

Employer Concentration and Outside Options

Gregor Schubert, Anna Stansbury, and Bledi Taska *

April 8, 2021

Abstract

We study the effect of within-occupation employer concentration and outside-occupation job options on wages in the US, identifying outside-occupation options using new occupational mobility data from 16 million resumes. Using shift-share instruments to identify plausibly exogenous local variation, we find that moving from the median to 95th percentile of employer concentration reduces wages by 3% on average and by 6% for workers in the lowest quartile of outward occupational mobility. Our findings imply that policymakers should take employer concentration seriously, but that measures of employer concentration – typically calculated for single occupations – should be considered alongside the availability of outside-occupation options.

* 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 titles “Monopsony and Outside Options”, and “Getting Labor Markets Right”. For their detailed comments and advice, we thank Justin Bloesch, Gabriel Chodorow-Reich, Karen Dynan, Richard Freeman, Xavier Gabaix, Ed Glaeser, Claudia Goldin, Benny Goldman, Larry Katz, Larry Summers, Maya Sen, and Betsey Stevenson. For helpful discussions, we thank David Balan, John Coglianese, Oren Danieli, David Deming, Martin Feldstein, Emma Harrington, Simon Jaeger, Robin Lee, Mike Lipsitz, Ioana Marinescu, Jeff Miron, Simon Mongey, Suresh Naidu, Nancy Rose, Isaac Sorkin, Elizabeth Weber Handwerker, Ron Yang, seminar participants at Harvard, MIT Sloan, Chicago Booth, Oxford, the Federal Trade Commission, the OECD, and the Peterson Institute for International Economics, and participants at the briq Workshop on Firms, Jobs and Inequality, the Federal Reserve System Community Development Research Conference, the IZA Summer School on Labor Economics, the IZA/CAIS Workshop on Matching in Online Labor Markets, the Urban Economic Association Meetings, the Wharton People and Organizations Conference, the Asian and Australasian Society of Labour Economics Conference, the Oxford NuCamp Virtual PhD Workshop, the Southern Economic Association meetings, and the Equitable Growth Grantee Conference. 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 recent years, concerns about employer concentration have increased. Employer concentration has been posited as a possible explanation for inequality, low pay, and stagnant pay growth. Antitrust authorities have been called on to consider employer concentration in the review of merger and acquisition cases. Concerns have been raised that employer concentration facilitates restrictions on competition like no-poaching agreements. And, since employer concentration can be a source of monopsony power,¹ concerns around high employer concentration have bolstered calls to raise minimum wages and strengthen collective bargaining.²

To assess whether – or in which cases – policy should respond to employer concentration, we need to understand the nature and effects of employer concentration in the US. In this paper, we seek to answer the question: *To what extent does employer concentration matter for US workers’ wages, and for whom does it matter the most?* We estimate the effects of employer concentration on average hourly wages across over 100,000 US SOC 6-digit occupation-by-metropolitan-area labor markets over 2013–2016, following Azar, Marinescu, Steinbaum and Taska (2020a) in measuring employer concentration with a Herfindahl-Hirschman Index (HHI) constructed from Burning Glass Technologies’ online job postings database. Our empirical strategy addresses two common empirical issues: endogeneity and market definition.

The first empirical issue is endogeneity. While recent research has documented a negative relationship between local employer concentration and wages, the extent to which this is causal – and the magnitude of any such causal effect – is unclear: employer concentration may be correlated with other local economic conditions which also affect wages, complicating the estimation of any underlying wage-concentration relationship.

To respond to this issue, we propose a new identification approach for the

¹Other possible sources of monopsony power include search frictions, switch costs, and worker and job heterogeneity (Robinson, 1933; Manning, 2003).

²Authors making the arguments in this paragraph include, variously, Bahn (2018); Shambaugh, Nunn, Breitwieser and Liu (2018); Krueger and Posner (2018); Naidu, Posner and Weyl (2018); Marinescu and Hovenkamp (2019); Marinescu and Posner (2020).

effects of employer concentration on wages, drawing on shift-share and granular IV methodology (Borusyak, Hull and Jaravel, 2018; Gabaix and Koijen, 2020). Specifically, we instrument for employer concentration within a particular local occupation with the predicted change in employer concentration, predicting each local employer’s hiring with its national hiring in that occupation (excluding the local area in question). This enables us to construct shocks to local employer concentration that are plausibly orthogonal to local productivity, with the key identifying assumption being that each large firm’s decision to increase or decrease its hiring nationwide is exogenous with respect to the local economic conditions in any specific local occupation.

The second empirical issue is market definition. Assessing the effect of local employer concentration on wages, and pinpointing the workers who are most affected by it, requires a good definition of the relevant local labor market for workers. Using new, highly-granular occupational mobility data constructed from 16 million US workers’ resumes (obtained by Burning Glass Technologies),³ we show that occupational mobility is high and highly heterogeneous across occupations. This suggests that regressing wages on within-occupation employer concentration – as much recent research does – without considering the availability of these outside occupation job options (1) may obscure heterogeneity, as some occupations are a better approximation of workers’ true labor market than others, and (2) may lead to biased estimates, as workers who are in high-concentration labor markets (*within* their local occupation) also tend to have poor local job options *outside* their occupation.

To respond to this issue, we introduce two new factors into our baseline regressions of wages on within-occupation employer concentration. First, we allow the estimated coefficient on within-occupation employer concentration to vary by occupations’ outward mobility. This enables us to estimate different effects of employer concentration on wages for low-mobility vs. high-mobility occupations (for whom the SOC 6-digit occupation is less likely to be a good

³The large sample size – an order of magnitude more than other data sources – enables us to estimate occupational transitions reliably between a large share of US occupations. We are making this new occupational mobility dataset publicly available.

approximation to their true labor market). Second, we develop a measure of the value of workers’ outside job options in other occupations – an “outside-occupation option index” – and estimate its effect on wages in our baseline regression alongside the effect of within-occupation employer concentration. Our outside-occupation option index is the weighted average of local wages in all occupations except the worker’s own, with each weight the product of: (i) occupational mobility flows to each outside occupation and (ii) the local relative employment share in each outside occupation.⁴ We use a shift-share IV approach to identify effects of changes in this outside-occupation option index on wages, instrumenting for local occupational wages with the leave-one-out national mean wage in outside option occupations.

How much does employer concentration matter for wages? Our baseline results suggest that moving from the median to the 95th percentile HHI (as faced by workers) results in 2.6 log points lower wages.⁵ This average masks substantial heterogeneity: within-occupation employer concentration matters substantially more for workers who are less able to find comparably good jobs in other occupations. For occupations in the bottom quartile of occupational mobility, like registered nurses and security guards, moving from the median to 95th percentile HHI is associated with between 4 and 8 log points lower wages; for occupations in the highest quartile of occupational mobility, like counter attendants or bank tellers, our point estimate is zero and the confidence interval rules out any decrease in wages greater than 1.8 log points. As expected, we find that regressions of wages on employer HHI suffer from omitted variable bias if the availability of outside-occupation options is not included in the analysis, with an upward bias in the coefficient size of about 30-40%. A back-of-the-envelope calculation, using our coefficient estimates, suggests that over 10% of the 110 million workers covered by our data experience wage suppression of 2% or more as a result of employer concentration.

⁴We show that this index can be derived from our conceptual framework in which workers’ wage depends on the expected value of their outside option.

⁵Our instrumental variable estimates are about 50% larger than our OLS estimates, suggesting that some combination of omitted variables or measurement error bias the coefficient towards zero in simple regressions of wages on employer concentration.

Many of the most-affected workers are healthcare workers, reflecting both high healthcare employment concentration and low occupational mobility.

We also find a positive and significant effect of an increase in the value of outside-occupation options, holding constant within-occupation employer concentration: for the median occupation, moving from the 25th to the 75th percentile value of outside-occupation options across cities is associated with 3.7 log points higher wages. These magnitudes are meaningful relative to the degree of geographic wage dispersion across cities: for the median occupation, moving from the 25th to the 75th percentile city by average wage is associated with a 21 log points higher wage. Overall, our paper demonstrates that the availability of job options outside workers' own firm – which is affected both by within-occupation employer concentration and the quality of outside-occupation job options – matters substantially for workers' own wages.

Related literature: We build on a growing body of work demonstrating an empirical relationship between wages and employer concentration, which began in recent years with Azar, Marinescu and Steinbaum (2020b), Azar et al. (2020a), Benmelech, Bergman and Kim (2018), and Rinz (2018),⁶ as well as a growing theoretical literature demonstrating an effect of employer concentration on wages (Berger, Herkenhoff and Mongey, 2019; Jarosch, Nimczik and Sorkin, 2019; Azkarate-Askasua and Zerecero, 2020). We make two contributions to this literature: we use a new instrument to estimate plausibly causal effects of employer concentration on wages, and we show that simple wage-concentration regressions which do not consider outside-occupation options are likely biased and obscure important heterogeneity.

