily improve upon binary labor market definitions, instead constructing 'probabilistic' labor market measures using data on worker transitions between occupations, industries, or locations. We hope that the tools and insights provided in this paper enable other researchers to use, and improve upon, methods like ours to ensure that the labor markets they are researching are the ones that workers are experiencing.

References

Acemoglu, Daron and David Autor, “Skills, tasks and technologies: Implications for employment and earnings,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 1043–1171.

Arnold, David, “Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes,” Technical Report, Working paper 2020.

Ashenfelter, Orley, Henry Farber, and Michael Ransom, “Modern Models of Monopsony in Labor Markets: A Brief Survey,” *IZA Discussion Paper*, 2013, (4915).

Autor, David H, Frank Levy, and Richard J Murnane, “The skill content of recent technological change: An empirical exploration,” *Quarterly journal of economics*, 2003, *118* (4), 1279–1333.

_ , Lawrence F Katz, and Melissa S Kearney, “The polarization of the US labor market,” *American economic review*, 2006, *96* (2), 189–194.

Azar, José A, Ioana Marinescu, Marshall I Steinbaum, and Bledi Taska, “Concentration in US Labor Markets: Evidence From Online Vacancy Data,” Technical Report, NBER 2018.

Azar, José, Ioana Marinescu, and Marshall I Steinbaum, “Labor Market Concentration,” Technical Report, NBER 2017.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, “Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies,” Technical Report 2019.

Beaudry, Paul, David A Green, and Benjamin Sand, “Does industrial composition matter for wages? A test of search and bargaining theory,” *Econometrica*, 2012, *80* (3), 1063–1104.

Benmelech, Efraim, Nittai Bergman, and Hyunseob Kim, “Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?,” Technical Report, NBER 2018.

Berger, David W, Kyle F Herkenhoff, and Simon Mongey, “Labor Market Power,” Technical Report, NBER 2019.

Boal, William M and Michael R Ransom, “Monopsony in the labor market,” *Journal of Economic Literature*, 1997, *35* (1), 86–112.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, “Quasi-experimental shift-share research designs,” Technical Report, NBER 2018.

Burdett, Kenneth and Dale T. Mortensen, “Search, Layoffs, and Labor Market Equilibrium,” *Journal of Political Economy*, 1980, *88* (4), 652–672.

Caldwell, Sydnee and Nikolaj Harmon, "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks," 2018. Working paper.

— **and Oren Danieli**, "Outside Options in the Labor Market," 2018. Working paper.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2016.

Chodorow-Reich, Gabriel and Johannes Wieland, "Secular labor reallocation and business cycles," *Journal of Political Economy*, forthcoming.

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

Gabaix, Xavier and Ralph Koijen, "Granular Instrumental Variables," Technical Report, Working paper 2020.

Gathmann, Christina and Uta Schönberg, "How general is human capital? A task-based approach," *Journal of Labor Economics*, 2010, *28* (1), 1–49.

Goldin, Claudia, "A grand gender convergence: Its last chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

Hershbein, Brad, Claudia Macaluso, and Chen Yeh, "Concentration in U.S. local labor markets: evidence from vacancy and employment data," *Working Paper*, 2019.

Jaeger, Simon, Benjamin Schoefer, Sammy Young, and Josef Zweimueller, "Wages and the Value of Nonemployment," *Working Paper*, 2019.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin, "Granular Search, Market Structure, and Wages," 2019. Working paper.

Kambourov, Gueorgui and Iourii Manovskii, "Rising occupational and industry mobility in the United States: 1968–97," *International Economic Review*, 2008, *49* (1), 41–79.

— **and** — , "A Cautionary Note on Using (March) Current Population Survey and Panel Study of Income Dynamics Data to Study Worker Mobility," *Macroeconomic Dynamics*, 2013, *17*, 172–194.

Lipsius, Ben, "Labor Market Concentration does not Explain the Falling Labor Share," Technical Report 2018. Working paper.

Macaluso, Claudia, "Skill remoteness and post-layoff labor market outcomes," 2019. Working paper.

Manning, Alan, *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press, 2003.

_ **and Barbara Petrongolo**, "How local are labor markets? Evidence from a spatial job search model," *American Economic Review*, 2017, *107* (10), 2877–2907.

Marinescu, Ioana and Herbert Hovenkamp, "Anticompetitive Mergers in Labor Markets," *Indiana Law Journal*, 2019.

McFadden, D, "Conditional Logit Analysis of Qualitative Choice Behavior," *Frontiers in Econometrics*, 1974, pp. 105–142.

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

Nimczik, Jan Sebastian, "Job Mobility Networks and Endogenous Labor Markets," 2018.

Nosnik, Israel P, Patricia Friedmann, Harris M Nagler, and Caner Z Dinlenc, "Resume fraud: unverifiable publications of urology training program applicants," *The Journal of urology*, 2010, *183* (4), 1520–1523.

Pissarides, Christopher A, *Equilibrium Unemployment Theory*, MIT Press, 2000.

Rinz, Kevin, "Labor Market Concentration, Earnings Inequality, and Earnings Mobility," *U.S. Census Bureau Carra Working Paper Series*, 2018, *2018-10.*

Shaw, Kathryn L, "Occupational change, employer change, and the transferability of skills," *Southern Economic Journal*, 1987, pp. 702–719.

Sloane, Arthur A, "Countering resume fraud within and beyond banking: no excuse for not doing more," *Labor Law Journal*, 1991, *42* (5), 303.

Sokolova, Anna and Todd Sorensen, "Monopsony in Labor Markets: A Meta-Analysis," 2020.

Webber, Douglas A, "Firm market power and the earnings distribution," *Labour Economics*, 2015.

Xu, Ming, "Understanding the Decline in Occupational Mobility," *Working Paper*, 2018.

7 Figures and Tables

Table 1: Summary statistics

Percentile	1	5	10	25	50	75	90	95	99
Number of obs. in the BGT occ. mobility data in '000s, by occ. (2002-2015)									
Obs.	0.6	1.1	1.6	4.9	20.8	112.3	466.8	853.9	3,471.9
Share leaving job and occupation, by occ. (2002-2015)									
Share in diff. job	0.30	0.35	0.37	0.40	0.45	0.52	0.61	0.66	0.74
Share leaving 6d. occ.	0.047	0.062	0.074	0.90	0.10	0.12	0.14	0.18	0.29
Leave share	0.09	0.11	0.14	0.19	0.24	0.28	0.33	0.38	0.69
Share of occupational transitions which cross SOC 2d boundary (2002-2015)									
All occ. transitions	0.55	0.65	0.70	0.79	0.87	0.93	0.97	0.98	1.00
Excl. management	0.40	0.48	0.51	0.59	0.67	0.75	0.80	0.83	0.87
Outside-occupation option index oo^{occs}									
oo^{occs} (2016)	1.3	2.0	2.5	3.4	4.7	6.5	8.8	10.6	16.1
$\frac{oo^{occs}}{wage}$	0.03	0.07	0.09	0.15	0.23	0.34	0.45	0.54	0.78
Occupation-city wages and employment									
Mean hourly wage (2016)	8.94	10.38	11.81	15.17	20.82	30.04	42.7	52.11	87.14
Mean hourly wage (1999)	6.12	6.97	7.83	9.93	13.7	19.4	25.5	29.59	39.95
Employment (2016)	30	40	50	80	190	580	1,770	3,530	13,140
Employment (1999)	3	21	36	82	248	800	2,400	4,675	17,117
Single-occupation HHI									
HHI (2016)	0.004	0.011	0.020	0.047	0.12	0.28	0.556	1	1

Notes: In Panel A, an observation is a person-year unit , that is also observed in the data in the following year. We exclude occupations with <500 observations. In Panel B, the leave share is conditional on leaving a job. Panel C shows the share by origin occupation of all SOC 6-digit occupational coincidences which also span SOC 2-digit boundaries. The second row of Panel C excludes all transitions to and from management occupations (SOC 11). Panel D shows the outside-occupation option index oo^{occs}, defined in section 4. Panel E shows HHI by occupation-CBSA cell. Summary statistics are calculated over all occupation-city-year cells for which we have wage data, an outside-occupation option index, and a vacancy HHI.

•

Table 2: Top five occupations with lowest leave shares and highest leave shares

Initial occupation	Leave share	Employment (2017)	Obs.	Modal new occupation
Dental hygienists	.062	211,600	17,458	Dental assistants
Nurse practitioners	.088	166,280	57,830	Registered nurses
Pharmacists	.09	309,330	121,887	Medical & health svc. mgrs.
Firefighters	.098	319,860	60,039	Emerg. med. tech. & paramed.
Self-enrichment educ. teachers	.1	238,710	169,369	Teachers & instructors, all other
...				
Bill and account collectors	.32	271,700	310,951	Customer service rep.
Tellers	.32	491,150	468,829	Customer service rep.
Machine setters, operators, tenders†	.32	154,860	6,805	Production workers, all other
Telemarketers	.36	189,670	47,409	Customer service rep.
Food servers, nonrestaurant	.45	264,630	13,199	Waiters and waitresses

This table shows the twenty large occupations with the lowest and the highest occupation leave shares - defined as the 1-year horizon probability of no longer working in their current occupation, conditional on leaving their job - in the BGT data over 2002-2015, as well as total national employment in that occupation in 2017 from the OES, the number of occupation-year observations in the BGT data ('obs.') and the most popular occupation that workers who leave the initial occupation move to ('modal new occupation'). Large occupations are defined as those with national employment over 150,000 in 2017 (roughly the 75th percentile of occupations when ranked by nationwide employment). † Full occupation title is "Molding, coremaking, and casting machine setters, operators, and tenders, metal and plastic." See Table A1 for a longer list.

Table 3: 10 thickest occupational transition paths for large occupations

Initial occupation	New occupation	Transition share	Employment (2017)	Obs.
Lic. practical & lic. vocat. nurses	Registered nurses	.3	702,700	254,787
Nurse practitioners	Registered nurses	.23	166,280	57,830
Construction managers	Managers, all other	.19	263,480	917,349
Sales rep., wholesale & mfg., techn. & scientific prod.	Sales rep., whls. and mfg., except techn. & scientific prod.	.19	327,190	198,337
Physicians and surgeons, all other	Medical & health svcs. managers	.19	355,460	59,630
Software dev., systems software	Software developers, applications	.19	394,590	53,322
Legal secretaries	Paralegals and legal assistants	.18	185,870	132,543
Accountants and auditors	Financial managers	.18	1,241,000	1,459,175
Registered nurses	Medical & health svcs. managers	.16	2,906,840	1,427,102
Cost estimators	Managers, all other	.16	210,900	124,646
Human resources specialists	Human resources managers	.16	553,950	2,035,604
Physical therapists	Medical & health svcs. managers	.16	225,420	44,314

This table shows the 'thickest' occupational transition paths from large occupations (defined as those with national employment greater than 150,000 in 2017). The transition share from occupation o to occupation p is defined as the share of all occupation leavers from the initial occupation o who move into that particular new occupation p. Only occupations with at least 500 observations in the BGT data are included, with the last column showing total observations. See Table A2 for a longer list.

Table 4: Two-stage least squares regression of wage on outside-occupation option index

Dependent variable:	Log wage			
	(1)	(2)	(3)	(4)
Panel A: OLS				
oo^{occs}	0.140***	0.082***	0.095***	0.044***
	(0.010)	(0.004)	(0.005)	(0.006)
Panel B: 2SLS				
oo^{occs}, instrumented	0.122***	0.070***	0.076***	0.030***
	(0.010)	(0.005)	(0.006)	(0.006)
First stage:				
Coeff. on $oo^{occs,inst}_{o,k,t}$	0.997***	0.868***	0.808***	1.010***
	(0.015)	(0.020)	(0.012)	(0.054)
1st-stage F-Stat.	2122	1838	4921	348
Fixed effects	Year	Occ-Year City	City-Year Occ-Year	Occ-Year Occ-City
Observations	1,944,370	1,944,370	1,944,370	1,944,370

* p<0.10, ** p<0.05, *** p<0.01

Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses. Units of observation are 6 digit SOC by city by year, for all observations with available data over 1999–2016 inclusive. The instrumented outside-occupation option index uses the national leave-one-out mean wage in outside option occupations to instrument for the local (city-level) wage, and the initial local employment share in outside option occupations to instrument for the current local employment share.

Table 5: Regressions incorporating employment share

Dependent var.:	Employment share		Log wage			
	(1)	(2)	(3)	(4)	(5)	(6)
oo^{occs}	-0.054**		0.044***	0.043***		
	(0.023)		(0.005)	(0.005)		
oo^{occs} , instr.		-0.226***			0.030***	0.026***
		(0.027)			(0.006)	(0.006)
Empl. share				-0.015***		-0.015***
				(0.001)		(0.001)
Fixed effects	Occ-City Occ-Year	Occ-City Occ-Year	Occ-City Occ-Year	Occ-City Occ-Year	Occ-City Occ-Year	Occ-City Occ-Year
Observations	1,931,901	1,931,901	1,931,901	1,931,901	1,931,901	1,931,901

* p<0.10, ** p<0.05, *** p<0.01

Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses. Units of observation are 6 digit SOC by city by year, for all observations with available data over 1999–2016 inclusive. The instrumented outside-occupation option index uses the national leave-one-out mean wage in outside option occupations to instrument for the local (city-level) wage, and the initial local employment share in outside option occupations to instrument for the current local employment share.

Table 6: Regression of wage on Vacancy HHI and outside options, by quartile of occupation leave share

Dependent variable:	Log wage				
	Full sample	By quartile of leave share			
	All quartiles	Q1	Q2	Q3	Q4
Panel A					
Log vacancy HHI	-0.014***	-0.021***	-0.010***	-0.011***	-0.008***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
Panel B					
Log vacancy HHI	-0.010***	-0.017***	-0.007***	-0.007***	-0.004***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Log outside-occ. options	0.083***	0.098***	0.061***	0.078***	0.105***
instrumented	(0.008)	(0.013)	(0.011)	(0.010)	(0.013)
Panel C					
Log vacancy HHI,	-0.032***	-0.057***	-0.026***	-0.022***	-0.022***
instrumented	(0.005)	(0.010)	(0.008)	(0.008)	(0.007)
Panel D					
Log vacancy HHI,	-0.027***	-0.055***	-0.023***	-0.016*	-0.016**
instrumented	(0.005)	(0.011)	(0.008)	(0.009)	(0.007)
Log outside-occ. options,	0.065***	0.059***	0.044***	0.068***	0.092***
instrumented	(0.009)	(0.016)	(0.013)	(0.014)	(0.015)
Fixed effects	Occ-Year	Occ-Year	Occ-Year	Occ-Year	Occ-Year
	City-Year	City-Year	City-Year	City-Year	City-Year
Observations	182,973	50,256	49,221	47,133	36,330

Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses. Units of observation are 6 digit SOC by city by year, for all observations with available data over 2013–2016 inclusive. Outside options index is instrumented using national leave-one out wages and beginning-of-period employment shares. HHI is instrumented using local exposure to national firm-level vacancy growth, squared (and regressions with instrumented HHIs also include a control for exposure to national firm-level vacancy growth). See text for detailed explanation of instrument. Regressions in Panels B through D are estimated using 2SLS. Regressions in Panels C and D control for total vacancy growth in the occupation-city cell. Occupations are split into quartiles by the average occupation leave share in the Burning Glass Technologies resume data (averaged over 2002–2015). Significance levels: * p<0.10, ** p<0.05, *** p<0.01.