Second, in estimating the effect of outside-occupation options on wages, we add to a literature on outside options in the labor market, including Beaudry, Green and Sand (2012), who show local spillovers from changes in industrial employment in the US; Caldwell and Danieli (2018), who find wage effects of workers' outside options in Germany, estimated from the diversity of jobs held

⁶As well as Lipsius (2018), Hershbein, Macaluso and Yeh (2019), Gibbons, Greenman, Norlander and Sørensen (2019), and Qiu and Sojourner (2019) in the US, Abel, Tenreyro and Thwaites (2018) in the UK, Marinescu, Ouss and Pape (2019) in France, Martins (2018) in Portugal, and Dodini, Lovenheim, Salvanes and Willén (2020) in Norway.

by similar workers; Macaluso (2019), who shows that the skill mix of local employment affects laid-off workers' outcomes in the US; and Alfaro-Urena, Manelici and Vasquez (2020), who estimate the outside option value of jobs at multinational corporations in Costa Rica. Third, in using occupational transitions to identify outside options we build on papers which use worker flows to identify the scope of workers' labor markets (Manning and Petrongolo, 2017; Nimczik, 2018), and to study skill similarity across occupations and industries (Shaw, 1987; Neffke, Otto and Weyh, 2017; Arnold, 2020).

Finally, we contribute to a broader literature on imperfect competition in labor markets, including both the literature on labor market monopsony and the elasticity of the labor supply curve to the firm (e.g. Boal and Ransom, 1997; Manning, 2003; Azar, Berry and Marinescu, 2019a; Berger et al., 2019; Azkarate-Askasua and Zerecero, 2020), and the large search-and-matching literature which features outside options in the worker-firm wage bargain (e.g. Burdett and Mortensen, 1980; Cahuc, Postel-Vinay and Robin, 2006).⁷

2 Conceptual framework

There are a number of models of the labor market in which employer concentration matters for wages. First, employer concentration can generate upward-sloping labor supply curves to individual firms, leading to wage markdowns (e.g. Berger et al., 2019).⁸ Second, the presence of a small number of firms facilitates collusion to suppress wages. Third, in a bargaining model of the labor market, employer concentration reduces the number of feasible outside options for workers, as the average worker in a given labor market has few distinct firms as alternative possible employers: this reduces workers' relative bargaining position and therefore the wage (Jarosch et al., 2019).⁹ In this

⁷Additional work on monopsony includes the empirical estimates of the elasticity of labor supply to the firm in Webber (2015) and Sokolova and Sorensen (2020), and empirical analyses of specific industries, firms, or worker classes in Hirsch and Schumacher (2005), Staiger, Spetz and Phibbs (2010), Ransom and Sims (2010), Ashenfelter, Farber and Ransom (2013), Matsudaira (2014), Naidu, Nyarko and Wang (2016), Bassier, Dube and Naidu (2019), Goolsbee and Syverson (2019), and Dube, Jacobs, Naidu and Suri (2020).

⁸One aspect of monopsony power as initially discussed by Robinson (1933).

⁹Several recent papers specifically show that an HHI of employer concentration is a relevant statistic for firms' labor market power. Berger et al. (2019) show in a general

paper, the conceptual framework guiding our empirical analysis is the third: **employer concentration worsens workers’ outside options**. While the main focus of our paper is empirical, we outline below a stylized framework where we formalize this intuition (developed further in Appendix A).

Wage bargaining. At the start of each period, each employed worker Nash-bargains with her employer i . The outcome is wage w_i , equal to the value of the worker’s outside option if she leaves her job, oo_i , plus share β of the match surplus (where p_i is the product of the match):¹⁰

$$w_i = \beta p_i + (1 - \beta) oo_i \quad (1)$$

Job search. Once bargaining with incumbent workers has concluded, firms post vacancies to fill empty positions – vacated either by labor force exit, or by bargaining breakdown with an existing worker. Each posted vacancy offers a wage equal to the wage the firm has bargained with its incumbent workers. Job seekers – workers whose previous wage bargain broke down, or who newly entered the labor force – are each paired randomly with a vacancy within their labor market. (For now, consider a situation where all workers and all jobs are perfectly substitutable.) Random matching means that the chance of any given worker receiving a job offer from a particular firm j is equal to firm j ’s share of vacancies in the labor market, σ_j (in the spirit of Burdett and Mortensen (1980) and Jarosch et al. (2019)). If a worker does not receive any job offers, she moves to unemployment for the period and receives b .

Outside option value. The outside option for an employed worker is to leave her current job and become a job seeker. Her outside option value is therefore a weighted average of the wages paid by each local firm (weighted

equilibrium oligopsony model that firms with higher market share have lower labor supply elasticities, and that the wage-bill HHI is a relevant statistic for assessing the welfare effects of firms’ labor market power. Azkarate-Askasua and Zerecero (2020) also show that firms’ employment shares affect the elasticity of labor supply to the firm, in a model which also features worker bargaining power. Arnold (2020) and Naidu and Posner (2020) show that an employer HHI is related to the size of the wage markdown under Cournot competition. Jarosch et al. (2019) show that in a search model with bargaining, the presence of large employers worsens workers’ outside option and so reduces wages, with the effect determined by a concentration index closely related to an HHI.

¹⁰We assume all workers at the same firm i have the same outside option.

by the probability of being matched with each feasible firm), and the unemployment benefit (weighted by the probability of receiving no job offers):¹¹

$$oo_i = \sum_j \sigma_j \cdot w_j + \sigma_i \cdot b. \quad (2)$$

Average wage. To a second order approximation, we can therefore express the average wage in the labor market ($\bar{w} = \sum_i \sigma_i w_i$) as a function of the sum of the squares of employer shares ($HHI = \sum_i \sigma_i^2$):

$$\bar{w} = \psi \cdot \bar{p} + (1 - \psi) \cdot b - \beta(1 - \beta) \sum_i \sigma_i^2 \hat{p}_i \quad (3)$$

where coefficient $\psi = 1 - (1 - \beta)HHI$ is a function of employer concentration HHI and worker bargaining power β , $\bar{p} = \sum_i \sigma_i p_i$ is average productivity across firms, and $\hat{p}_j = p_j - \bar{p}$ is the difference between firm j 's productivity and the market average. This expression illustrates that (1) the wage declines as employer concentration increases, since workers are less likely to receive job offers from other firms, and (2) there is an interaction between employer concentration and worker bargaining power, since the outside option matters less if workers have more bargaining power over the match surplus.¹²

Outside-occupation job options. The framework above assumes a clearly-delineated labor market, with all workers and jobs equally substitutable. This is typically the approach taken in theoretical and empirical work on employer concentration and wages (and the direct analog of the market definition approach in antitrust in product markets).¹³ This is not, however, how most

¹¹Note that this implies the wage at firm i will depend *negatively* on its share of the labor market σ_i . Why? In a labor market with atomistic firms, every job seeker would be matched with a feasible employer each period. But in a labor market with some large employers, the chance that a worker is re-matched with the firm she just quit is non-zero, and we assume this firm does not re-hire her. This means that the probability that a worker at firm i receives a job offer from *any other* firm if she leaves her job is proportional to the share of all jobs which are outside her firm, $(1 - \sigma_i)$. This intuition can be translated to a less stylized setting with on-the-job search, where workers can only receive outside job offers from firms which are *not* their own firm.

¹²The third term depends on the joint distribution of employment shares and productivity. If their correlation is sufficiently small the wage is simply a concentration- and bargaining power-weighted average of productivity p and unemployment benefit b .