Figure 1: Coefficients and 95% confidence intervals from regression of 2002-2015 average probability of moving into another occupation (conditional on any job move) on occupational characteristics of the form $\pi_{o \to p} = \alpha_o + \beta f(X_{occ\ o \to p}) + \gamma f(\Delta w_{o \to p}) + \epsilon_{op}$ where the function $f(\cdot)$ represents the absolute (left side) or relative difference (right side, target minus origin) in its argument when moving from occupation o to p, and α_o is an origin occupation fixed effect. All regressions also include a constant, and absolute (top panel) or relative (bottom panel) avg. hourly wage differences - except for the amenities regressions, where wage differences are omitted. Standard errors are clustered at the origin occupation level.

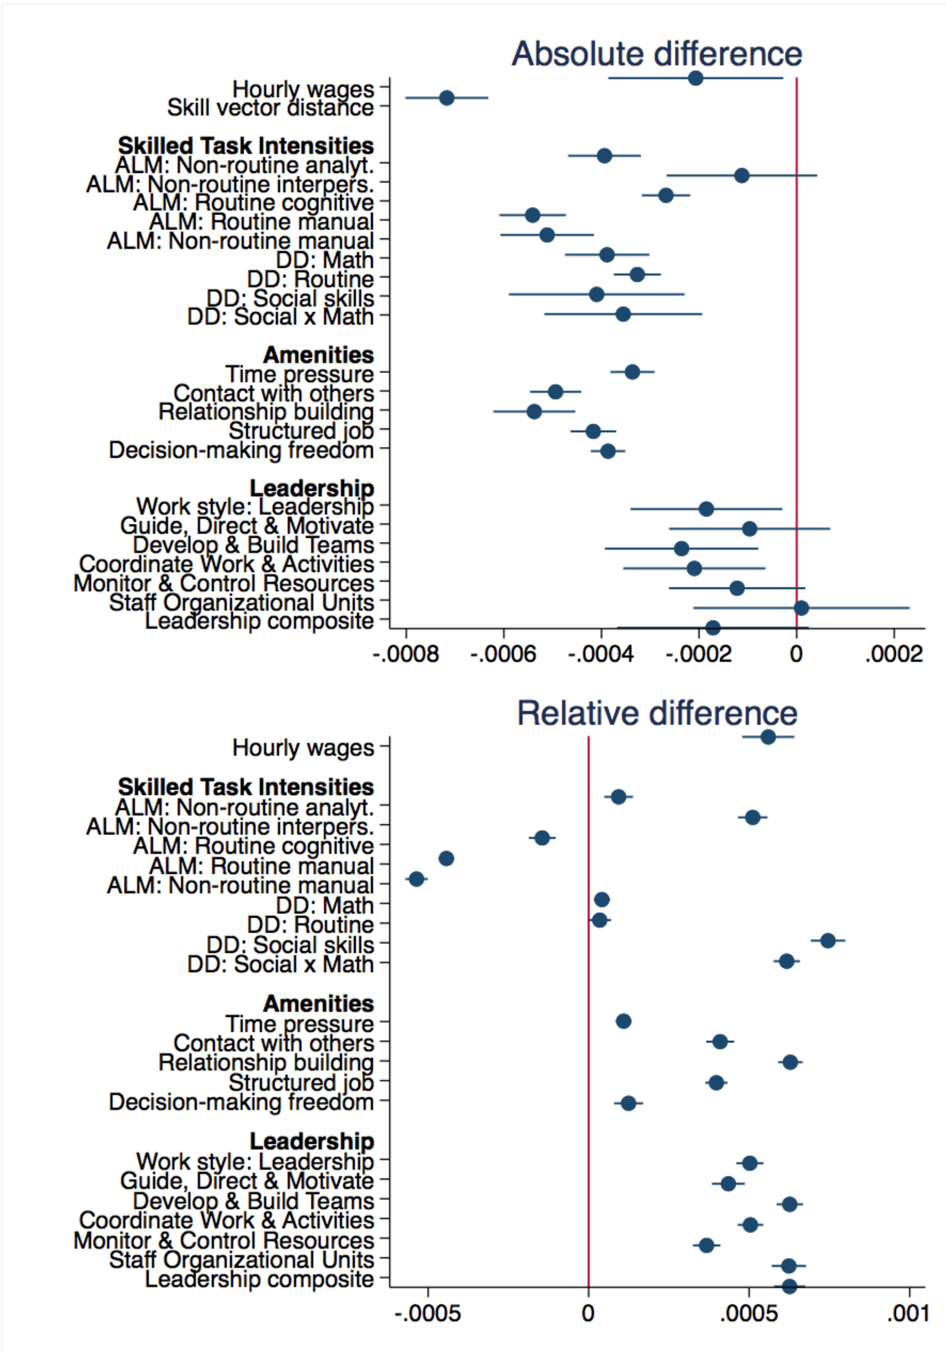

For Online Publication: APPENDIX

A Burning Glass Technologies Resume Data

This section contains further information about our resume data set from Burning Glass Technologies ("BGT"), discussed more briefly in Section 3. This is a new proprietary data set of 16 million unique U.S. resumes spanning years over 2002–2018.

Resumes were sourced from a variety of BGT partners, including recruitment and staffing agencies, workforce agencies, and job boards. Since we have all data that people have listed on their resumes, we are able to observe individual workers' job histories and education up until the point where they submit their resume, effectively making it a longitudinal data set.

A.1 Data cleaning and transition data construction

We apply a number of different filters to the Burning Glass resume data before calculating our occupational mobility matrices: First, we retain only resumes that are from the U.S. Next, we keep only jobs on these resumes that last for longer than 6 months to ensure that we are only capturing actual jobs rather than short-term internships, workshops etc. We also apply a number of filters to minimize the potential for mis-parsed jobs, by eliminating all jobs that lasted longer than 70 years. Moreover, we impute the ages of workers based on their first job start date and education and limit our sample to resumes submitted by workers between the ages of 16 and 100. As we are interested in occupational transitions during the last two decades, we then restrict the data set to jobs held after 2001. The final number of resumes that contain at least two years of job data under these restrictions is 15.8 million. The main job information retained for each resume are the occupation and duration of each job held.

For each of these resumes, we start by extracting separate observations for each occupation that the worker was observed in, in each year. These observations are then matched to all other occupation-year observations on the same resume. We retain all matches that are in sequential years - either in the same occupation or in different occupations. For instance, if a worker was a Purchasing Manager in the period 2003-2005, and a Compliance Officer in 2005-2007, we would record 1-year horizon sequential occupation patterns of the form shown in Table 7.

Table 7: Illustrative example of sequential job holding data.

Year:	2004	2005	2006
Current Occ.	*1-Year Horizon Occ.*		
Purchasing Mgr. (11-3061)	11-3061 13-1040		
Compliance Off. (13-1040)		13-1040	13-1040

In our data we have 80.2 million job observations. This results in 178.5 million observations of year-to-year occupation coincidences (including year-to-year pairs where workers are observed in the same occupation in both years). Below, we describe the characteristics of this data and how it compares to other data sets - with all statistics referring to this final set of filtered sequence observations, or the 15.8 million resumes, unless otherwise noted.

We use these occupation coincidence pairs to construct our measures of occupational mobility as follows. For each pair of (different) occupations o to p, we count the total number of year-to-year occupation coincidence pairs where the worker is observed in occupation o at any point in year t and is observed in occupation p at any point in year $t + 1$. We then divide this by the total number of workers in occupation o in year t who are still observed in the sample in the following year $t + 1$.

Since our data is not fully representative on age within occupations, we compute these occupation transition shares separately for different age categories (24 and under, 25 to 34, 35 to 44, 45 to 54, and 55 and over). We then aggregate them, reweighting by the average proportion of employment in each of these age categories in that occupation in the U.S. labor force over 2012–2017 (from the BLS Occupational Employment Statistics). Our aggregate occupational mobility matrix has therefore been reweighted to correspond to the empirical within-occupation age distribution in the labor force, eliminating any potential bias from the skewed age distribution of our sample.

A.2 Summary statistics

Below, we describe the characteristics of this data and how it compares to other data sets. All statistics referring to the final set of 15.8 million filtered resumes, or 178.5 million observations of year-to-year occupation coincidences ('observations') from these resumes, unless otherwise noted.

Job number and duration: The median number of jobs on a resume is 4, and more than 95% of the resumes list 10 or fewer jobs (note that a change of job under our definition could include a change of job title or occupation under the same employer). The median length job was 2 years, with the 25th percentile just under 1 year and the 75th percentile 4 years. The median span of years we observe on a resume (from date started first job to date ended last job) is 12 years. Table 8 shows more information on the distribution of job incidences and job durations on our resumes.

Table 8: Distribution of number of jobs on resume and duration of jobs in BGT data set.

Percentile	10th	25th	50th	75th	90th
# Jobs on resume	2	3	4	6	9
Job duration (months)	4	12	24	48	98

Gender: BGT imputes gender to the resumes using a probabilistic algorithm based on the names of those submitting the resumes. Of our observations, 88% are on resumes where BGT was able to impute a gender probabilistically. According to this imputation, precisely 50% of our observations are imputed to come from males and 50% are more likely to be female. This suggests that relative to the employed labor force, women are very slightly over-represented in our data. According to the BLS, 46.9% of employed people were women in 2018.

Education: 141.3 million of our observations are on resumes containing some information about education. The breakdown of education in our data for these data points is as follows: the highest educational level is postgraduate for 25%, bachelor's degree for 48%, some college for 19%, high school for 8% and below high school for less than 1%. This substantially overrepresents bachelor's degree-holders and post-college qualifications: only 40% of the labor force in 2017 had a bachelor's degree or higher according to the BLS, compared to 73% in this sample (full comparisons to the labor force are shown in Figure A1). It is to be expected that the sample of the resumes which *provide* educational information are biased towards those with tertiary qualifications, because it is uncommon to put high school on a resume. Imputing high school only education for all resumes which are missing educational information substantially reduces the overrepresentation of those with a BA and higher: by this metric, only 58% of the BGT sample have a bachelor's degree or higher. This remains an overrepresentation, but this is to be expected: a sample drawn from online resume submissions is likely to draw a more highly-educated population than the national labor force average both because many jobs requiring little formal education also do not require online applications, and because we expect online applications to be used more heavily by younger workers, who on average have more formal education. As long as we have enough data to compute mobility patterns for each occupation, and workers of different education levels *within* occupations do not have substantially different mobility patterns, this should therefore not be a reason for concern.

Age: We impute individuals' birth year from their educational information and from the date they started their first job which was longer than 6 months (to exclude internships and temporary jobs). Specifically, we calculate the imputed birth year as the year when a worker started their first job, minus the number of years the worker's maximum educational qualification requires, minus 6 years. High school is assumed to require 12 years, BA 16 years, etc. For those who do not list any educational qualification on their resume, we impute that they have high school only, i.e. 12 years of education. Since we effectively observe these individuals longitudinally - over the entire period covered in their resume - we impute their age for each year covered in their resume.

As a representativeness check, we compared the imputed age of the people corresponding to our 2002-2018 sample of sequential job observations in the BGT sample to the age distribution of the labor force in 2018, as computed by the BLS. The BGT data of job observations substantially overrepresents workers between 25 and 40 and underrepresents the other groups, particularly workers over 55. 55% of observations in the BGT sample would have been for workers 25-40 in 2017, compared to 33% of the US labor force - see Figure A2 for the full distribution. One would expect a sample drawn from online resume submissions to overweight younger workers for three reasons: (1) because younger workers may be more familiar with and likely to use online application systems, (2) because older workers are less likely to switch jobs than younger workers, and (3) because the method for job search for more experienced (older) workers is more likely to be through direct recruitment or networks rather than online applications. Moreover, by the nature of a longitudinal work history sample, young observations will be overweighted, as older workers will include work experiences when they are young on their resumes, whereas younger workers, of course, will never be able to include work experiences when they are old on their current resumes. Therefore, even if the distribution of resumes was not skewed in its age distribution, the sample of job observations would still skew younger.

As noted above, we directly address this issue by computing occupational mobility only after reweighting observations to adjust the relative prevalence of different ages in our sample relative to the labor force. For instance, this means that we overweight our observations for 45-49 year olds, as this age category is underrepresented in our sample relative to the labor force.

Occupation: The BGT automatic resume parser imputes the 6-digit SOC occupation for each job in the dataset, based on the job title. Of 178.5 million useable observations in the data set, 169.6 million could be coded into non-military 6-digit SOC occupations by the BGT parser. 833 of the 840 6-digit SOC occupations are present, some with few observations and some with very many. Ranking occupations by the number of observations,[47] the 10th percentile is 1,226 observations, 25th percentile is 4,173, the median is 20,526, 75th percentile is 117,538, and the 90th percentile is 495,699. We observe 216 occupations with more than 100,000 observations, 83 occupations with more than 500,000 observations, and 19 occupations with more than 2 million observations.[48]

[47] As defined above, for our purposes, an observation is a person-occupation-year observation for which we also observe another occupation in the following year: i.e. the start of a year-to-year occupation coincidence sequence.

[48] The occupations with more than 2 million observations are: General and Operations Managers; Sales Managers;

Figure A3 compares the prevalence of occupations at the 2-digit SOC level in our BGT data to the share of employment in that occupation group in the labor force according to the BLS in 2017. As the figure shows, at a 2-digit SOC level, management occupations, business and finance, and computer-related occupations are substantially overweight in the BGT data relative to the labor force overall, while manual occupations, healthcare and education are substantially underrepresented. However, this does not bias our results, as we compute mobility at the occupation-level.

Location: Since not all workers list the location where they work at their current job, we assign workers a location based on the address they list at the top of their resume. 115.4 million of our observations come from resumes that list an address in the 50 U.S. states or District of Columbia. Comparing the proportion of our data from different U.S. states to the proportion of workers in different U.S. states in the BLS OES data, we find that our data is broadly representative by geography. As shown in figure A4, New Jersey, Maryland and Delaware, for instance, are 1.5-2x as prevalent in our data as they are in the overall U.S. labor force (probably partly because our identification of location is based on residence and the BLS OES data is based on workplace), while Nebraska, Montana, South Dakota, Alaska, Idaho and Wyoming are less than half as prevalent in our data as they are in the overall U.S. labor force. However, the figure also suggests that the broad patterns of the demographic distribution of populations across the U.S. is reflected in our sample. Aggregating the state data to the Census region level, the Northeast, Midwest, South, and West regions represent 24%, 22%, 38%, and 16% of our BGT sample, while they constitute 18%, 22%, 37%, and 24% of the BLS labor force. This shows that our sample is very close to representative for the Midwest and South regions, and somewhat overweights the Northeast, while underweighting workers from the West region.

A.3 Advantages over other datasets

As a large, nationally-representative sample with information about labor market history over the past year, the Current Population Survey is often used to study annual occupational mobility. Kambourov and Manovskii (2013) argue however that the CPS should be used with caution to study occupational mobility. First, the coding is often characterized by substantial measurement error. This is particularly a concern for measuring mobility from one year to the next, as independent coding is often used when there are changes in employers, changes in duties, or proxy responses, and this raises the likelihood of an occupational switch being incorrectly identified when in fact

Managers, All Other; Human Resources Specialists; Management Analysts; Software Developers, Applications; Computer User Support Specialists; Computer Occupations, All Other; First-Line Supervisors of Retail Sales Workers; Retail Salespersons; Sales Representatives, Wholesale and Manufacturing, Except Technical and Scientific Products; First-Line Supervisors of Office and Administrative Support Workers; Customer Service Representatives; Secretaries and Administrative Assistants, Except Legal, Medical, and Executive; Office Clerks, General; Heavy and Tractor-Trailer Truck Drivers; Financial Managers; Food Service Managers; Medical and Health Services Managers.

the occupation remained the same.

Due to its structure, the CPS is also only able to identify occupational mobility at an annual or shorter frequency. The PSID is another data source frequently used to study occupational mobility. As a truly longitudinal dataset it is able to capture truly annual mobility (or mobility over longer horizons), but its small sample size means that it is unable to provide a more granular picture of mobility between different pairs of occupations.

The BGT dataset allows us to circumvent some of these concerns. Its key advantage is its sample size: with 16 million resumes (after our parsing) covering over 80 million job-year observations, we are able to observe a very large number of job transitions and therefore also to observe a very large number of transitions between different pairs of occupations. Our sample of job-year observations is more than an order of magnitude larger than that which would be available from the CPS when pooling over the same time period we use (2002–2018). In addition, since individuals list the dates they worked in specific jobs on their resumes, we are able to observe occupational transitions at the desired frequency, whether that is annual or longer.[49] And individuals listing their own jobs means that there is less of a risk of independent coding falsely identifying an occupational switch when none occurred. In addition, the length of many work histories in the data allows for inferring a broader range of latent occupational similarities by seeing the same individual work across different occupations, even when the jobs are decades apart.

A.4 Caveats and concerns

The BGT dataset does, however, have other features which should be noted as caveats to the analysis.

1/ Sample selection: There are three areas of concern over sample selection: first, our data is likely to over-sample people who are more mobile between jobs, as the data is collected only when people apply for jobs; second, our data is likely to over-sample the types of people who are likely to apply for jobs online rather than through other means; and third, our data is likely to over-sample the types of people who apply for the types of jobs which are listed through online applications.

2/ Individuals choose what to put on their resume: We only observe whatever individuals have chosen to put on their resume. To the extent that people try to present the best possible picture of their education and employment history, and even sometimes lie, we may not observe certain jobs or education histories, and we may be more likely to observe "good" jobs and education histories than "bad" ones. The implication of this concern for our measure of job opportunities depends on the exact nature of this distortion. If workers generally inflate the level of occupation that they

[49] Since many individuals list only the year in which they started or ended a job, rather than the specific date, measuring transitions at a sub-annual frequency is too noisy.

worked at, this would not necessarily distort our estimates of job transitions systematically, unless transition probabilities across occupations vary systematically with the social status / level of otherwise similar jobs. At the same time, if workers choose to highlight the consistency of their experiences by describing their jobs as more similar than they truly were, we may underestimate the ability of workers to transition across occupations. Conversely, if workers exaggerate the breadth of their experience, the occupational range of transitions would be overestimated. In any case, this issue is only likely to be significant, if these types of distortions exist for many observed workers, do not cancel out, and differ systematically between workers in different occupations.

We are only aware of a very limited number of studies directly trying to estimate the incidence of misrepresentations on resumes. For instance, Sloane (1991) surveys HR exexcutives in banking and finds that 51 responding executives are jointly aware of a total of 17 incidences of meaningfully falsified job titles, which, given the presumably large number of resumes and contained job listings that would have been processed under these executives seems small. All but one of the respondents estimated the incidence of falsification of *any* part of the resume to be below 20%, with most opting for lower estimates. Note that this study was done before online search made verification of basic resume information much faster and more affordable. More recently, Nosnik et al. (2010) found that 7% of the publications listed by a sample of urology residency applicants on their resumes could not be verified.

While such low rates of misrepresentation seem unlikely to introduce systematic bias into our data, it is also important to keep in mind that we are trying to estimate the *plausibility* in a bargaining setting of other jobs constituting relevant outside options. If the skills of a job that they haven't actually held are plausibly consistent with *other* jobs on their resume in the eyes of jobseekers - and ultimately of employers - then this still constitutes evidence that these jobs are perceived as pertaining to the same labor market.

3/ Parsing error: Given the size of the dataset, BGT relies on an algorithmic parser to extract data on job titles, firms, occupations, education and time periods in different jobs and in education. Since there are not always standard procedures for listing job titles, education, dates etc. on resumes, some parsing error is likely to exist in the data. For example, the database states that 25,000 resumes list the end date of the most recent job as 1900.

4/ Possible duplicates: The resume data is collected from online job applications. If a worker over the course of her career has submitted multiple online job applications, it is possible that her resume appears twice in the raw database. BGT deduplicates the resume data based on matching name and address on the resume, but it is possible that there are people who have changed address between job applications. In these cases, we may observe the career history of the same person more than once in the data. Preliminary checks suggest that this is unlikely to be a major issue.

B OES Occupational Code Crosswalk

To construct our data set of wages and employment at the occupation-CBSA level over 1999 - 2016 - as described in Section 3 - we need to create a crosswalk for OES occupational codes from SOC 2000 to SOC 2010.

We start from the crosswalk provided by the BLS for matching occupation codes. The crosswalk is based on an exact match if a SOC 2000 code corresponds to exactly one SOC 2010 code.

When SOC 2000 codes map into multiple SOC 2010 codes,or vice versa, we create a probabilistic mapping. This mapping is based on relative employment shares between the target occupation codes as of 2009 and 2012, obtained at a national level from the BLS.

When one SOC 2000 code splits into multiple SOC 2010 codes, its employees are split based on the relative employment shares in the resulting SOC 2010 codes as of 2012.

When there are multiple SOC 2000 codes mapping into multiple SOC 2010 codes, the number of employees in 2009 and 2012 are counted for the whole cluster of ambiguous assignments. Then, unique assignments within the cluster are made based on the ratio of total 2012 to 2009 employees in the cluster. The remaining employees are apportioned based on their relative share in the remainder. For 2010 and 2011 numbers, the OES combines data collected under both the old and new classification system, and grouped them under either SOC 2010 codes or hybrid identifiers.[50] Where this combination did not result in ambiguity with regard to the meaning of the SOC 2010 code used, this difference in collection methods was ignored and the content of the OES 2010 code transferred one-to-one into the applicable SOC 2010.[51]

Where the OES 2010 code is more aggregated than the SOC 2010 code, it was split based on 2012 employment shares in the target codes.[52]

Using these occupational crosswalks, we can stack the OES occupational employment and wage data by city provided by the BLS, creating an unbalanced panel of 2.5 million occupation-by-city-by-year data points of employment and mean hourly and annual wages for the years 1999-2016. Out of these data, the 2005-2016 panel that was originally coded by the BLS using the new CBSA definition has about 1.8 million data points.

[50]Detailed breakdown of the affected codes available at: `https://www.bls.gov/oes2010_and_2011_oes_classification.xls`

[51]This was the case for the following OES 2010 codes: 11-9013, 15-1799, 51-9151

[52]This was the case for the following OES 2010 codes: 13-1078, 15-1150, 15-1179, 21-1798, 25-2041, 25-3999, 29-1111, 29-1128, 29-2037, 29-2799, 31-1012, 31-9799, 39-4831, 41-9799, 43-9799, 47-4799, 49-9799, 51-9399.

C Alternative approaches to estimating occupational similarity

In Section 3 of this paper, we argue that using occupational transitions is a better way to identify workers' outside options than the two other commonly-used methods of estimating occupational similarity: skill-and task-based similarity measures, and demographic- and qualification-based similarity measures. We expand on this argument below.

Skill- and task-based occupational similarity measures define two occupations as more similar, the more similar the skills and tasks are that they require.[53] While this captures many dimensions of the feasibility of an occupational transition, a skill- or task-based measure has several weaknesses relative to a transition-based measure. First it cannot capture non-skill-related aspects which affect the availability of outside options, such as occupational licensing barriers. Second, it cannot capture the desirability of moving from one occupation to another: it may be that two occupations are very similar in terms of the skills and tasks that they require, but the amenities may differ – so that the kind of people that work in one occupation may not want to work in the other. Third, skill- or task-based similarity measures are (usually) symmetric between occupation pairs, whereas transitions data can capture the asymmetry of the value of different occupations as outside options for each other: occupation p may be a relevant outside option for occupation o but not the other way around, perhaps because of generalist/specialist skill differentials, differences in job hierarchy or status, or specific requirements for experience, training or certification. And fourth, skill- or task-based similarity measures require substantial assumptions as to how skill and task data should be combined to create a similarity measure, and the measures can be highly sensitive to these assumptions. (In contrast, a transition-based measure has the advantage of being non-parametric).[54]

Demographic- and qualification-based occupational similarity measures define two occupations as more similar, the more similar are their workers based on their observable demographic and educational characteristics.[55] This type of measure can capture occupational similarity in terms of the skills required, based on workers' inherent characteristics and education/training, and in terms of preferences determined by these factors. It also has the advantage of requiring substantially fewer assumptions than a skill- and task-based measure, since it uses workers' actual labor market choices to reveal their outside options. Since it does not consider career paths, however, a demographic- and qualification-based occupational similarity measure cannot capture the role of occupation-specific experience and learning, or obstacles to occupational transitions, in determin-

[53] For example, Macaluso (2019) measures occupational skill similarity using the vector difference of occupational skill content, and Gathmann and Schönberg (2010) use the angular separation of occupations' task content vectors.

[54] This allows us to capture the equilibrium job choice policy function without having to impose a particular model of how workers and firms choose to offer and accept jobs, or about equilibrium play.

[55] A simplified version of the approach used by Caldwell and Danieli (2018), who probabilistically identify workers' outside options using the distribution of other similar workers across jobs and locations.

ing future employment options. In that sense, a demographic- and qualification-based measure of occupational similarity can be thought of as a static approach to defining a 'revealed' labor market, whereas a transition-based measure can be thought of as a dynamic approach. In addition, as with skill- and task-based approaches, this approach in practice requires assumptions on which observables are relevant for job choices and parametric assumptions on the functional form of the choice function.

Our transitions-based measure does have a major potential drawback relative to a skill- or task-based measure: off-equilibrium outside options are not observed if bargaining is efficient. It may be the case that another occupation is very feasible but slightly less desirable, which makes it a relevant outside option for a worker but one that is rarely exercised in equilibrium. However, if the number of workers and firms is large enough to observe rare transitions, worker preferences are continuous, and idiosyncratic shocks have enough variance to induce many workers to change occupations, these off-equilibrium options will on average still be revealed by the transition data - and we believe these conditions hold for job transitions.[56]

D Determinants of occupational mobility

This section expands on the results on determinants of occupational mobility discussed more concisely in section 3. To use worker transitions to infer the network of worker outside options, we must assume that the empirical occupational transitions we observe reflect the underlying feasibility and/or desirability of an occupation as an outside option. It is possible, however, that our empirical occupational transition shares simply reflect something idiosyncratic in our data, or short-run contractions or expansions of different occupations (though the size, time horizon, and relative representativeness of our data should do something to assuage these concerns). In this section, we describe in more detail our analysis on the degree to which our measure of occupational transitions reflects similarities between different occupations in terms of task requirements, wages, amenities, and leadership responsibilities.

[56] More specifically, there are three conditions under which the above concern about off-equilibrium options in the 'revealed labor market' approach based on observed occupational transitions is not significant. First, there is a continuous distribution of worker heterogeneity with regard to preferences over different firms, and so any given worker's closest outside options (off-equilibrium option) are revealed by the actual equilibrium paths of similar workers (similar to the way that choice probabilities map to expected value functions in discrete choice models with i.i.d. preference shocks (McFadden, 1974)). Second, there has to be a sufficient number of similar workers and firms to observe these transitions. Third, that the only *relevant* off-equilibrium outside options for workers in the wage bargaining process are those which are quite similar to their existing job or skill set in expected match quality (i.e. that cashier jobs are not relevant outside options for engineers), such that the variance of worker preferences beyond the expected match quality is large enough to manifest in different job matches for all relevant outside options. If these conditions are satisfied, the expected relevant off-equilibrium options for workers in a given occupation can be inferred by the equilibrium choices of other workers in the same occupation.

D.1 Occupation characteristics: measures

Task requirements. To measure occupational similarity in terms of tasks required, we use two different approaches from prior literature.

First, we use the vector difference between the importance scores for "Skill" task content items provided by the O*Net database of occupational characteristics, as proposed by Macaluso (2019).[57] Our measure of average task distance $\bar{D}_{op}$ between occupations o and p is defined as:

$$\bar{D}_{op} = \frac{1}{35} \sum_{k=1}^{35} |S_{k,occ\ p} - S_{k,occ\ o}|,$$

where $S_{k,occ\ p}$ is the standardized skill k measure for occupation p.