¹³Labor markets have typically been defined as a single occupation or industry within a

labor markets work in practice. We therefore now adapt our earlier framework, defining the baseline labor market as a worker's local occupation and incorporating workers' option to switch occupation. The value of the outside option is still the weighted average of the wage in each other local firm, but the weight is now a product of two factors: (i) the probability that a worker from occupation o will receive *any* job offer from a firm in occupation p , which we denote $Prob(o \rightarrow p)$, and (ii) the vacancy share of firm j in occupation p , $\sigma_{j,p}$. The outside option for workers in firm i and occupation o is therefore:¹⁴

$$oo_{i,o} = \underbrace{Prob(o \rightarrow o) \cdot \sum_{j \neq i} \sigma_{j,o} \cdot w_{j,o}}_{\text{own-occ options}} + \underbrace{\sum_{p \neq o}^{N_{occs}} Prob(o \rightarrow p) \sum_l^{N_p} \sigma_{l,p} \cdot w_{l,p}}_{\text{outside-occ options}} + \underbrace{Prob(o \rightarrow o) \sigma_{i,o} \cdot b}_{\text{unemployment}}$$

Define $oo^{occs} = \sum_{p \neq o}^{N_{occs}} Prob(o \rightarrow p) \sum_l^{N_p} \sigma_{l,p} \cdot w_{l,p}$, the value of outside-occupation job options. We can then express the average wage in occupation o , to a second order approximation, as:

$$\bar{w}_o = \tilde{\psi}_o (\alpha \bar{p}_o + (1 - \alpha) oo^{occs}) + (1 - \tilde{\psi}_o) b - \tilde{p}_o \quad (4)$$

where $\tilde{\psi}_o = 1 - (1 - \beta) Prob(o \rightarrow o) HHI_o$, $\alpha = \frac{\beta}{1 - Prob(o \rightarrow o)(1 - \beta)}$, and $\tilde{p}_o = \beta(1 - \beta) Prob(o \rightarrow o) \sum_i \sigma_i^2 \hat{p}_{i,o}$. This expression suggests that the average wage in occupation o is, roughly, a weighted average of average productivity in occupation o (\bar{p}_o), the value of jobs outside occupation o (oo_o^{occs}), and unemployment benefits (b), where the weights depend on employer concentration within workers' occupation, the likelihood a worker will remain in her own occupation, and worker bargaining power.¹⁵ As before, higher within-occupation employer concentration reduces the wage. There is also an interaction with $Prob(o \rightarrow o)$: the less likely it is that a worker can find a job in a different

given local area (commuting zone, metropolitan area, or county), and debate has focused on how narrow an occupational or industrial definition to draw (e.g. Azar et al., 2020b,a). Jarosch et al. (2019) and Dodini et al. (2020) define local labor markets more flexibly as clusters of firms inferred using worker flows or common skill requirements (respectively), but still use a binary concept of the labor market.

¹⁴This assumes that employment decisions are taken at the firm-by-occupation level.

¹⁵Ignoring the final term \tilde{p}_o , which is once again small if the average productivity of individual firms is not strongly correlated with their employment shares

occupation, the more employer concentration in her own occupation matters.

Implications for empirical analysis. Typically, regressions of wages on employer concentration at the level of a local labor market take the binary market definition approach: they do not take into account workers’ (differential) ability to switch occupation. The discussion above illustrates two problems that this could cause when estimating the effect of concentration on wages. First, one should expect the effect of within-occupation employer concentration on wages to be different for low-mobility vs. high-mobility occupations. Second, if the degree of employer concentration within a local occupation (HHI) is correlated with the quality of outside options *outside* the local occupation (oo^{occs}), then estimation of the effect of concentration on the wage may be biased without controlling for outside-occupation options.

In our empirical analysis, we take both of these into account: we regress local average occupational wages on *both* within-occupation employer concentration *and* a measure of the value of outside-occupation options, allowing the coefficients to vary according to an occupation’s degree of mobility.

3 Measuring employer concentration

In our empirical analysis, we will jointly estimate the effect of both within-occupation employer concentration and outside-occupation job options on wages. We follow Azar et al. (2020a) and Hershbein et al. (2019) in using Burning Glass Technologies’ (“BGT”) database of online vacancy postings to measure employer concentration.¹⁶ We calculate the Herfindahl-Hirschman Index (HHI) of each employer’s share of vacancy postings within individual SOC 6-digit occupations and metropolitan areas, in each year 2013–2016:

$$HHI_{o,k,t} = \sum_{i=1}^N \left(\frac{v_{i,o,k,t}}{\sum_{i=1}^N v_{i,o,k,t}} \right)^2 \quad (5)$$

¹⁶Why use vacancies rather than employment data? First, we are not able to obtain firm-level employment data within local occupations. Second, vacancies may be a better reflection of workers’ feasible outside options than employment. In equilibrium, one would expect vacancy and employment HHIs to be highly correlated. Indeed, Marinescu et al. (2019) show in France that an HHI of employment flows (reflecting filled vacancies) is highly correlated with employment HHIs.

where $v_{i,o,k,t}$ denotes the number of vacancy postings from employer i in occupation o and metropolitan area k in year t . The BGT vacancy posting data covers the near-universe of online job postings, drawn from over 40,000 distinct sources including company websites and online job boards, with no more than 5% of vacancies from any one source (Hazell and Taska, 2019).¹⁷

Representativeness. Since the Burning Glass Technologies vacancy data covers the near-universe of online job postings, it is relatively representative of the vacancies which are advertised online. There are however two reasons for concern. First, not all vacancies are posted online. Azar et al. (2020a) estimate that in 2016, the BGT vacancy database captured around 85% of all job vacancies both online and offline (as measured from the Help Wanted Online database), but this is likely substantially lower for certain occupations where a large share of jobs are advertised offline or informally.¹⁸ Second, in occupations where firms tend to hire many workers for each posted vacancy, our estimates of employer concentration will be biased to the degree that larger firms may hire a higher number of people per vacancy posting.¹⁹

To understand the degree to which each of these might be an issue, we calculate a measure of ‘represented-ness’ of each occupation in the BGT data: the occupation’s share of vacancy postings in the BGT database relative to the occupation’s share of total employment (as per BLS OES). By this metric, occupations which are particularly underrepresented include low-wage food service jobs, cleaners, home health aides, laborers, and cashiers. In our estimates of the effect of employer concentration on wages, we carry out a number of

¹⁷Each vacancy posting contains the job title, company name, location, date, and job description. Using proprietary parsing technology, BGT imputes a SOC 6-digit occupation code. More details on the process by which BGT obtains, parses, and deduplicates this data can be found in Carnevale, Jayasundera and Repnikov (2014). To identify jobs at the same employer, we largely group jobs by employer name. We discuss the data and our process for identifying employers in detail in Appendix B.

¹⁸If the missing vacancies disproportionately come from small firms or households, which seems likely, we will overestimate employer concentration for underrepresented occupations.

¹⁹If there is a big difference in the ratio of hires per job posting between large and small firms, we will underestimate concentration in labor markets with skewed employer size distributions, relative to those with more symmetric employer size distributions. Our measures of employer concentration may therefore be less reliable for occupations for which there are many large employers who hire a lot of workers for undifferentiated job roles.

sensitivity checks to account for underrepresentation of certain occupations.²⁰ For further discussion of the BGT vacancy data, see Appendix B.

4 Identifying outside-occupation options

To jointly estimate the effect of *both* within-occupation employer concentration *and* outside-occupation job options on wages, we also need a measure of outside-occupation options. We use data on workers’ mobility patterns between occupations to identify these.²¹

BGT resume data. Since there is no existing US occupational mobility data with high enough granularity to study transitions between SOC 6-digit occupations, we construct a new data set of occupational transitions using 16 million unique US resumes, which enable us to observe longitudinal snapshots of workers’ job histories over 2002–2018.²² This resume data was collected by labor market analytics company Burning Glass Technologies (“BGT”), who sourced the resumes from a variety of BGT partners, including recruitment and staffing agencies, workforce agencies, and job boards.

Transition share. We use this data to construct our measure of occupational transitions $\pi_{o \rightarrow p}$, which approximates the probability of a worker moving from occupation o to occupation p conditional on leaving her job:²³

$$\begin{aligned}\pi_{o \rightarrow p} &= \frac{\# \text{ in occ } o \text{ in year } t \text{ observed in occ } p \text{ in year } t + 1}{\# \text{ in occ } o \text{ in year } t \text{ observed in a new job in year } t + 1} \\ &\approx \text{Prob}(\text{move from occ } o \text{ to occ } p | \text{leave job})\end{aligned}\tag{6}$$

²⁰We also control for occupation-by-year fixed effects, which should assuage concerns about the relative representativeness of the data for different occupations.

²¹We see occupational mobility patterns as a transparent, non-parametric way to capture the value of a different occupation as an outside option, since they capture a combination of both feasibility and desirability. In Appendix E we compare our approach to approaches based on task or skill similarity.

²²The CPS has at least an order of magnitude fewer occupational transition observations than our data over the same time period. This matters: with 705,600 possible transition pairs between the SOC 6-digit occupations, data sets with even a few million observations are not big enough to capture many transition paths.

²³Specifically, $\pi_{o \rightarrow p}$ is the share of people observed in occupation o at some point in year t who are also observed in occupation p at some point in year $t + 1$, as a fraction of all those in occupation o in t who are observed in a new job at some point in $t + 1$. We exclude jobs lasting 6 months or less. Our measure includes people with jobs in two different occupations at the same time – implicitly assuming that this indicates viability as an outside option.