Second, we use composite task measures from recent literature relating occupational task content to important economic outcomes. We consider six task composites (denoted "ALM") first introduced in Autor et al. (2003) and updated to the most recent O*Net version in Acemoglu and Autor (2011). These composites mainly capture the distinction between cognitive vs. manual and routine vs. non-routine task contents. We also consider a categorization by Deming (2017) (denoted "DD"), which recasts the occupational task composites and also introduces a composite capturing social skill-related task intensity.[58]

Job amenities. We measure similarity in the "temporal flexibility" of different occupations[59] using the 5 O*Net occupation characteristics that Goldin (2014) identifies as proxies for the ability to have flexibility on the job: time pressure, contact with others, establishing and maintaining interpersonal relationships, structured vs. unstructured work, and the freedom to make decisions.[60]

Leadership responsibility. Another reason for observing occupational transitions may be career advancement (which is often reflected in a change of occupation). To study whether this appears in our data, we identify occupational characteristics measuring leadership responsibilities

[57] In our measure, as in Macaluso (2019), dissimilarity is measured as the average difference in importance scores (scaled to lie between zero and ten) across the full set of 35 tasks. For a similar notion of task distance, see (Gathmann and Schönberg, 2010).

[58] We update the task composites from Deming (2017) by using the latest source for task contents on O*Net, and computing the composites at the level of SOC 2010 occupational codes.

[59] These amenities are particularly important because, as Goldin (2014) notes, "certain occupations impose heavy penalties on employees who want fewer hours and more flexible employment" (p. 1106), which in turn may contribute to gender gaps in earnings.

[60] Note that higher scores in each of these domains imply more rigid time demands as a result of business needs and make it less likely that workers are able to step away from their job whenever they need to. The five characteristics correspond the following O*Net survey items: IV.C.3.d.1 - How often does this job require the worker to meet strict deadlines?; IV.C.1.a.4 - How much does this job require the worker to be in contact with others (face-to-face, by telephone, or otherwise) in order to perform it?; IV.A.4.a.4 - Developing constructive and cooperative working relationships with others; IV.C.3.b.8 - To what extent is this job structured for the worker, rather than allowing the worker to determine tasks, priorities, and goals?; IV.C.3.a.4 - Indicate the amount of freedom the worker has to make decisions without supervision.

from the O*Net database, and create a new "leadership" composite measure defined at the level of each SOC 6-digit occupation. The measure incorporates the six characteristics most associated with leadership positions in the O*Net data, alongside the O*Net work style category for leadership.[61]

D.2 Occupational similarity and mobility

To evaluate whether workers are more likely to move to occupations that have similar characteristics to their current occupation, we estimate the following regression:

$$\pi_{o\to p} = \alpha_o + \beta^{abs}|X_{occ\,p} - X_{occ\,o}| + \gamma|\Delta w_{o\to p}| + \epsilon_{op}. \tag{16}$$

where $\pi_{o\to p}$ is the share of job changers in the origin occupation o that move into target occupation p, $|X_{occ\,p} - X_{occ\,o}|$ is the absolute difference between the target and the origin occupation in each of the occupational characteristics X_o defined above, and α_o are origin occupation fixed effects to control for differences in outward mobility across occupations. We control for absolute wage differences between the occupations in all regressions except for those estimating the effect of wages or amenity differences on occupational mobility,[62] but note that the results are qualitatively similar without the wage controls.

We would expect the coefficient on the absolute difference in characteristics to be negative: the greater the difference between two occupations, the less likely we should be to observe the worker moving from one into the other. Our results bear this out: in every regression of pairwise occupational mobility on the absolute difference in characteristics, the coefficients are significantly negative or statistically insignificant, as shown in figure 1.[63]

[61] We used the following algorithm to determine which characteristics measure leadership responsibilities: On the O*Net website, we looked at the work activity characteristics that describe "Interacting with Others". For each of them, we considered the list of top 20 occupations with the highest level of that characteristic and counted how many of them are managerial positions, as evidenced by the words "supervisor", "manager", "director", or equivalents, in the occupation title. We selected all the characteristics for which the share of managerial positions among the top 20 occupations was greater than half, as these characteristics seem to be associated with "leadership" in some sense; we also added the O*Net work style category for leadership. The final list of characteristics contains the following O*Net items: I.C.2.b. - Leadership work style: job requires a willingness to lead, take charge, and offer opinions and direction; IV.A.4.a.2. - Communicating with Supervisors, Peers, or Subordinates; IV.A.4.b.1. - Coordinating the Work and Activities of Others; IV.A.4.b.2. - Developing and Building Teams; IV.A.4.b.4. - Guiding, Directing, and Motivating Subordinates; IV.A.4.c.3. - Monitoring and Controlling Resources; IV.A.4.c.2. - Staffing Organizational Units (We were reassured to note that for 6 of these 7 characteristics, "Chief Executives" are among the Top 20 occupations in terms of importance of this measure.). We use the mean score across these 7 characteristics as our "leadership" composite. All variables are converted into standardized Z-scores before including them in regressions, so coefficients represent the effect of a one standard deviation difference in the characteristic on the outcome variable.

[62] Amenities are most likely to be priced into wages (Goldin, 2014) and controlling for the latter would therefore be inappropriate.

[63] Our findings build on Macaluso (2019), who showed that greater skill distance between SOC 2-digit occupations

The previous results impose symmetry on the likelihood of occupational transitions – but between many pairs of occupations, the probability of moving in one direction is likely to be different than the probability of moving in the other direction. To study whether differences in characteristics also predict the direction of occupational flows, we estimate a similar regression equation to that shown in equation 16, but now using the *relative* (target minus origin) difference in occupational characteristics as the independent variable:

$$\pi_{o \to p} = \alpha_o + \beta^{rel}(X_{occ\,p} - X_{occ\,o}) + \gamma \Delta w_{o \to p} + \epsilon_{op}. \tag{17}$$

Again, we include origin occupation fixed effects and now control for relative wage differences between the occupations in all regressions except for the amenity differences and the wage regression. The β^{rel} coefficients obtained from estimating equation 17 for the different measures are shown in Figure 1.[64]

A number of our predictions are borne out in the data: we find (1) that workers are more likely to move towards jobs with higher wages; (2) that workers transition on average *towards* jobs that require more leadership responsibility - as would be expected from moves up the career ladder; (3) that occupational transitions have on average been *towards* occupations that have higher analytical content and require more social skills, and out of occupations with more routine task requirements;[65] and (4) that workers have on average been moving into occupations that require more contact and working relationships with others (and so have less time flexibility).

While occupational transitions therefore do reflect similarity in tasks, temporal flexibility, and leadership requirements, we note that there is substantial variation in occupational transitions which is not captured by these other occupational similarity measures. Appendix Table A3 shows the adjusted R-squared statistics from regressions of $\pi_{o \to p}$ on our measures of skill distance, wage difference, amenity difference (temporal flexibility), leadership difference, and a composite skill measure. In all of these cases, while the correlation is strong and positive, the explanatory power is relatively low.[66]

is associated with lower occupational flows between these occupations: we demonstrate this relationship at the SOC 6-digit level with a larger variety of task and skill measures, and show that differences between occupations in temporal flexibility and leadership responsibilities also appear to determine workers' likelihood of moving between them.

[64]Note that this analysis involves directed relationships between occupations, so if the same share of moves in each direction is observed for an given occupation pair, the estimated effect of differences between them would be zero.

[65]These patterns could be in line both with career progression for individual workers, and/or with the aggregate decline of routine occupations over the same time period documented by Autor et al. (2006), and the increasing demand for social skills documented by Deming (2017).

[66]This contrasts with results in Macaluso (2019), who shows that at a 2-digit level, skill distance can explain $\sim$23% of the variation in flows between occupational groups. The difference in these results shows that while skill distance may be a good predictor of mobility for more aggregate occupational groupings, for the more detailed analysis in this paper it cannot capture much of the variation the sparse matrix of mobility between 6-digit occupational pairs. The failure of skill similarity measures to explain many occupational transitions can be illustrated by a few cases from

E Derivation of equilibrium wage expression and estimating equation

This section expands on our theory in section 2, showing how average wages in a local occupation can be expressed as a function of local concentration, productivity and outside options in a way that takes into account the bargaining links between occupations. The derivation involves iteratively substituting the Nash bargaining expression from equation (5). We will suppress indices for location and time here to make the exposition clearer, but all expressions are assumed to apply to an occupation in a particular location and time period.

However, before we can iteratively substitute, it will help to define some expressions that will allow us to simplify the exposition. Define different concentration indices of the form

$$H_x = \sum_i \sigma_i^x$$

where H_2 would correspond to a conventional Herfindahl-Hirschman Index. Moreover, define different aggregation levels recursively by

$$\begin{aligned}
\Omega_0 &= \sum_i \sigma_i = H_1 \\
\Omega_1 &= \sum_i \sigma_i (1-\sigma_i) = H_1 - H_2 \\
\Omega_2 &= \sum_i \sigma_i \sum_{j \neq i} \sigma_j (1-\sigma_j) = \Omega_1 - (H_2 - H_3) \\
\Omega_3 &= \sum_i \sigma_i \sum_{j \neq i} \sigma_j \sum_{l \neq j} \sigma_l (1-\sigma_l) = \Omega_2 - \Omega_1 H_2 + (H_3 - H_4) \\
&\vdots \\
\Omega_n &= \sum_i \sigma_i \sum_{j \neq i} \sigma_j \sum_{l \neq j} \sigma_l \cdots \sum_{z \neq y} \sigma_z (1-\sigma_z) = \sum_{j=1}^{n-1} H_j \Omega_{n-j} (-1)^{j-1} + (-1)^{n+1} (H_n - H_{n-1})
\end{aligned}$$

our data. First, consider some occupation pairs that are very similar on a skill distance metric (in the lowest distance decile), but where our data shows almost no (less than 0.01%) chance of moving from one to the other when switching jobs, in either direction: Surveyors and Medical & clinical laboratory technologists; Carpenters and Dental assistants; Travel agents and Police, fire & ambulance dispatchers. In all of these occupational pairs it is intuitively clear why they may look similar in terms of an abstract description of the tasks involved, but in practice this skill distance does not make them relevant outside options for one another because of differences in other job characteristics or requirements. Second, consider another pair of occupations which are very similar on the skill distance metric (again, in the lowest distance decile): Pediatricians and Management analysts. When pediatricians change jobs, 8.7% of them become management analysts, but less than 0.01% of management analysts switching jobs become pediatricians. The skill distance metric misses the fact that one of these occupations requires extensive training and licensing which means that, in practice, the occupational move is only possible in one direction.

Now, by definition, average wages in local occupation o are given by:

$$\bar{w}_o = \sum_i \sigma_i w_{i,o}.$$

Iteratively substituting for the wage using the expression in equation (5), we obtain

$$\begin{aligned}
\bar{w}_o &= \sum_i \sigma_i \left(\beta MPL_o + (1-\beta) \left(Prob_{o\to o} \cdot \sum_{j\neq i}^{N_{i,o}} \sigma_{j,o} \cdot w_{j,o} \right) + (1-\beta) oo_{i,o}^{occs} \right) \\
&= \left(\beta MPL_o + (1-\beta) oo_o^{occs}\right) \left(1 + (1-\beta) \cdot Prob_{o\to o} \cdot \sum_i \sigma_i \sum_{j\neq i}^{N_{i,o}} \sigma_{j,o} \right) \\
&+ (1-\beta)^2 \cdot Prob_{o\to o}^2 \cdot \sum_i \sigma_i \sum_{j\neq i}^{N_{i,o}} \sigma_{j,o} \cdot \sum_{l\neq j}^{N_{j,o}} \sigma_{l,o} w_{l,o} \\
&= \left(\beta MPL_o + (1-\beta) oo_o^{occs}\right) \left(1 + (1-\beta) \cdot Prob_{o\to o} \cdot \sum_i \sigma_i \sum_{j\neq i}^{N_{i,o}} \sigma_{j,o} + (1-\beta)^2 \cdot Prob_{o\to o}^2 \cdot \sum_i \sigma_i \sum_{j\neq i}^{N_{i,o}} \sigma_{j,o} \right. \\
&+ (1-\beta)^3 \cdot Prob_{o\to o}^3 \cdot \sum_i \sigma_i \sum_{j\neq i}^{N_{i,o}} \sigma_{j,o} \cdot \sum_{l\neq j}^{N_{j,o}} \sigma_{l,o} \cdot \sum_{m\neq l}^{N_{j,o}} \sigma_{m,o} w_{m,o} \\
&\vdots \\
&= \left(\beta MPL_o + (1-\beta) oo_o^{occs}\right) \left(1 + (1-\beta) Prob_{o\to o} \Omega_1 + (1-\beta)^2 Prob_{o\to o}^2 \Omega_2 + (1-\beta)^3 Prob_{o\to o}^3 \Omega_3 \ldots \right) \\
&= \left(\beta MPL_o + (1-\beta) oo_o^{occs}\right) (\psi_o),
\end{aligned}$$

where the final term ψ_o is a function of the order r concentration indices we defined above, Ω_r:

$$\psi_o = \sum_{r=0}^{\infty} (1-\beta)^r Prob_{o\to o}^r \Omega_r.$$

Then, log average wages are given by

$$\ln \bar{w}_o = \ln \left(\beta MPL_o + (1-\beta) oo_o^{occs}\right) + \ln \psi_o,$$

which is the expression shown in equation (6) in the main theory section.

For the first-order approximation to the ψ_o term, note that if $r = 1$,

$$\begin{aligned}\psi_o &= \Omega_0 + (1-\beta)Prob_{o\to o}\Omega_1 \\ &= H_1 + (1-\beta)Prob_{o\to o}(H_1 - H_2) \\ &= 1 + (1-\beta)Prob_{o\to o}(1 - HHI),\end{aligned}$$

which gives us the second term in equation (7) once we add the indices for occupation and location to the concentration index.