Leave share. We also construct the ‘occupation leave share’, approximating the share of people who leave their occupation when they leave their job:²⁴

$$\begin{aligned} \text{leave share}_o &= \frac{\# \text{ in occ } o \text{ in year } t \text{ \& no longer in occ } o \text{ in year } t + 1}{\# \text{ in occ } o \text{ in year } t \text{ \& in a new job in year } t + 1} \\ &\approx \text{Prob}(\text{leave occ } o | \text{leave job}) \end{aligned} \quad (7)$$

We calculate these as averages across the whole US and across all years in our data, to capture the underlying degree of occupational similarity rather than transitory fluctuations from year to year.²⁵

Representativeness. 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 use this data set to estimate occupational transitions paths from one occupation to another, 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. The over-representation of younger workers, however, might be a concern if younger workers tend to be more mobile or to have different occupational mobility patterns than older workers. We therefore adjust for the over-representation by age by re-weighting our observed occupational transitions to match the distribution of employment by age within each U.S. occupation, provided by the BLS for 2012-2017. We discuss the BGT resume data in more detail in Appendix C.

Five facts about occupational mobility. Is it sensible to use occupational transitions to infer the scope of workers’ labor markets? Are job options outside the narrow occupation even an important part of most workers’ labor markets? Using the BGT data, we document five stylized facts about occupa-

²⁴Specifically, this measure captures the share of people observed in occupation o in year t who are *no longer* observed in occupation o at any point in year $t + 1$, as a share of those observed in occupation o in year t who are observed in some new job in year $t + 1$.

²⁵We estimate transition probabilities $\pi_{o \rightarrow p}$ and leave shares 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.

tional mobility which suggest the answer to both questions is yes.

1. Occupational mobility is high, suggesting that the SOC 6-digit occupation fails to capture many workers' true labor markets: the average probability of a worker leaving her 6-digit occupation when she leaves her job - the "occupation leave share" defined above - is 23%.
2. Mobility is heterogeneous across occupations, suggesting that the SOC 6-digit occupation is a better approximation of the labor market for some occupations than others: a quarter of occupations have a leave share lower than 19%, and a quarter higher than 28% (Table 1, Figure 2).²⁶
3. Aggregating up the SOC classification hierarchy - which groups ostensibly similar occupations - still fails to capture most occupational transitions, suggesting that this cannot solve the market definition problem.²⁷
4. The occupational transition matrix is sparse, suggesting that workers' relevant labor markets are mostly comprised of only a few occupations, and is highly asymmetric, suggesting that the relevance of occupations as outside options is not symmetric across occupation pairs (unlike in many task- and skill-based measures of occupational similarity).²⁸
5. Empirical occupation transitions reflect similarities between occupations in terms of their task requirements, wages, amenities, and leadership responsibilities, suggesting that occupational transitions do indeed reflect the underlying feasibility of an occupation as an outside option.²⁹

²⁶Almost all of the occupations with low leave shares are highly specialized, including various medical, legal and educational occupations (see Appendix Table A4). In contrast, many high leave share occupations require more general skills, including restaurant hosts/hostesses, cashiers, tellers, counter attendants, and food preparation workers.

²⁷For 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). For example, only 39% of systems software developers leave their 2-digit occupation group when they move across 6-digit occupations, compared to 95% of flight attendants. Note that management roles are often considered a separate 2-digit occupational group from non-management roles in the same field. Excluding transitions to and from management, at the median 67% of SOC 6-digit occupational transitions cross SOC 2-digit boundaries.

²⁸See Appendix Figure A7 and Appendix Table A5). The asymmetry partly reflects the fact workers in an occupation with specialized skills may be able to move to occupations which require generalist skills (e.g. retail salespersons) but the reverse flow is less feasible.

²⁹To show this, we regress our measure of occupational transitions on a number of different occupational characteristics derived from the O*Net database: the vector difference in

Measuring outside-occupation options. The facts above suggest that we can use occupational transitions to infer workers’ outside option occupations. Informed by our conceptual framework, we therefore define a measure of the outside option value of jobs in other occupations - the **outside-occupation option index** oo^{occs} - as a weighted average of the wage in each alternative occupation, weighted by a measure of the likelihood that the worker will move to a job in each of those alternative occupations if she leaves her current job: $oo_{o,k,t}^{occs} = \sum_{p \neq o}^{N_{occs}} Prob(o \rightarrow p)_{o,k,t} \cdot w_{p,k,t}$, where subscripts refers to the city (k), current occupation (o), possible destination occupations (p), and year (t). To proxy for $Prob(o \rightarrow p)$, we use the product of two variables: (1) the national average empirical occupation transition share $\pi_{o \rightarrow p}$, and (2) the relative employment share of occupation p in city k compared to the national average, $\frac{s_{p,k}}{s_p}$.³⁰ Our empirical outside-occupation option index is therefore:

$$oo_{o,k,t}^{occs} = \sum_{p \neq o}^{N_{occs}} \pi_{o \rightarrow p} \cdot \frac{s_{p,k,t}}{s_{p,t}} \cdot \bar{w}_{p,k,t} \quad (8)$$

We construct this outside-occupation option index for each year 1999-2016 for as many SOC 6-digit occupations and cities as our data allows (using the BLS Occupational Employment Statistics (OES) to obtain relative employment shares $\frac{s_{p,k,t}}{s_{p,t}}$ and average wages $\bar{w}_{p,k,t}$).³¹

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, Levy and Murnane (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. In every pairwise regression of occupational mobility on the absolute difference in characteristics (controlling for the difference in wages), the coefficients are significantly negative or statistically insignificant, as shown in Figure 3. Similarly, Macaluso (2019) finds that mobility between U.S. SOC 2-digit occupations is highly correlated with task similarity. See Appendix F for more details on our analysis.

³⁰The national occupation transition share proxies for the likelihood that, nationwide, the average worker’s best job option outside her occupation would be in each other occupation p ; the local relative employment share adjusts this for the local availability of jobs in each occupation p .

³¹See summary stats of our index in Table 2. Note: 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. Of the possible 786,335 occupation-city

5 Empirical Approach

Our analysis thus far suggests four testable predictions: (1) higher employer concentration reduces wages, (2) better outside-occupation options increase wages, (3) the wage-HHI relationship is stronger for occupations with limited outward mobility, and (4) the estimated wage-HHI relationship may be biased if within-occupation HHI is correlated with outside-occupation job options. In this section, we lay out our approach to evaluate these predictions.

Regression specification. Our baseline specification regresses the log of the average hourly wage in a SOC 6-digit occupation, metropolitan area (“city”), and year, on the log of the HHI of employer concentration ($HHI_{o,k,t}$), the log of our outside-occupation option index ($oo_{o,k,t}^{occs}$), and a set of occupation-by-year and city-by-year fixed effects ($\alpha_{o,t}$, $\alpha_{k,t}$):

$$\ln \bar{w}_{o,k,t} = \alpha + \alpha_{o,t} + \alpha_{k,t} + \gamma_1 \ln HHI_{o,k,t} + \gamma_2 \ln oo_{o,k,t}^{occs} + \xi_{o,k,t} \quad (9)$$

We also run this regression allowing the coefficients γ_1 and γ_2 on the HHI and outside-occupation option index to vary according to the occupation’s degree of outward mobility in our BGT resume data (the “occupation leave share”), interacting $HHI_{o,k,t}$ and $oo_{o,k,t}^{occs}$ with an indicator variable for the applicable quartile of outward mobility of occupation o .

We run these regressions across the largest possible subset of U.S. occupation-city-year cells for which we can obtain all our data: Our full data set for our baseline regressions over 2013–2016 comprises 212,417 occupation-city-year observations.³² We use BLS OES data for average hourly wages by occupation, city and year for the dependent variable $\bar{w}_{o,k,t}$. As discussed above, the HHI is constructed from BGT vacancy posting data and the outside-occupation option index is constructed from wage and employment data from BLS OES,

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. To create a consistent panel of occupations over time we crosswalk SOC classifications over time: see Appendix D.

³²This includes 387 cities and 715 occupations, with 94,809 occupation-city labor markets appearing in at least one year from 2013–2016. We have data on the wage, HHI, and outside-occupation option index (*but not* the instruments) for a larger set of occupation-city-year labor markets. We calculate summary statistics and counterfactuals on this larger set.

and occupational transition shares from BGT resume data.³³

Identification: employer concentration. Endogeneity issues may bias the estimated coefficients on the HHI. The direction of the bias is ambiguous: an increase in employer concentration could reflect the expansion of a highly productive large firm, which would result in higher employer concentration (expected to reduce wages) but also higher average productivity (expected to increase wages). Or, an increase in employer concentration could reflect a lack of local dynamism, with few new firms, which may lead to higher employer concentration alongside falling productivity.³⁴

We therefore instrument for local labor market concentration, creating an instrumental variable which leverages differential local occupation-level exposure to national firm-level hiring, in a strategy which builds on both 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), and on the shift-share ‘Bartik’ approach. Our strategy is based on the facts that (a) increases in local employer concentration are often driven by individual large firms growing, (b) these firms usually operate across many labor markets, (c) local labor markets are differentially exposed to different large firms, and (d) the employment growth of these large firms nationally is likely orthogonal to economic conditions in a specific local occupation.