In order to derive our estimating equation (13) we totally differentiate the log average wage expression from equation (7) for an assumed reference city with median characteristics $\bar{oo}_o^{occs}$, $M\bar{P}L_o$, $\bar{w}_o$ and $H\bar{H}I_o$. Thus, we obtain

$$\begin{aligned}\Delta \ln \bar{w}_{o,k} = &\left(\frac{(1-\beta)\bar{oo}_o^{occs}}{\beta M\bar{P}L_o + (1-\beta)\bar{oo}_o^{occs}}\right)\Delta \ln oo_{o,k}^{occs} + \left(\frac{\beta M\bar{P}L_o}{\beta M\bar{P}L_o + (1-\beta)\bar{oo}_o^{occs}}\right)\Delta \ln MPL_{o,k} \\ &- \frac{(1-\beta)Prob_{o\to o}H\bar{H}I_o}{1+(1-\beta)Prob_{o\to o}(1-H\bar{H}I_o)}\Delta \ln HHI_{o,k}\end{aligned}$$

where the Δ operator denotes differentials relative to the reference city. Writing the differential explicitly as the difference between the city k, occupation o cell and the reference city-occupation and gathering all the reference cell terms, we obtain

$$\ln \bar{w}_{o,k} = \alpha + \gamma_1 \ln oo_{o,k}^{occs} + \gamma_2 \ln HHI_{o,k} + \epsilon_{o,k}.$$

where

$$\begin{aligned}\gamma_1 &= \left(\frac{(1-\beta)\bar{oo}_o^{occs}}{\beta M\bar{P}L_o + (1-\beta)\bar{oo}_o^{occs}}\right) \\ \gamma_2 &= -\frac{(1-\beta)Prob_{o\to o}H\bar{H}I_o}{1+(1-\beta)Prob_{o\to o}(1-H\bar{H}I_o)} \\ \gamma_1 &= \left(\frac{\beta M\bar{P}L_o}{\beta M\bar{P}L_o + (1-\beta)\bar{oo}_o^{occs}}\right) \\ \alpha &= \bar{w}_o - \gamma_1 \ln \bar{oo}_o^{occs} - \gamma_2 \ln H\bar{H}I_o - \gamma_3 \ln M\bar{P}L_o \\ \epsilon_{o,k} &= \gamma_3 \ln MPL_{o,k} \\ (\text{and } \Delta \ln \bar{w}_{o,k} &\equiv \ln \bar{w}_{o,k} - \bar{w}_o).\end{aligned}$$

Note that the productivity residual $\epsilon_{o,k}$ represents the main identification concern as it is likely correlated with other local occupation characteristics. We can project this productivity term on fixed effects in order to remove some of the confounding variation, and -for our preferred specification -

obtain the occupation-year and city-year demeaned productivity residual $\xi_{o,k}$.

Adding time subscripts (with the reference city-occupation now representing an average across time) and the city-year and occupation-year fixed effects into the wage equation, we obtain the baseline estimating equation (equation 13):

$$\ln \bar{w}_{o,k,t} = \alpha + \alpha_{o,t} + \alpha_{k,t} + \gamma_1 \ln oo^{occs}_{o,k,t} + \gamma_2 \ln HHI_{o,k,t} + \xi_{o,k,t}.$$

F Theory extensions

Simple matching model

In section 2, we outline a simple search model which generates a transition-weighted average wage as a measure of the value of outside options. The transition-weighted average wage can also be justified as a measure of outside options using a simple matching model, which we outline below. In our matching model, there are no search frictions. Instead, heterogeneous workers all work at the firm at which they are most productive. Workers have job offers from other firms, so they know the value of their next-best outside option. This next-best option could be a job in the worker's own occupation, or in a different occupation. As in the search model, the worker and firm Nash bargain, so that the worker's wage is a weighted average of her marginal product in the job and her outside option:

$$w_i = \beta(MPL_i - oo_i) + oo_i = \beta MPL_i + (1-\beta)oo_i \tag{18}$$

Each worker's wage is different, since each worker has a different best outside job option (and the worker and firm both know its value). To construct the average wage in a given occupation and city, we can segment the workers within that occupation and city into five groups: those whose best outside job option is in their own occupation, those whose best outside job option is in another occupation, and those whose best outside option is unemployment. Within each of these labor markets, assume that workers are offered a wage equal to the average wage in that labor market. This gives us an expression for the average value of the outside option in occupation o and city k:

$$\bar{oo}_{o,k} = \varsigma_{o,k}\bar{w}_{o,k} + \sum_{p \neq o}^{N_{occs}} \varsigma_{p,k}\bar{w}_{p,k} + \left(1 - \sum_{p}^{N_{occs}} \varsigma_{p,k}\right) b$$

where $\varsigma_{p,l}$ is the share of workers in occupation o & city k with best outside option in occupation p. Assume that workers' actual occupational moves reflect moves to their best outside job option - either because they were involuntarily displaced and had to find their next-best job, or because they left their job by choice after a preference shock. Then, assume that the distribution of best outside options for workers who *remain* in their jobs in occupation o and city k is equal to the distribution of best outside offers for workers who used to be in jobs in occupation o and city k. Under these assumptions, we can use occupational transitions to approximate for $\varsigma_{p,k}$ in the outside options expression, meaning that the average value of the best outside-labor-market option for workers in occupation o and city k is the weighted average of wages in other occupations, weighted by the proportion of workers who moved from occupation o and city k to each of the other relevant labor markets.

Outside options in different cities

In section 2 we restrict our analysis of outside options to job options within workers' current city. In practice, for some workers there will be relevant job options outside their city or local area. The intuition of our theory from section 2 can easily be applied to incorporate job options outside workers' current city, as follows. Define each labor market as an occupation p-by-city l cell. A worker's outside options can then be segmented into the value of their outside job options in each occupation p in each city l as follows:

$$oo_{i,o,k} = \underbrace{Prob_{o,k\to o,k}\cdot\sum_{j\neq i}^{N_{o,k}}\sigma_{j,o,k}\cdot w_{j,o,k}}_{\text{options in own occ and city}} + \underbrace{\sum_{p\neq o}^{N_{\text{occs,k}}} Prob_{o,k\to p,k}\cdot\sum_{m=1}^{N_p}\sigma_{m,p,k}\cdot w_{m,p,k}}_{\text{options in other occs, own city}} + \underbrace{\sum_{l\neq k}^{N_{\text{o,cities}}} Prob_{o,k\to o,l}\cdot\sum_{n=1}^{N_l}\sigma_{n,o,l}\cdot w_{n}}_{\text{options in own occ, other cities}} \qquad (19)$$

where $Prob_{o,k\to p,l}$ is the probability that, for a worker currently in occupation o and city k, her best outside job option is in occupation p and city l. As we do with occupational mobility, in theory mobility between occupation-city pairs could be used to approximate $Prob_{o,k\to p,l}$ if sufficient data is available.

G Identification assumptions in IV analysis

This section provides more formal details on the assumptions required for identification of the outside-occupation options effect on wages using the instrumental variables strategy based on national leave-one out mean wages (in Section 5).

To avoid endogeneity concerns over the local employment shares, we instrument for the local relative employment share in each occupation using the initial employment share in that occupation in 1999, the first year in the data.[67] Our instrument for the oo^{occs} index, $oo^{occs,inst}$, therefore becomes the weighted average of national leave-one out mean wages in occupation p, $\bar{w}_{p,\cancel{k},t}$, where

[67] Or we use the first year the occupation-city cell is in the data, if it is not present in 1999

the weights are the product of the year 1999 relative employment share in each of those occupations in the worker's own city, $\frac{s_{p,k,1999}}{s_{p,1999}}$, and the national occupation transition shares from the worker's occupation o to each of the other occupations, $\pi_{o\to p}$.

$$oo^{occs,inst}_{o,k,t} = \sum_{p}^{N_{occs}} \left(\pi_{o\to p} \cdot \frac{s_{p,k,1999}}{s_{p,1999}} \cdot \bar{w}_{p,\cancel{k},t}\right) \tag{20}$$

To make the assumptions transparent under which this wage instrument identifies β^{oo} in equation (13), we follow the framework presented in Borusyak et al. (2018). For simplicity, assume that outside options and the concentration index are not correlated - but the intuition for the identification does not depend on that. Note that we can write the instrument as

$$oo^{occs,inst}_{o,k,t} = \sum_{p=1}^{N^{occs}} s_{okp}\bar{w}_{p,\cancel{k},t}$$

where $s_{okp} = \pi_{o\to p} \cdot \frac{s_{p,k,1999}}{s_{p,1999}}$ is a measure of predicted local exposure to the shock. In our fixed effects IV estimation of equation (13), the exclusion restriction for the instrument for outside-occupation options is then equivalent to

$$Cov[oo^{occs,inst}_{o,k,t}, \xi_{o,k,t}|\Gamma_{kt}, \Gamma_{ot}] = \sum_{t=1}^{T}\sum_{p=1}^{N^{occs}} \bar{s}_{okp} w^{\perp}_{p,\cancel{k},t}\phi^{oo}_{pt} \to 0$$

where $\bar{s}_{okp} = \mathbb{E}[s_{okp}]$ is the average exposure to occupation p, and $\phi^{oo}_{pt} \equiv \mathbb{E}[s_{okp}\xi_{o,k,t}]/\mathbb{E}[s_{okp}]$ is an exposure-weighted expectation of the structural wage residuals. Moreover, $w^{\perp}_{p,\cancel{k},t}$ represents $\bar{w}_{p,\cancel{k},t}$ after it has been residualized with regard to city-k-by-year-t fixed effects Γ_{kt} and occupation-o-by-year-t fixed effects Γ_{ot}, as well as the concentration index..

Borusyak et al. (2018) show that this orthogonality condition holds under two assumptions. First, we require that the national occupation-level shocks are quasi-randomly assigned conditional on local exposure to structural wage shocks ϕ_{pt} and the fixed effects Γ_{kt} and Γ_{ot}. That is,

$$\mathbb{E}[\bar{w}_{p,\cancel{k},t}|\phi^{oo}_{pt}, \Gamma_{kt}, \Gamma_{ot}] = \tau_1\Gamma_{kt} + \tau_2\Gamma_{ot} \quad \forall p \in N^{occs}$$

for some constant parameters τ_1 and τ_2. Second, there needs to be a large number of independent occupational shocks, that is,

$$\mathbb{E}[(\bar{w}_{p,\cancel{k},t} - \mu)(\bar{w}_{j,\cancel{k},t} - \mu)|\phi_{pt}, \phi^{oo}_{jt}, \Gamma_{kt}, \Gamma_{ot}] = 0$$

for all $p, j \in N^{occs}$ if $p \neq j$, and also $\sum_{p=1}^{N^{occs}} \bar{s}_p^2 \to 0$.

The first assumption requires that the national leave-one-out mean wage $\bar{w}_{p,\not{k},t}$ in outside option occupation p is correlated with the local wage of occupation o in location k (relevance condition), but does not affect the local wage in initial occupation o through a direct channel other than increasing the quality of local outside options $oo_{o,k,t}^{occs,inst}$. However, this lack of a direct effect only needs to hold *conditional* on controlling for fixed effects that include the national wage trend in occupation o itself and wage trends that are common to all occupations in city k. The inclusion of these fixed effects increases our confidence that the assumptions for instrument validity hold.

We can proceed similarly for the granular IV identification of the concentration effects. For details on the granular IV identification approach more generally, see Gabaix and Koijen (2020). Here, we will focus on our specific application to local concentration indices, and again assume outside options and concentration are not correlated to simplify the exposition.

We can rewrite the concentration instrument as

$$\begin{aligned}\Delta HHI_{o,k,t}^{inst} &= \sum_j \sigma_{j,o,k,t-1}^2 \left(\frac{(1+\tilde{g}_{j,t})^2}{(1+g_{o,k,t})^2} - 1 \right) \\ &= \sum_j \sigma_{j,o,k,t-1}^2 \tilde{G}_{j,o,k,t}\end{aligned}$$

where $\tilde{G}_{j,o,k,t} = \frac{(1+\tilde{g}_{j,t})^2}{(1+g_{o,k,t})^2} - 1$ is the predicted firm-level excess local growth relative to the average local occupation growth - the time-varying shock - and $\sigma_{j,o,k,t-1}^2$ is the exposure of the local concentration index to that shock.

In our fixed effects IV estimation of equation (13), the exclusion restriction for the instrument on the HHI concentration index is then equivalent to

$$Cov[HHI_{o,k,t}^{inst}, \xi_{o,k,t} | \Gamma_{kt}, \Gamma_{ot}, \tilde{g}_{o,k,t}] = \mathbb{E}\left[\sum_{t=1}^{T} \sum_{p=1}^{N^{occs}} \sigma_{j,o,k,t-1}^2 \tilde{G}_{j,o,k,t}^{\perp} \xi_{o,k,t} \right] \to 0$$

Here, $\tilde{G}_{j,o,k,t}^{\perp}$ represents $\tilde{G}_{j,o,k,t}$ after it has been residualized with regard to city-k-by-year-t fixed effects Γ_{kt} and occupation-o-by-year-t fixed effects Γ_{ot}, as well as $\tilde{g}_{o,k,t}$, the local linear share-weighted exposure to the firm growth instrument.

This orthogonality condition holds under two assumptions. First, we require that the national firm-level growth shocks are quasi-randomly assigned conditional on local exposure to structural wage shocks $\xi_{o,k,t}$, the fixed effects Γ_{kt} and Γ_{ot}, and predicted average local vacancy growth $\tilde{g}_{o,k,t}$. That is,

$$\mathbb{E}[\sum_j \sigma_{j,o,k,t-1}^2 \tilde{G}_{j,o,k,t} | \xi_{o,k,t}, \Gamma_{kt}, \Gamma_{ot}, \tilde{g}_{o,k,t}] = \tau_1^{HHI} \Gamma_{kt} + \tau_2^{HHI} \Gamma_{ot} + \tau_3^{HHI} \tilde{g}_{o,k,t} \quad \forall p \in N^{occs}$$

for some constant parameters τ_1^{HHI} and τ_2^{HHI} and τ_3^{HHI}. That is, once we account for the control variables, expected local squared exposure to excess national firm-level growth needs to be random in expectation.

Second, there needs to be a large number of independent firm-level shocks, that is,

$$\mathbb{E}[(\tilde{G}_{j,o,k,t} - E[\tilde{G}_{j,o,k,t}])(\tilde{G}_{j,o,k,t} - E[\tilde{G}_{j,o,k,t}])|\phi_{pt}, \phi_{jt}, \Gamma_{kt}, \Gamma_{ot}] = 0$$

for all $p, j \in N^{occs}$ if $p \neq j$.

The first assumption requires that the local size-weighted exposure to national firm-level employment shocks does not affect the local wage in occupation o through a direct channel other than increasing the local labor market concentration $HHI_{o,k,t}$, conditional on the control variables. Note that this allows for different local occupations to have different average expected average growth rates based on national firm growth. It only requires that whether this growth is driven by the national growth of locally large firms vs. small firms varies across local occupations in a way that is uncorrelated with local wage residuals.

To be concrete, note again the example from the main text: Imagine there are two large coffee chains in in Portland, OR, and in Hartford, CT, respectively: Starbucks, and Dunkin'. In Portland, Starbucks is relatively more dominant than Dunkin', whereas in Hartford it is the other way around, meaning that the employment share for Starbucks is much larger in Hartford than in Portland.

We noted before, that, in years where Starbucks grows substantially faster than Dunkin' nationwide, employer concentration of baristas will grow by more in Portland than in Hartford. Moreover, our granular IV identification approach controls for local growth rates of overall barista employment in both cities. Thus, it allows for each city to be exposed differently to overall trends in the demand for coffee and baristas. The identification only requires that once we account for overall city exposure to barista demand, whether that demand was driven by the city's major chain or smaller stores is not correlated with local idiosyncratic barista wage shocks.

H Appendix: Figures

BGT Data: education relative to 2018 labor force

Figure A1: Comparison of distribution of highest educational attainment in the labor force, according to BLS data, to distribution in BGT data. Two versions are shown: BGT 1 excludes all resumes missing educational information, while BGT 2 assumes all resumes missing educational information have high school education but no college

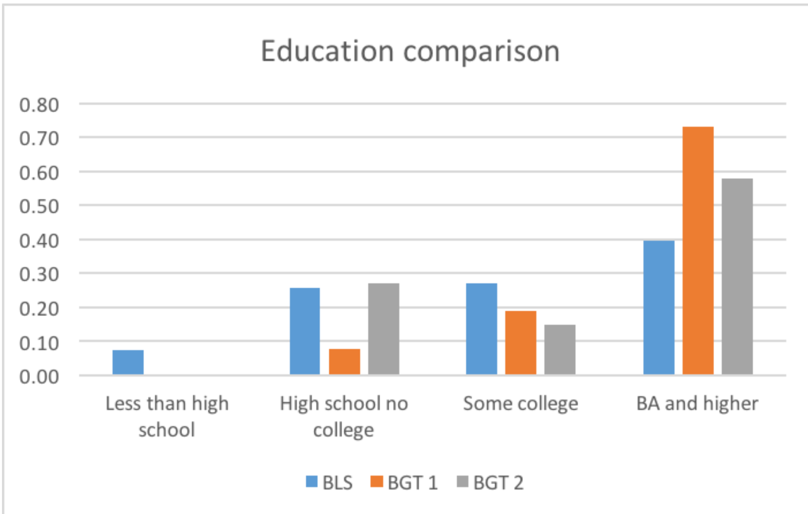

BGT Data: age distribution relative to 2018 labor force

Figure A2: Comparison of distribution of age in the labor force, according to 2018 BLS data, to distribution of imputed worker ages in BGT job sequence data.

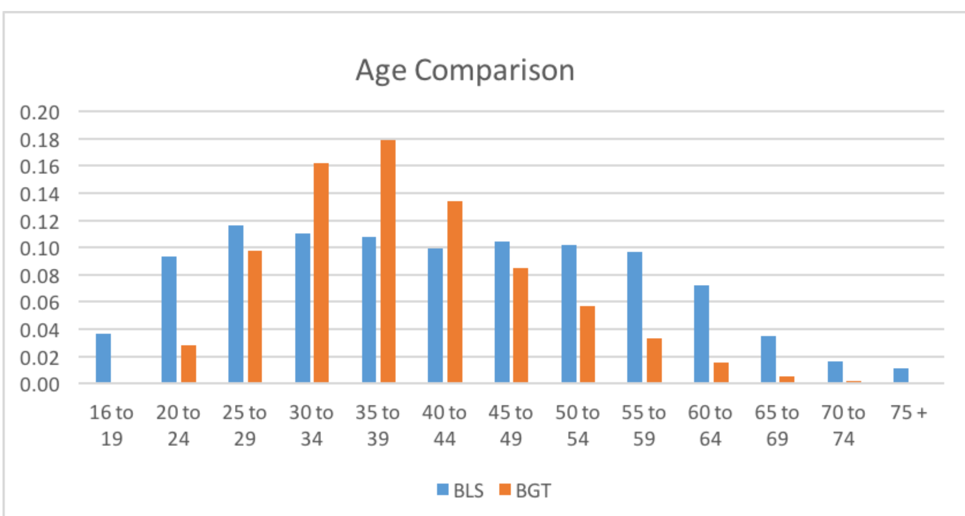

BGT Data: occupations relative to 2017 labor force

Figure A3: Comparison of distribution of 2-digit SOC occupations in the labor force, according to 2017 BLS data, to distribution of occupations in BGT job sequence data.

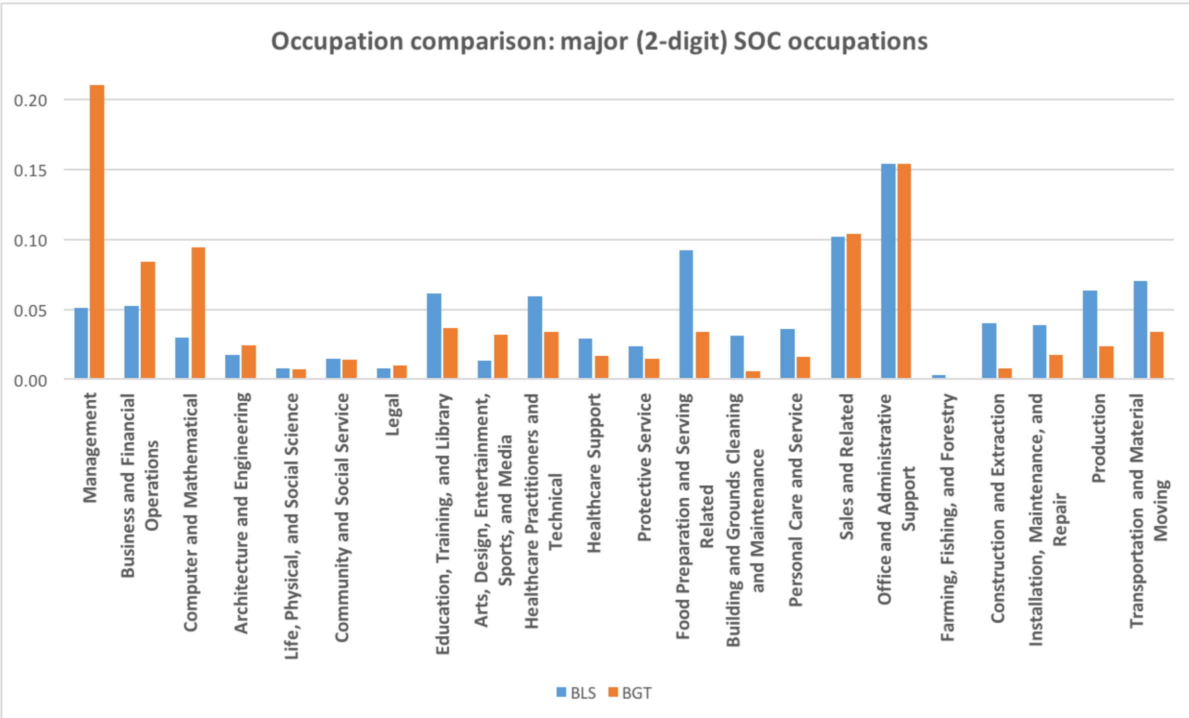

BGT Data: locations relative to 2017 labor force

Figure A4: Comparison of distribution of employment by U.S. state, according to 2017 BLS data, to distribution of resume addresses in BGT job sequence data. Graph shows share of total in each state.

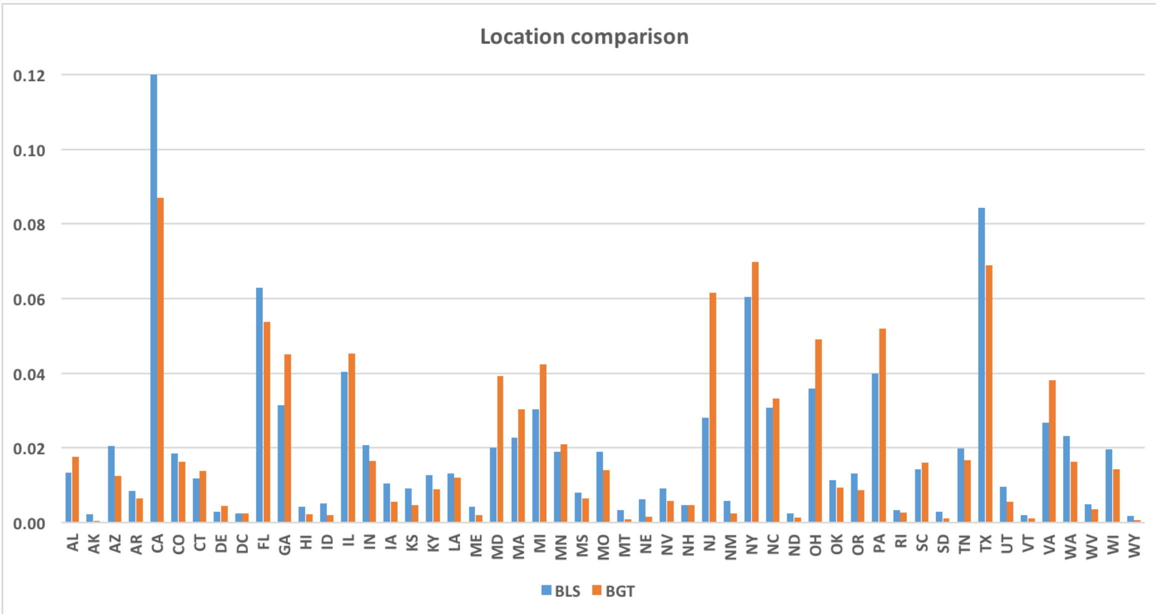

Outward occupational mobility from SOC 6-digit occupations

Figure A5: Distribution of the "occupation leave share" – the probability that a worker will leave their occupation conditional on leaving their job – by occupation. Occupation leave share is calculated from BGT resume data for 2002-2015 period. Histogram shows 786 occupations, with dashed line indicating the sample mean.

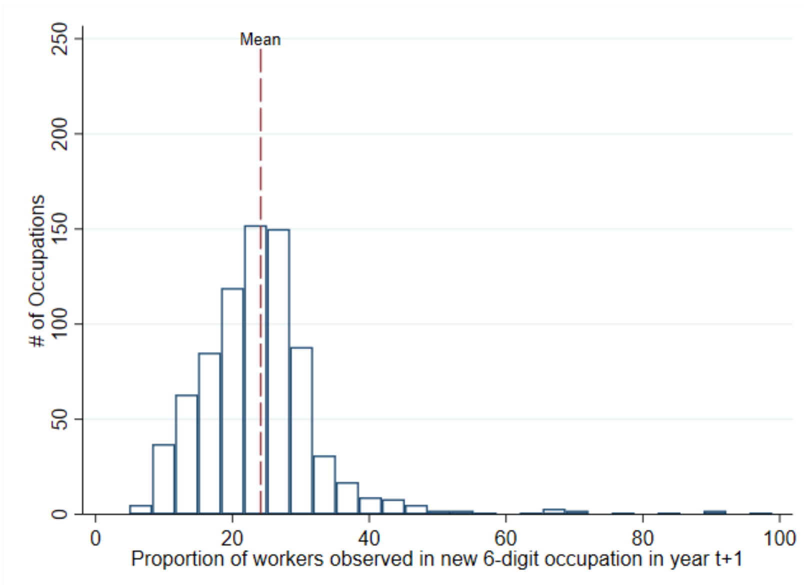

6-digit moves that are also 2-digit moves

Figure A6: Distribution of the proportion of workers moving 6-digit SOC occupation who *also* move 2-digit SOC occupation, by occupation, calculated from BGT resume data for 2002-2015 period. Histogram shows 786 occupations.

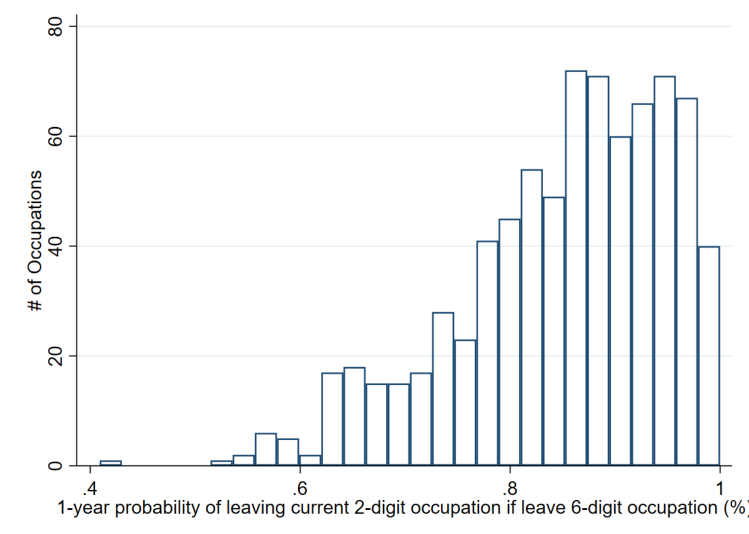

6-digit SOC occupational transition matrix

Figure A7: Occupational transition matrix showing transition probability between 6-digit SOC occupations conditional on leaving the initial job. Occupations are sorted in SOC numerical order. Cells colored black have a transition probability of 1% or greater conditional on leaving the initial job. Transitions to own occupation are excluded. Data computed from BGT resume data set for 2002-2015. The annotation points out certain common destination occupations, which show up as darker vertical lines on the heatmap.

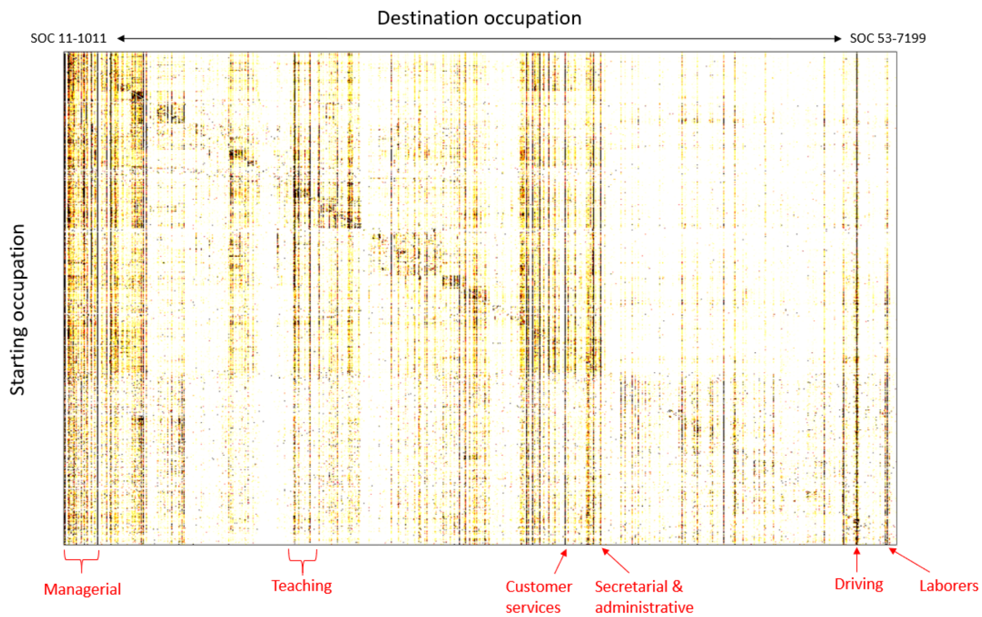

2-digit SOC occupational transition matrix

Figure A8: Occupational transition matrix showing transition probability between 2-digit SOC occupation groups conditional on leaving the initial job. Cells colored black have a transition probability of 25% or greater conditional on leaving the initial job. Job transitions within an occupation group are excluded. Data computed from BGT resume data set for 2002-2015.