Specifically, we note that the growth in local employer concentration in occupation o is a function of the growth in local occupational employment for

³³Our estimates of the effect of employer concentration on wages could inform analyses based either on bargaining models or models where employer concentration affects the wage via the elasticity of labor supply to the firm. However, our outside-occupation option index is informed specifically by a bargaining model. (Note also: a bargaining framework has different implications from a monopsonistic framework for employment).

³⁴Concerns like these are raised in many of the critiques of the empirical literature which finds a negative correlation between local employer concentration and wages, including Berry, Gaynor and Scott Morton (2019) and Rose (2019). Rose (2019) argues that empirical strategies attempting to identify a causal effect of employer concentration on wages must isolate the effect of employer concentration from changes in labor demand; our identification strategy attempts to do this. Hsieh and Rossi-Hansberg (2019) show that over recent decades large national firms have expanded into more local labor markets, reducing local employer concentration and possibly increasing productivity.

each employer j , $g_{j,o,k,t}$ (leaving aside firm entry): $\Delta HHI_{o,k,t} = \sum_j \sigma_{j,o,k,t}^2 - \sum_j \sigma_{j,o,k,t-1}^2 = \sum_j \sigma_{j,o,k,t-1}^2 \left(\frac{(1+g_{j,o,k,t})^2}{(1+g_{o,k,t})^2} - 1 \right)$. The increase in local occupational employer 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 in the labor market $g_{o,k,t}$. We use the leave-one-out firm-level national vacancy growth in occupation o (which we denote $\tilde{g}_{j,o,t}$) to instrument for local firm-level vacancy growth in occupation o ($g_{j,o,k,t}$), constructing our instrument for the HHI, $Z_{o,k,t}^{HHI}$, as:

$$Z_{o,k,t}^{HHI} = \log \left(\sum_j \sigma_{j,o,k,t-1}^2 \left(\frac{(1 + \tilde{g}_{j,o,t})^2}{(1 + \tilde{g}_{o,k,t})^2} - 1 \right) \right) \quad (10)$$

where $\tilde{g}_{o,k,t} = \sum_j \sigma_{j,o,k,t-1} \tilde{g}_{j,o,t}$ is the predicted local growth rate in vacancies, as predicted from the national (leave-one-out) growth of hiring in occupation o by each local firm j .³⁵ The key assumptions for this instrument to be valid are that the firm's national leave-one-out vacancy growth is (i) correlated with its local vacancy growth, but (ii) uncorrelated with the determinants of occupation-specific productivity growth in any given local labor market k . Through the lens of shift-share instruments (Borusyak et al., 2018), our instrument features plausibly exogenous 'shocks' (a function of firms' national hiring growth), and possibly endogenous exposure 'shares' (the last-period local occupational vacancy shares of each of those firms).³⁶

One concern with this instrumental variable is that differential local exposure to national firms' growth may differentially affect total labor demand, not just employer concentration. In our model, the effect of a large firm's growth on local labor market concentration is quadratic, whereas the effect of a large

³⁵Note that by taking the log of the instrument, we implicitly exclude observations where the predicted change in HHI based on national firm-level growth is negative. Note also that we are instrumenting for the local level of the HHI with an instrument derived from an expression for the change in the HHI.

³⁶For intuition, a hypothetical example: assume that in Bloomington IL, State Farm has a large employment share of insurance sales agents, while in Amarillo TX employment is dispersed across several insurance companies. In years where State Farm grows substantially faster than other insurance companies nationwide, under most assumptions about how that growth is allocated geographically, employer concentration of insurance sales agents will grow by more in Bloomington IL than in Amarillo TX.

firm’s growth on local labor demand or productivity is linear. Thus, we control for (1) the growth rate of local vacancies in the occupation-city labor market ($g_{o,k,t} = \sum_j \sigma_{j,o,k,t-1} g_{j,o,k,t}$), and (2) the predicted growth rate of local vacancies based on large firms’ national growth (i.e. the direct linear analog to our concentration index: $\tilde{g}_{o,k,t} = \sum_j \sigma_{j,o,k,t-1} \tilde{g}_{j,o,t}$ as defined above). With these controls, 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 linear effect on local labor demand or productivity.³⁷

A second concern is bias due to the fact that our exposure ‘shares’ do not sum to one. As such, following Borusyak et al. (2018) we introduce an “exposure control”: the sum of the squared vacancy shares of local firms j which have *any* growth in other metro areas $\sum_j \sigma_{j,o,k,t-1}^2 \cdot \mathbb{1}[\tilde{g}_{j,o,t} \neq 0]$.³⁸ We further discuss identification conditions in Appendix G.

Third, our instrument is unlikely to be strong for small changes in employer concentration in initially unconcentrated labor markets – if each firm has only a trivial share of local employment, even substantial hiring growth will not much change local employer concentration. We therefore apply our estimates of the effect of employer concentration on wages only to local labor markets with above-median employer concentration.

Notwithstanding these caveats, we see our approach as a substantial new contribution with regard to the problem of estimating the effect of employer concentration on wages. Some recent empirical work instruments for changes in employer concentration in a given local occupation with changes in (the inverse of) the number of employers in the same occupation in other local areas

³⁷Controlling for national trend exposure directly to prevent it from confounding a non-linear IV is similar to the “double Bartik” approach in Chodorow-Reich and Wieland (forthcoming). While the assumption that these linear terms capture demand effects is relatively strong, note that their inclusion does not affect our baseline coefficient estimates (Appendix Tables A9 and A11). The most plausible demand-related threat to our identification is that the growth of a large national firm, locally, pushes up wages by more in areas with inelastic labor supply to the local occupation than in areas with elastic labor supply to the local occupation, and that the elasticity of labor supply to the local occupation is correlated with the initial local employment share of the large national firm in that occupation.

³⁸This controls for the fact that different local occupations may have different initial shares of employment accounted for by large national firms.

(e.g. Azar et al. (2020a,b); Rinz (2018); Qiu and Sojourner (2019); Marinescu et al. (2019); Gibbons et al. (2019)). This circumvents some endogeneity concerns, but a concern remains that national occupation trends in concentration may be correlated with unobservable national trends in occupational productivity, demand, or supply, which could confound estimated wage effects.³⁹ Our strategy allows us to use occupation-year fixed effects to control for national occupation-level factors which affect wages. Other recent empirical work uses M&A activity to generate plausibly exogenous variation in local labor market concentration, including Arnold (2020) for all industries, and Prager and Schmitt (2019) for hospital mergers. This avoids endogeneity concerns about the cause of the change in concentration, but reflects one specific source of concentration (M&A activity accounts for less than 2% of changes in local employer concentration (Arnold, 2020)) and cannot fully isolate the effects of employer concentration from other local economic effects of the M&A activity. Our approach allows us to examine the effects of various sources of variation across broad swathes of the US labor market, and to control at least somewhat for effects on local labor demand.⁴⁰ Ultimately, we believe that this set of complementary identification approaches – based off different variation, and with different strengths – can together provide a useful picture of the effects of employer concentration on wages.

Identification: outside-occupation options. Endogeneity issues may also bias the coefficients on our outside-occupation option index: a positive local demand shock for an occupation similar to a worker’s own may come at the same time 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 problem: if occupation p and occupation o are good outside options for each other, then a wage increase in o will increase wages in p and vice versa. To identify causal effects, we need exogenous shocks to the wages in workers’ outside-occupation options which do not affect, and

³⁹The authors control for variables like labor market tightness to address this.

⁴⁰Dodini et al. (2020) adopt an additional different strategy, demonstrating that workers laid-off in mass layoffs see larger wage losses in more concentrated labor markets in Norway.

are not affected by, the local wages in their own occupation.