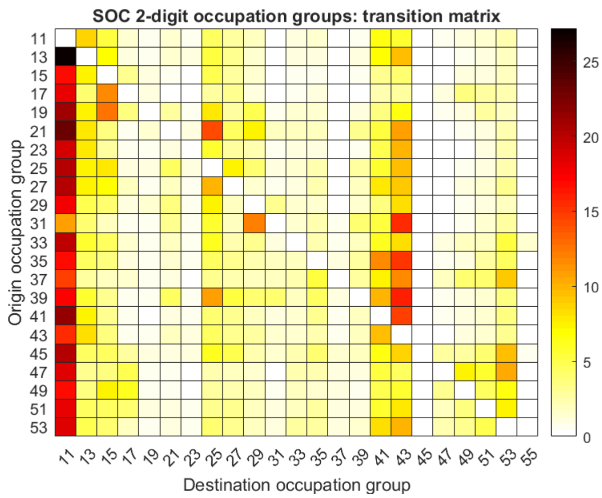

Examples of probabilistic labor markets

The graphs below illustrates the most common occupation transition paths for counter attendants and registered nurses, respectively. For both of these occupations, the majority of people who leave their SOC 6-digit occupation also leave their SOC 2-digit occupation group, but the pattern is very different. Counter attendants' outside-occupation job options are very diverse, and are mostly lateral moves into jobs in sales, office & administrative work, and food preparation and service. In contrast, almost all registered nurses who leave their occupation do so through a promotion, becoming medical and health service managers.

Figure A9: Occupational transitions for counter attendants in the food industry. Each bubble is a SOC 6-digit occupation, and the colors represent SOC 2-digit occupational groups. The size of each bubble is proportional to the share of counter attendants in the BGT data who are observed in each destination occupation in the following year.

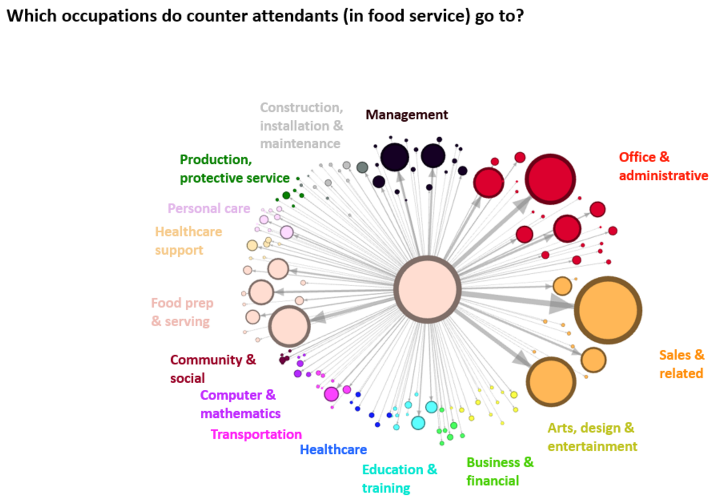

Figure A10: Occupational transitions for registered nurses. Each bubble is a SOC 6-digit occupation, and the colors represent SOC 2-digit occupational groups. The size of each bubble is proportional to the share of registered nurses in the BGT data who are observed in each destination occupation in the following year.

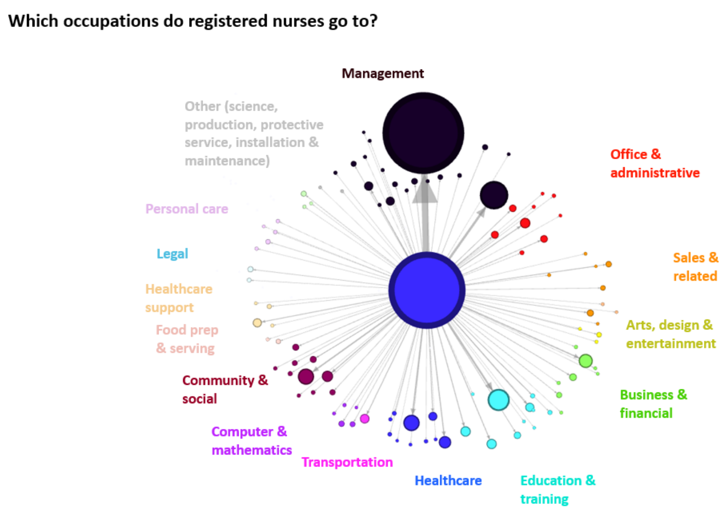

Coefficients on wage-HHI regressions, with and without controls for outside-occupation options

Figure A11: Coefficients on HHI and 95% confidence intervals from regressions of occupation-CBSA wages on employer HHI, without controlling for outside-occupation job options (black circles), and with controls for outside-occupation job options (blue diamonds), for full sample and segmenting by quartile of the occupation leave share (a proxy for outward occupational mobility). Regressions span 2013-2016 and include occupation-year and CBSA fixed effects. Standard errors are clustered at the CBSA level.

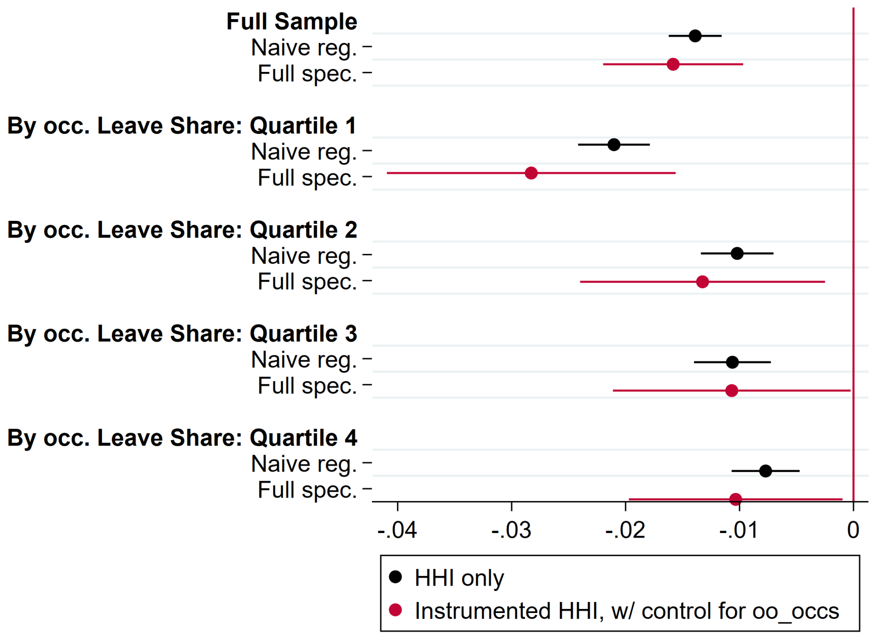

Correlations between wages, within-occupation HHI, and outside-occupation options

Figure A12: Binned scatter plots of the correlation between average wages, within-occupation HHI index, and outside-occupation option index for occupation-CBSA cells in 2016.

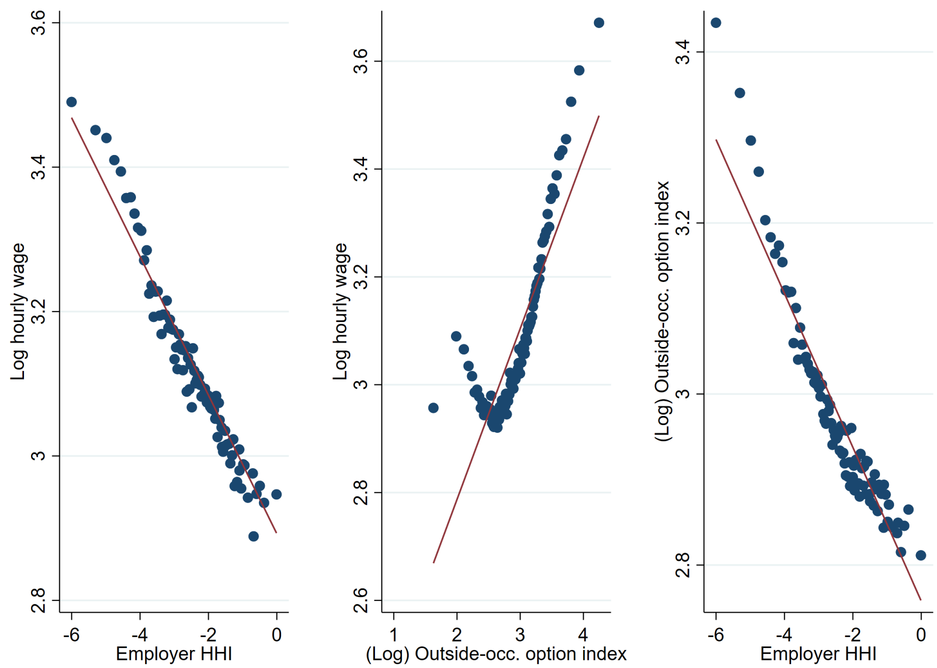

Residualized relationship between wages, within-occupation HHI, and outside-occupation options

Figure A13: Binned scatter plots of the log wage on log of outside-occupation options, controlling for log HHI and partialling out occupation-year and city-year fixed effects

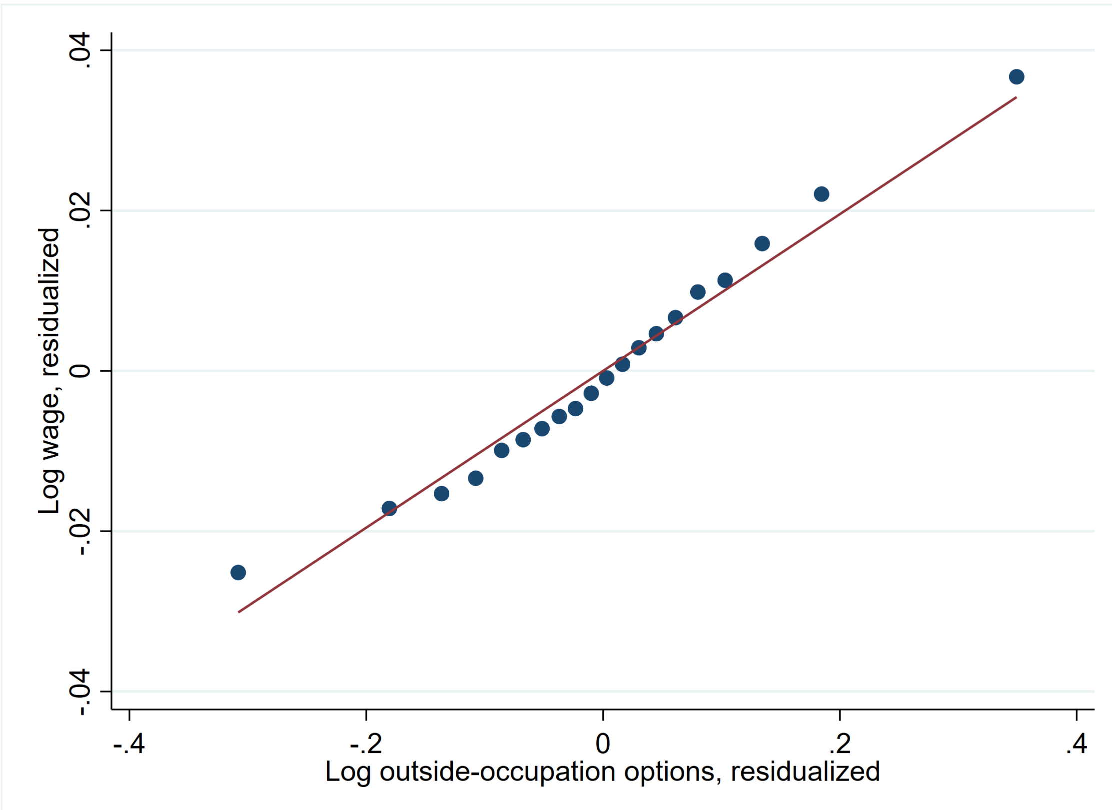

Figure A14: Binned scatter plots of the log wage on log HHI, controlling for the log of outside-occupation options and partialling out occupation-year and city-year fixed effects

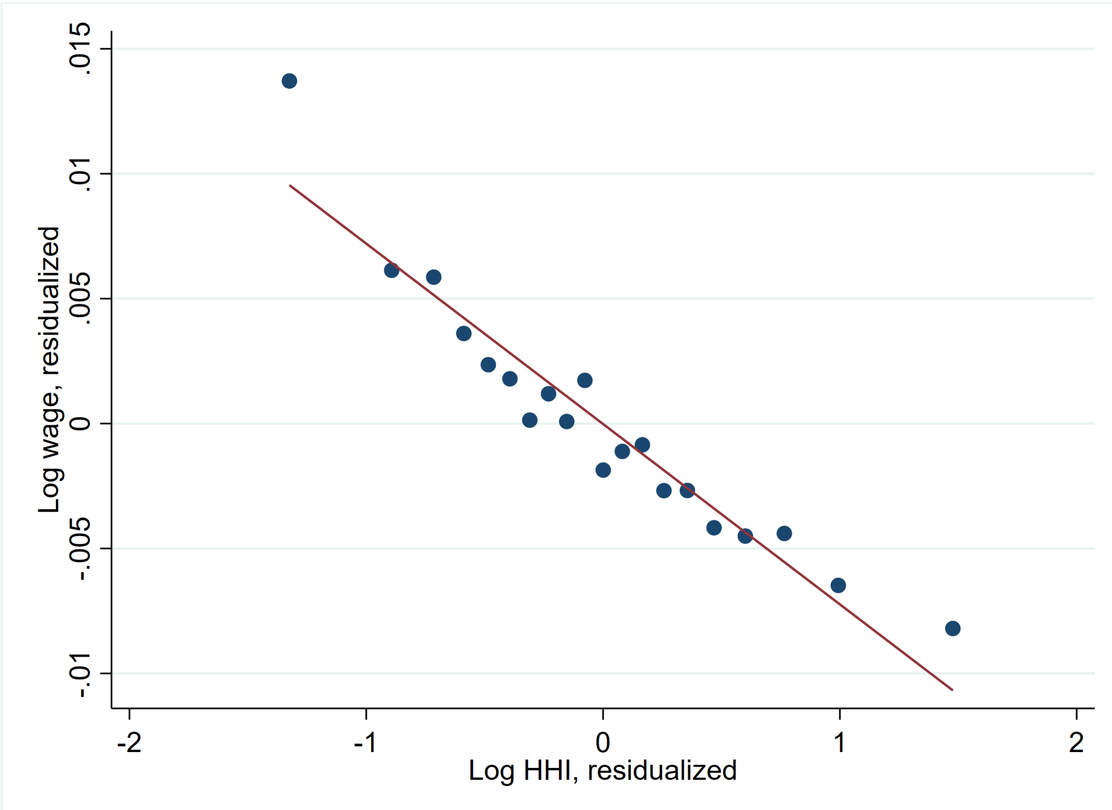

I Appendix: Tables

Table A1: Twenty large occupations with lowest leave shares and highest leave shares