We use a ‘Bartik’ shift-share approach, instrumenting for local wages in each outside option occupation p in city k with the leave-one-out national mean wage for occupation p excluding its wage in city k ($\bar{w}_{p,k,t}$). We also instrument for the local relative employment share in each occupation using the initial employment share in that occupation in 1999, the first year for which we have data ($\frac{s_{p,k,1999}}{s_{p,1999}}$).⁴¹ Our instrument for the oo^{occs} index is:

$$Z_{o,k,t}^{oo} = \sum_p^{N_{occs}} \left(\pi_{o \rightarrow p} \cdot \frac{s_{p,k,1999}}{s_{p,1999}} \cdot \bar{w}_{p,k,t} \right) \quad (11)$$

The identifying variation within a given occupation across different cities comes from differences in each city’s initial occupational employment composition. 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.⁴² For our instrument to be valid, the national leave-one-out mean wage $\bar{w}_{p,k,t}$ in outside option occupation p must be correlated with the local wage of occupation p in location k (relevance condition), but must not affect the local wage in initial occupation o through a direct channel other than increasing the quality of local outside-occupation options (*conditional* on controlling for occupation-year and city-year fixed effects).⁴³ We discuss conditions for identification further in Appendix G, following the approach to shift-share IVs of Borusyak et al. (2018).

⁴¹Or the first year the occupation-city is in the data, if it is not present in 1999.

⁴²That 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). This instrumental variable strategy is closely related to that of Beaudry et al. (2012), who use national industry wage premia to substitute for city-level industry wages when estimating spillover effects of cities’ industrial composition.

⁴³The inclusion of these fixed effects means that differences in city-level trends or national productivity of different occupations do not represent an issue for our identification strategy. An additional concern may be that groups of local occupations that share similar labor markets experience similar location-specific industry shocks. We show that our results are robust to controlling for common exposure to industry shocks (see Appendix Table A9).

6 Results: employer concentration

There is a robust negative correlation between log vacancy HHIs and log wages at the occupation-city level (Figure 4). In a regression with occupation-year and city-year fixed effects, the OLS relationship is strongly statistically significant, with a coefficient of -0.010 (Table 3, column *a*).⁴⁴

When instrumenting for the HHI, the coefficient magnitude increases by around 40% relative to the OLS specification (Table 3, column *c*). This suggests some combination of omitted variable bias or measurement error biasing the coefficient toward zero in simple OLS regressions of wages on HHI.⁴⁵ We also introduce a control for our outside-occupation option index (Table 3, columns *b* and *d* for OLS and 2SLS respectively). The coefficient on the instrumented outside-occupation option index is positive and highly statistically significant, confirming that outside-occupation job options matter for wages. After introducing the outside-occupation option index the coefficient on the HHI falls by about a third in the OLS regressions and about a fifth in the IV regressions, consistent with omitted variable bias. This is because the vacancy HHI is negatively correlated with workers' outside-occupation options: workers with worse options *within* their occupation also have worse options *outside* their occupation (as illustrated in Appendix Figure A13).

How big is the average effect of employer concentration on wages? Our baseline coefficient estimates in Table 3, column *d* – instrumenting for both employer concentration and outside-occupation job options – suggest that going from the HHI faced by the median worker to the HHI faced by the worker at the 95th percentile (from an HHI of 221 to 2,256) would be associated with a 2.6 log points lower hourly wage.⁴⁶ This is at the low end of the estimates reviewed in Marinescu and Hovenkamp (2019).⁴⁷ As noted, we caution against

⁴⁴Similarly Hershbein et al. (2019) find a coefficient of -0.014, regressing wages on vacancy concentration in local SOC 6-digit occupations over 2010–2017. Other papers' estimates are not directly comparable because of different labor market definitions or wage measures.

⁴⁵The first stage is shown in Table A7, column (a).

⁴⁶Calculated as $(\ln(2256) - \ln(221)) \cdot -0.011 = -0.026$.

⁴⁷Reviewing existing evidence, they suggest that a 10% increase in employer concentration (at the SOC 6-digit occupation by commuting zone level) leads to a 0.3% to 1.3% decrease in wages. Our point estimate in Table 3, column *d* suggests a 10% increase in concentration

applying these coefficient estimates to labor markets with very low initial levels of employer concentration, since our instrument is weak in these cases.⁴⁸

Heterogeneity by occupational mobility. Re-running our baseline regression, but allowing the coefficients on the HHI and outside-occupation option index to vary for occupations with different degrees of outward occupational mobility, we find that the average effect of within-occupation employer concentration on wages conceals substantial heterogeneity (Table 4, Figure 5).⁴⁹ For the quartile of occupations with the lowest outward mobility, as proxied by our occupation “leave share”, our coefficient estimate suggests that going from the median to the 95th percentile HHI faced by workers would be associated with 5.8 log points lower wages.⁵⁰ For the quartile of occupations with the highest outward mobility, the point estimate is zero, and the confidence interval suggests that an equivalent increase in the HHI would be associated with at most a 1.8 log point lower wage.⁵¹

Robustness checks. We explore a number of additional variations on our

(at the SOC 6-digit occupation by MSA level) leads to a 0.11% decrease in wages on average.

⁴⁸Figure A15 illustrates that the correlation between our (log) HHI instrument and the (log) HHI for occupation-city labor markets in 2016 somewhat breaks down for occupation-city cells with a very low value of the HHI or of our HHI instrument. Similarly, regressing wages only on occupation-city cells with very low HHIs results in a much weaker instrument (with, for example, a Kleibergen-Paap F-statistic of 14 for occupation-city cells with HHIs less than 50) and no significant relationship between wages and HHI.

⁴⁹In Appendix Tables A14 and A15, we explore whether there is evidence for heterogeneity of the effect of concentration on wages for different quartiles of the wage distribution, or for different occupation groups, but do not find patterns which are clearly statistically significantly different across these groups. We also do not find conclusive evidence of non-linearity in the concentration-wage relationship when we estimate the effect separately by quartile of initial employer HHI (Appendix Table A13).

⁵⁰Calculated as $(\ln(2256) - \ln(221)) \cdot -0.025 = -0.058$.

⁵¹Calculated as: $(\ln(2256) - \ln(221)) \cdot 1.96 \cdot -0.004 = -0.018$. The pattern of our results are consistent with Prager and Schmitt (2019), who find that hospital mergers which induce large increases in concentration reduce nursing and pharmacy workers’ wages substantially, somewhat suppress wages of non-medical hospital professionals, and have no detectable effect on wages for the remainder of hospital workers (in maintenance and repairs, operations, housekeeping, catering, and medical records). They interpret these differentials as reflecting the degree to which workers have industry-specific skills. Our estimates would similarly suggest that nursing and pharmacy workers would experience substantially higher wage effects of employer concentration than maintenance, housekeeping, and catering workers, since the former tend to have lower occupational mobility.

baseline analyses (Figure 6, Appendix Tables A9-A12). First, we show that our coefficient estimates are similar if we remove the controls for vacancy growth (Tables A9 and A11 column *a*), if we remove the exposure control (column *c*), or if we follow Gabaix and Koijen (2020) in adding an additional control for the equal-weighted vacancy growth of local firms ($g_{o,k,t}^e = \frac{1}{N} \sum_j g_{j,o,k,t}$) to reflect common local occupation-specific shocks (column *b*). Second, we control for an industry Bartik shock to control for common exposure to national industry trends (column *d*), and similarly find little change in our estimates. Third, we re-run our baseline regressions with fixed effects for occupation-city and year, rather than occupation-year and city-year (column *e*). Here, identifying variation comes from year-to-year changes in employer concentration within the same occupation-city labor market. In the full sample, we see a statistically significant negative point estimate of the same order of magnitude as our baseline estimates, but estimates by quartile of outward occupational mobility are noisy (which may not be surprising given the short time period). Fourth, we show that our coefficient estimates are similar if we weight each occupation-city cell by employment or log employment (Appendix Tables A10 and A12, columns *a* and *b*), assuaging concern that our results may have been driven by small occupations or cities. Fifth, to address concerns about representativeness in our BGT vacancy data, we re-run our baseline regression estimates *excluding* any occupations which are substantially underestimated in the vacancy data (column *c*),⁵² finding coefficient estimates almost identical to those in our baseline regressions, and also weight each occupation or metropolitan area respectively by its represented-ness in the BGT data (columns *d* and *e*).

7 Results: outside-occupation options

We can also use our baseline regressions to ask: how big are the effects of outside-occupation options on wages? Our baseline coefficient estimate (Table 3 column *d*) suggests that a 10 log point higher outside-occupation option index is associated with roughly 1 log point higher wages in workers' own occupation. This implies that moving from the 25th to the 75th percentile

⁵²We exclude occupations with a 'represented-ness' less than 0.5 in the BGT vacancy data, corresponding to about one third of occupations.

value of outside-occupation options across cities for the median occupation leads to 3.7 log points higher wages.⁵³ This is quite large in the context of the geographic variation of wages: for the median occupation, the interquartile range of average wages across cities in 2016 was 21 log points.⁵⁴ In addition, finding a large, significant, and positive effect of shocks to outside-occupation options on wages reinforces our conclusions that workers’ true labor markets are broader than their narrow 6-digit SOC occupations, and that our “probabilistic” method of identifying relevant outside options can capture workers’ true labor markets relatively well.⁵⁵

Robustness checks. There may be concerns our coefficients are biased by exposure to correlated industry shocks which affect both a workers’ own occupation and her outside option occupations.⁵⁶ To control for this possibility we construct a shift-share “industry Bartik” shock and include it in our baseline regressions.⁵⁷ The industry Bartik shock has a significantly positive effect on wages, but our coefficients on the outside-occupation option index are only slightly attenuated (Appendix Tables A9 and A11, column *d*).