Initial occupation	Leave share	Employment (2017)	Obs. (BGT)	Modal new occupation
Dental hygienists	.062	211,600	17,458	Dental assistants
Nurse practitioners	.088	166,280	57,830	Registered nurses
Pharmacists	.09	309,330	121,887	Medical and health services managers
Firefighters	.098	319,860	60,039	Emergency medical technicians and paramedics
Self-enrichment education teachers	.1	238,710	169,369	Teachers and instructors, all other
Physical therapists	.11	225,420	44,314	Medical and health services managers
Postsecondary teachers, all other	.11	189,270	825,879	Managers, all other
Graphic designers	.12	217,170	439,953	Art directors
Emergency medical technicians and paramedics	.12	251,860	111,180	Managers, all other
Fitness trainers and aerobics instructors	.13	280,080	281,903	Managers, all other
Licensed practical and licensed vocational nurses	.13	702,700	254,787	Registered nurses
Lawyers	.13	628,370	667,960	General and operations managers
Registered nurses	.13	2,906,840	1,427,102	Medical and health services managers
Health specialties teachers, postsecondary	.13	194,610	41,963	Medical and health services managers
Physicians and surgeons, all other	.14	355,460	59,630	Medical and health services managers
Heavy and tractor-trailer truck drivers	.14	1,748,140	2,174,486	Managers, all other
Radiologic technologists	.14	201,200	80,347	Magnetic resonance imaging technologists
Hairdressers, hairstylists, and cosmetologists	.14	351,910	107,167	Managers, all other
Coaches and scouts	.14	235,400	533,082	Managers, all other
Chief executives	.15	210,160	1,425,400	General and operations managers
...				
Installation, maintenance, and repair workers, all other	.29	153,850	60,742	Maintenance and repair workers, general
Parts salespersons	.29	252,770	34,038	First-line supervisors of retail sales workers
Billing and posting clerks	.29	476,010	274,963	Bookkeeping, accounting, and auditing clerks
Data entry keyers	.29	180,100	288,523	Customer service representatives
Cashiers	.29	3,564,920	1,753,947	Customer service representatives
Insurance claims and policy processing clerks	.3	277,130	235,763	Claims adjusters, examiners, and investigators
Stock clerks and order fillers	.3	2,046,040	597,137	Laborers and freight, stock, and material movers, hand
Packers and packagers, hand	.3	700,560	101,025	Laborers and freight, stock, and material movers, hand
Cooks, institution and cafeteria	.3	404,120	5,174	Cooks, restaurant
Helpers–production workers	.31	402,140	112,759	Production workers, all other
Sales rep., wholesale & mfg., tech. & scient. products	.31	327,190	198,337	Sales rep., wholesale & mfg., exc. techn. & scient. products
Hosts and hostesses, restaurant, lounge, and coffee shop	.31	414,540	159,098	Waiters and waitresses
Shipping, receiving, and traffic clerks	.31	671,780	318,080	Laborers and freight, stock, and material movers, hand
Loan interviewers and clerks	.32	227,430	234,933	Loan officers
Counter attendants, cafeteria, food concession, and coffee shop	.32	476,940	118,131	Retail salespersons
Bill and account collectors	.32	271,700	310,951	Customer service representatives
Tellers	.32	491,150	468,829	Customer service representatives
Machine setters, operators, and tenders†	.32	154,860	6,805	Production workers, all other
Telemarketers	.36	189,670	47,409	Customer service representatives
Food servers, nonrestaurant	.45	264,630	13,199	Waiters and waitresses

This table shows the twenty large occupations with the lowest and the highest occupation leave shares - defined as the 1-year horizon probability of no longer working in their current occupation, conditional on leaving their job - in the BGT data over 2002-2015, as well as total national employment in that occupation in 2017 from the OES, the number of occupation-year observations in the BGT data ('obs.') and the most popular occupation that workers who leave the initial occupation move to ('modal new occupation'). Large occupations are defined as those with national employment over 150,000 in 2017 (roughly the 75th percentile of occupations when ranked by nationwide employment). † Full occupation title is "Molding, coremaking, and casting machine setters, operators, and tenders, metal and plastic."

Table A2: Forty thickest occupational transition paths for large occupations

Initial occupation	New occupation	Transition share	Employment (2017)	Obs. (BGT data)
Licensed practical and licensed vocational nurses	Registered nurses	.3	702,700	254,787
Nurse practitioners	Registered nurses	.23	166,280	57,830
Construction managers	Managers, all other	.19	263,480	917,349
Sales rep., wholesale & mfg., tech. & scient. products	Sales rep., wholesale & mfg., exc. tech. & scient. products	.19	327,190	198,337
Physicians and surgeons, all other	Medical and health services managers	.19	355,460	59,630
Software developers, systems software	Software developers, applications	.19	394,590	53,322
Legal secretaries	Paralegals and legal assistants	.18	185,870	132,543
Accountants and auditors	Financial managers	.18	1,241,000	1,459,175
Registered nurses	Medical and health services managers	.16	2,906,840	1,427,102
Cost estimators	Managers, all other	.16	210,900	124,646
Human resources specialists	Human resources managers	.16	553,950	2,035,604
Physical therapists	Medical and health services managers	.16	225,420	44,314
Architectural and engineering managers	Managers, all other	.15	179,990	749,670
Computer programmers	Software developers, applications	.15	247,690	533,764
Software developers, applications	Computer occupations, all other	.15	849,230	2,110,229
Computer network architects	Computer occupations, all other	.15	157,830	407,591
Cooks, short order	Cooks, restaurant	.15	174,230	39,906
Cooks, institution and cafeteria	Cooks, restaurant	.14	404,120	5,174
First-line supervisors of construction trades and extraction workers	Construction managers	.14	556,300	186,747
Computer systems analysts	Computer occupations, all other	.14	581,960	1,152,614
Sales rep., wholesale & mfg., exc. tech. & scient. products	Sales managers	.13	1,391,400	4,377,654
Light truck or delivery services drivers	Heavy and tractor-trailer truck drivers	.13	877,670	226,349
Computer occupations, all other	Managers, all other	.13	315,830	3,515,188
Health specialties teachers, postsecondary	Medical and health services managers	.13	194,610	41,963
Meat, poultry, and fish cutters and trimmers	Heavy and tractor-trailer truck drivers	.13	153,280	2,383
Sales rep., wholesale & mfg., tech. & scient. products	Sales managers	.13	327,190	198,337
Operating engineers and other construction equipment operators	Heavy and tractor-trailer truck drivers	.13	365,300	55,317
Sales managers	Sales rep., wholesale & mfg., exc. tech. & scient. products	.13	371,410	3,471,904
Health specialties teachers, postsecondary	Registered nurses	.13	194,610	41,963
Industrial engineers	Engineers, all other	.13	265,520	171,358
Network and computer systems administrators	Computer occupations, all other	.13	375,040	1,103,700
Industrial production managers	Managers, all other	.12	171,520	750,609
Computer network support specialists	Computer user support specialists	.12	186,230	237,766
Software developers, systems software	Computer occupations, all other	.12	394,590	53,322
Financial analysts	Financial managers	.12	294,110	664,903
Legal secretaries	Secretaries and admin. assistants, except legal, medical, & exec.	.12	185,870	132,543
Mechanical engineers	Architectural and engineering managers	.12	291,290	408,178
Food batchmakers	Industrial production managers	.12	151,950	12,729
Licensed practical and licensed vocational nurses	Medical and health services managers	.11	702,700	254,787
Food batchmakers	Heavy and tractor-trailer truck drivers	.11	151,950	12,729

This table shows the ‘thickest’ occupational transition paths from large occupations (defined as those with national employment greater than 150,000 in 2017). The transition share from occupation o to occupation p is defined as the share of all occupation leavers from the initial occupation o who move into that particular new occupation p. Only occupations with at least 500 observations in the BGT data and 2017 OES employment data are shown.

Table A3: Adj. R-squared from regressions of occupational relevance on characteristics

Dependent variable:	$\pi_{o \to p}$	
Included characteristic	No FE	Incl. origin SOC FE
Skill distance	0.011	0.025
Wages	0.003	0.021
Job amenities	0.021	0.039
Leadership	0.017	0.033
Skill composites	0.035	0.058

Table shows adjusted R-squared from regressions of the form

$$\pi_{o \to p} = \kappa + \alpha_o + \beta \Delta X_{occ\, p-o} + \epsilon_{op}.$$

Here, $\pi_{o \to p}$ is the share of job changers in the origin occupation o that move into target occupation p, and α_o are origin occupation fixed effects (included only in the second column). All regressions contain a constant. The variable $\Delta X_{occ\, p-o}$ represents the group of included characteristic differences noted in the table, which are included in relative target-minus-origin form and as absolute distances, with the exception of skill distance. All regressions are weighted by the average 2002-2015 national employment in the origin SOC. Note that the underlying occupational transition matrix is sparse, with many cells that show zero transitions, which is why the linear regression fit yields a relatively small R-squared.

Table A4: Regressions of wage on outside-occupation option index: aggregated occupation codes

Dependent variable:	Log wage			
	(1)	(2)	(3)	(4)
Panel A: Minor SOC Group (3-digit) regressions:				
OLS: oo^{occs}	0.184***	0.091***	0.106***	0.071***
	(0.013)	(0.007)	(0.011)	(0.009)
IV: oo^{occs}, instrumented	0.143***	0.113***	0.105***	0.125***
	(0.018)	(0.012)	(0.015)	(0.014)
Observations	486,487	486,481	486,487	485,815
Panel B: Major SOC Group (2-digit) regressions:				
OLS: oo^{occs}	-0.063**	0.081***	0.002	0.080***
	(0.030)	(0.009)	(0.024)	(0.009)
IV: oo^{occs}, instrumented	-0.182***	0.079***	0.060**	0.327***
	(0.038)	(0.028)	(0.029)	(0.113)
Observations	137,650	137,650	137,650	137,609
Fixed effects	Year	Occ-Year City	City-Year Occ-Year	Occ-Year Occ-City

Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses: $*p < .1, **p < .05, *** p < .01$. Units of observation are 2-digit or 3-digit SOC by city by year, for all observations with available data over 1999–2016 inclusive. As noted in the paper, 'cities' refers to CBSAs (metropolitan and micropolitan statistical areas) or NECTAs (New England city and town areas). Each cell reports the coefficient for the variable of interest in one specification, with included fixed effects held constant within each column.

Table A5: Regressions of wage on outside-occupation option index, sample split into three periods

Dependent variable:	Log wage			
	(1)	(2)	(3)	(4)
Panel A: 1999–2006:				
OLS: oo^{occs}	0.132***	0.071***	0.085***	0.029***
	(0.011)	(0.004)	(0.005)	(0.005)
IV: oo^{occs}, instrumented	0.100***	0.053***	0.060***	0.018***
	(0.011)	(0.005)	(0.006)	(0.004)
Observations	788,519	788,463	788,519	772,025
Panel B: 2007–2011:				
OLS: oo^{occs}	0.141***	0.091***	0.097***	0.035***
	(0.010)	(0.005)	(0.006)	(0.005)
IV: oo^{occs}, instrumented	0.130***	0.080***	0.084***	0.013***
	(0.010)	(0.006)	(0.007)	(0.005)
Observations	579,283	579,242	579,283	565,394
Panel C: 2012–2016:				
OLS: oo^{occs}	0.149***	0.096***	0.105***	0.026***
	(0.010)	(0.005)	(0.006)	(0.008)
IV: oo^{occs}, instrumented	0.145***	0.090***	0.094***	0.011
	(0.012)	(0.007)	(0.007)	(0.009)
Observations	576,568	576,525	576,568	562,592
Fixed effects	Year	Occ-Year City	City-Year Occ-Year	Occ-Year Occ-City

Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses: $*p < .1$,$**p < .05$,$***\ p < .01$. Units of observation are 6-digit SOC by city by year, for all observations with available data over 2002–2016 inclusive (split into three five-year periods). As noted in the paper, 'cities' refers to CBSAs (metropolitan and micropolitan statistical areas) or NECTAs (New England city and town areas). Each cell reports the coefficient for the variable of interest (outside-occupation option index) in one regression specification, with included fixed effects held constant within each column.

Table A6: Regressions of wage on outside-occ. option index, *employment-weighted*

Dependent variable:	Log wage			
	(1)	(2)	(3)	(4)
Panel A: OLS				
oo^{occs}	0.382***	0.094***	0.132***	0.027***
	(0.026)	(0.010)	(0.009)	(0.008)
Panel B: 2SLS				
oo^{occs}, instrumented	0.396***	0.096***	0.106***	0.026***
	(0.025)	(0.013)	(0.015)	(0.010)
First stage:				
Coeff. on $oo^{occs,inst}_{o,k,t}$	1.106***	0.886***	0.874***	0.923***
	(0.052)	(0.020)	(0.020)	(0.081)
1st-stage F-Stat.	451	1929	1943	129
Fixed effects	Year	Occ-Year	City-Year	Occ-Year
		City	Occ-Year	Occ-City
Observations	1,944,370	1,944,230	1,944,230	1,931,901

Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses: $*p < .1$,$**p < .05$,$*** \ p < .01$. Units of observation are 6 digit SOC by city by year, for all observations with available data over 1999–2016 inclusive. The instrumented outside-occupation option index uses the national leave-one-out mean wage in outside option occupations to instrument for the local (city-level) wage, and the initial local employment share in outside option occupations to instrument for the current local employment share. Observations are weighted by the average employment of their occupation-city over the sample period. Each cell reports the coefficient for the variable of interest (outside-occupation option index) in one regression specification, with included fixed effects held constant within each column.