⁵³We calculate this by estimating the interquartile range of the outside-occupation option index for each occupation across cities in 2016, taking the median across occupations, and applying our coefficient estimate of 0.95. With the same exercise for employer concentration, we find that for the median occupation, moving from the 25th to the 75th percentile HHI across cities would result in a wage increase of 1.4 log points.

⁵⁴For a specific example where outside-occupation options might be relevant, consider Baltimore, MD, and Houston, TX. They are a similar size with a similar average hourly wage, but statisticians in Baltimore earned 26% more than statisticians in Houston in 2016. Applying our baseline coefficient estimate suggests that around 20% of this difference may be attributable to differential availability of outside-occupation job options.

⁵⁵While we do not consider the effects of outside-city options on wages in this paper, our methodology could easily be extended to do so. The wage effect of local employer concentration and outside-occupation options is limited by workers’ option to move.

⁵⁶For example, if (1) the finance industry and the tech industry are disproportionately likely to employ both accountants and data scientists, (2) San Francisco has a large share of tech employment while New York has a large share of finance employment, and (3) being a data scientist is a good outside option occupation for an accountant, then in years where tech is booming nationwide, this will impact SF more than NY. Accountants in SF will see wages rising by more than accountants in NY but this may be driven simply because more accountants in SF already work in tech.

⁵⁷The shock is constructed such that the exposure of occupation o in city k to each industry ι is defined as the employment share of industry ι in occupation o nationwide, multiplied by the employment share of industry ι in city k . See Appendix G.

In addition, while our HHI data only covers 2013–2016, we can calculate our outside-occupation option index from 1999 onwards. Over this longer period, we find large, positive, and significant effects of outside-occupation options on wages, even with both occupation-by-city and occupation-by-year fixed effects (Appendix Table A16). We also find large effects if we calculate the outside-occupation option index using occupational mobility at the SOC 2-digit or 3-digit level instead of 6-digit level (Appendix Table A17); if we split our analysis into three time periods (Appendix Table A18); if we weight by employment (Appendix Table A19), or if we control for local occupational employment (Appendix Table A20).

8 Discussion and Implications

What might our results suggest about the aggregate effects of employer concentration? We use our coefficient estimates in a rough back-of-the-envelope quantification of the “wage effect” of employer concentration in each above-median HHI labor market in 2016, relative to a scenario where their HHI is reduced to 200 (roughly the median in our data in 2016),⁵⁸ as follows:

$$\text{wage effect}_{o,k,t} = (\log(HHI)_{o,k,t} - \log(200)) \cdot \gamma_1^q \quad (12)$$

where γ_1^q denotes the estimated coefficient on the $\log(HHI)$ in our baseline regression specification in Table 4 column d , for the appropriate quartile q of outward occupational mobility. Note that this exercise considers the effect of changes in employer concentration *holding all else constant*, including local productivity. It can illustrate the degree to which wages may be marked down from local occupational productivity as a result of employer concentration, but cannot necessarily illustrate what would happen if a specific policy or business decision were to change local employer concentration (as it might also change local productivity). It also rests on the assumption that we can apply our estimated coefficients linearly.

⁵⁸An HHI of 200 might correspond to roughly 50 equal-sized employers, or two large employers each with 10% of workers and an atomistic ‘fringe’ of firms employing the rest. This level of concentration is not typically thought to be a concern in product markets. Note: There may be monopsony power even in unconcentrated labor markets arising from employer heterogeneity or search frictions (Naidu and Posner, 2020).

Roughly 49 million of the 110 million workers in our data set are in occupation-city labor markets with an HHI greater than 200. Of these, our counterfactual wage exercise suggests that roughly 3 million workers have wages which are 5% lower or more as a result of above-median employer concentration, and a further 9 million have wages which are 2%-5% lower. This suggests that perhaps 10% of the US labor market sees wage suppression of 2% or more as a result of employer concentration.⁵⁹ We show in Table 5 the average estimated wage effect for different combinations of employer concentration and outward occupational mobility. Our estimates suggest that employer concentration slightly widens inequality: the share of workers affected by employer concentration is smaller than average in the highest quartile of the wage distribution and in high-wage cities (Appendix Figures A18 and A19). Note that our estimates focus only on wages, but employer concentration may also affect non-wage benefits and workplace amenities.⁶⁰

Which occupations are most affected by employer concentration? In Table 6, we list the twenty-five occupations with the largest number of workers who see an estimated wage effect of 2% or greater in their local occupational labor market (excluding occupations which are substantially under-represented in the BGT vacancy data). A large share of the most-affected occupations are in healthcare, including more than one million registered nurses, licensed practical and vocational nurses, and nursing assistants, and around 200,000 pharmacists and pharmacy technicians.⁶¹ According to our estimates, large

⁵⁹This may be an overestimate since the figure include some occupations which are underrepresented in the BGT vacancy posting data. Excluding all occupations with a ‘represented-ness’ of less than 0.5 in the BGT vacancy data, there remain 7.1 million workers (6.5% of workers in our data) whose wages are suppressed by at least 2% as a result of employer concentration. On the other hand, our data only covers 110 million of the 140 million nonfarm employees in 2016, and those not represented in our data are disproportionately in non-metropolitan areas or small occupations. One might expect these workers to face greater wage suppression from employer concentration than the average in our data.

⁶⁰Qiu and Sojourner (2019) and Marinescu, Qiu and Sojourner (2020) find negative relationships between employer concentration and the receipt of employment-based health insurance, and labor rights violations respectively.

⁶¹This is in keeping with recent work that has found large effects of hospital mergers on wages of nursing and pharmacy workers (Prager and Schmitt, 2019), and a low elasticity of the labor supply of registered nurses to individual hospitals (Staiger et al., 2010).

numbers of security guards and hairdressers, hairstylists, and cosmetologists are also affected by employer concentration, as large shares of their local labor markets are comprised of employment by a few large companies or chains (although, note that both occupations are somewhat underrepresented in our vacancy data).⁶²

Our back-of-the-envelope exercise suggests that while employer concentration suppresses wages for several million workers, the majority of American workers likely do not experience significant wage suppression as a result of employer concentration. Thus, policymakers should focus attention on the subset of workers who face both concentrated labor markets within their occupation and limited opportunities for occupational mobility.

Implications: antitrust. One area where this can be done is antitrust.⁶³ Marinescu and Hovenkamp (2019) and Naidu et al. (2018) argue that antitrust authorities should use measures of employer concentration as a preliminary screen for anticompetitive effects of mergers in labor markets (as they already do in product markets). Our analysis suggests that this screen should involve two variables: the HHI in a local 6-digit SOC occupation, and the degree of outward mobility from that occupation.⁶⁴

⁶²Note also that we consider employer concentration at the level of a salon chain, many of which are franchised: one might argue that it is better to consider employer concentration at the level of individual franchised salons, though non-compete and no-poaching agreements may make this distinction moot in practice.

⁶³Several scholars have called for antitrust authorities to pay attention to employer concentration (Marinescu and Hovenkamp, 2019; Naidu et al., 2018; Hemphill and Rose, 2017; Steinbaum and Stucke, 2020; Hovenkamp, 2018; Krueger and Posner, 2018). Historically antitrust authorities paid little attention to employer concentration (though monopsony is referred to in the 1992 DoJ-FTC Horizontal Merger Guidelines (Phillips, 2019)), but this has changed in recent years: the topic has featured in FTC and DoJ hearings, the FTC is expanding its retrospective merger review to scrutinize labor market power, and the FTC raised concerns about wage suppression for nurses in a September 2020 public comment on a proposed hospital merger in Hendrick TX.

⁶⁴The screen should also evaluate whether employer concentration in outside-occupation options will be affected by the merger – a concern in occupations whose outside options are predominantly in the same industry, like healthcare. Our proposal differs slightly from Marinescu and Hovenkamp (2019), who argue that antitrust authorities should screen for anti-competitive effects of mergers based only on the HHI in a local SOC 6-digit occupation. Our finding that the wage suppressive effects of employer concentration are at least four times higher for low outward mobility occupations suggests that screening based only on

However, it is important to note that our findings *do not* tell us that *all* increases in employer concentration reduce wages. In some cases, if higher employer concentration comes alongside higher productivity, workers’ wages may be higher in the high-concentration high-productivity scenario than a low-concentration lower-productivity scenario, so seeking to reduce employer concentration may not be the best response: close scrutiny of individual cases, and industry- and occupation-specific studies, are necessary to understand whether antitrust action would be appropriate in any specific circumstance.⁶⁵ In addition, while increased antitrust scrutiny of labor markets is important, it is unlikely to affect the majority of workers impacted by employer concentration (Naidu and Posner, 2020), since most changes in employer concentration are not caused by mergers and acquisitions and many concentrated labor markets do not feature illegal anti-competitive practices.

Implications: policy to raise wages. In many cases, rather than seeking to reduce employer concentration it may be more appropriate to recognize the fact that employer concentration may give large firms scope to pay a wage which is marked down relative to productivity – and to design labor market policies to counteract this. One such way to do this might be equipping workers with countervailing power by bolstering support for collective bargaining.⁶⁶ An

local within-occupation HHI without considering outward occupational mobility will lead to some mergers being scrutinized which may have little effect on wages, while others which may have serious anti-competitive effects may go unnoticed.

⁶⁵As emphasized by Hovenkamp (2018), Berger et al. (2019), and Arnold (2020). Naidu et al. (2018) argue that antitrust authorities should permit mergers where the incremental increase in workers’ wages because of increased productivity would *outweigh* any incremental decrease in workers’ wages induced by the increase in employer concentration.

⁶⁶In our conceptual framework outlined in section 2, higher worker bargaining power β reduces the weight placed on the outside option in the wage bargain and therefore reduces the importance of employer concentration in wage determination. There is some suggestive empirical evidence that labor markets with higher unionization rates see smaller effects of employer concentration on wages. When we re-run our baseline regression of the effect of employer concentration on wages with an interaction with states’ right-to-work status – a proxy for the ease of forming a union – we find larger effects of both employer concentration and outside-occupation options on wages in right-to-work states (Appendix Table A21). Prager and Schmitt (2019) find larger effects of hospital mergers on nursing wages when nursing unionization rates are lower and in right-to-work states, and Benmelech et al. (2018) find a stronger relationship between employer concentration and wages in U.S. manufacturing firms where unionization rates are lower.

alternative might be strengthening minimum wages or benefits standards in local labor markets characterized by high employer concentration.⁶⁷

Implications: promoting mobility. Our results suggest that employer concentration within a local occupation matters less if workers can find similarly good jobs outside their occupation – and, most likely, if workers can easily move geographically. This suggests that policies which make it easier to switch occupation and/or to work in different geographic areas may – by increasing workers’ outside options – reduce the degree to which employer concentration can suppress wages.⁶⁸ These could include reducing any disproportionate barriers to acquiring training, licensing, or certification in occupations, increasing reciprocal recognition of state-specific licenses and certifications, and increasing affordable housing supply in high-cost cities. In addition, restrictions on worker mobility *within* an occupation (like non-compete clauses) could exacerbate the effects of employer concentration on wages.⁶⁹

Incidence of employer concentration effects. Our estimates suggest that increases in employer concentration reduce local wages, but cannot tell us whether the incidence of these wage reductions falls on firms in the form of higher profits, or consumers in the form of lower prices (and the balance likely depends on the nature of product market competition). Similarly, Kahn and Tracy (2019) argue that the ultimate incidence of local labor market concentration falls to a large extent on local landowners as lower local wages reduce local rents and house prices. Understanding the ultimate incidence of these effects is important to determine the appropriate policy response.

⁶⁷Indeed, higher minimum wages, would be expected to have less of a negative effect on employment in labor markets where employers have monopsony power. Azar, Huet-Vaughn, Marinescu, Taska and Von Wachter (2019b) find that US labor markets with higher employer concentration see smaller employment effects of minimum wage increases.

⁶⁸Indeed, the decline in occupational and geographic mobility in the U.S., which may partly reflect an increase in costs of mobility, could be acting to increase the effects of employer concentration (Molloy, Smith and Wozniak, 2011; Xu, 2018).

⁶⁹See Johnson and Kleiner (2020) on the effect of state licensing standards on mobility, Ganong and Shoag (2017) on the effect of housing costs on mobility, and Starr, Prescott and Bishara (2019); Johnson, Lavetti and Lipsitz (2020) on the effect of non-competes.

9 Conclusion

Our findings point to a middle ground between two prominent views about the effects of employer concentration in the US labor market. On the one hand, employer concentration is *not* a niche issue confined to a few factory towns: we find large, negative, and significant effects of employer concentration on wages when estimated using nuanced market definitions and plausibly exogenous variation across the majority of the US labor market, and our back-of-the-envelope calculations suggest that perhaps 10% of the U.S. private sector workforce experiences non-trivial wage effects of employer concentration. On the other hand, most workers are not in highly concentrated labor markets, and the effects of employer concentration therefore do not seem big enough to have a substantial effect on the aggregate wage level or degree of income inequality in the U.S. economy (though other sources of monopsony power may still be important).⁷⁰ The fact that employer concentration affects wages for several million American workers suggests that increased policy attention to this issue is appropriate, in terms of antitrust, policies to raise wages, and policies to increase worker mobility. For these policy decisions, our work underscores that the definition of the labor market is vitally important.

⁷⁰Similarly, Rinz (2018), Berger et al. (2019), and Lipsius (2018) show that employer concentration has fallen over recent decades in most local industries, casting doubt on the argument that changing employer concentration can explain median pay stagnation or rising income inequality. It is possible, however, that the decline in countervailing worker power has exposed firms' latent monopsony power, meaning that employer concentration (and other sources of monopsony power) have greater wage effects than in the past (Erickson and Mitchell, 2007; Naidu et al., 2018; Stansbury and Summers, 2020).

References

- Abel, Will, Silvana Tenreyro, and Gregory Thwaites, “Monopsony in the UK,” 2018.
- Alfaro-Urena, Alonso, Isabela Manelici, and Jose P. Vasquez, “The Effects of Multinationals on Workers: Evidence from Costa Rica,” 2020.
- Arnold, David, “Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes,” 2020. Working paper.
- 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.
- Azar, José, Steven Berry, and Ioana Marinescu, “Estimating Labor Market Power,” 2019.
- Azar, José A, Ioana Marinescu, Marshall I Steinbaum, and Bledi Taska, “Concentration in US Labor Markets: Evidence From Online Vacancy Data,” *Labour Economics*, 2020, 66, 101886.
- Azar, José, Emiliano Huet-Vaughn, Ioana Marinescu, Bledi Taska, and Till Von Wachter, “Minimum wage employment effects and labor market concentration,” Technical Report, National Bureau of Economic Research 2019.
- , Ioana Marinescu, and Marshall Steinbaum, “Labor market concentration,” *Journal of Human Resources*, 2020, pp. 1218–9914R1.
- Azkarate-Askasua, Miren and Miguel Zerecero, “The Aggregate Effects of Labor Market Concentration,” *Unpublished Working Paper*, 2020.
- Bahn, Kate, “Understanding the importance of monopsony power in the U.S. labor market,” Technical Report, Washington Center for Equitable Growth 2018.
- Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, “Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies,” 2019. Working paper.
- 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?,” 2018.
- Berger, David W, Kyle F Herkenhoff, and Simon Mongey, “Labor Market Power,” 2019.
- Berry, Steven, Martin Gaynor, and Fiona Scott Morton, “Do increasing markups matter? lessons from empirical industrial organization,” *Journal of Economic Perspectives*, 2019, 33 (3), 44–68.
- 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.
- Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, “Wage bargaining with on-the-job search: A structural econometric model,” *Econometrica*, 2006, 74 (2), 323–364.
- Caldwell, Sydnee and Oren Danieli, “Outside Options in the Labor Market,” 2018. Working paper.
- Carnevale, Anthony P., Tamara Jayasundera, and Dmitri Repnikov, “Understanding Online Job Ads Data: A Technical Report,” 2014.
- 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.
- Dodini, Samuel, Michael F. Lovenheim, Kjell G. Salvanes, and Alexander Willén**, “Monopsony, Skills, and Labor Market Concentration,” *Unpublished Working Paper*, 2020.
- Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri**, “Monopsony in Online Labor Markets,” *American Economic Review: Insights*, 2020, 2 (1), 33–46.
- Erickson, Christopher L and Daniel JB Mitchell**, “Monopsony as a metaphor for the emerging post-union labour market,” *International Labour Review*, 2007, 146 (3-4), 163–187.
- Gabaix, Xavier and Ralph Koijen**, “Granular Instrumental Variables,” 2020. Working paper.
- Ganong, Peter and Daniel Shoag**, “Why has regional income convergence in the US declined?,” *Journal of Urban Economics*, 2017, 102, 76–90.
- Gibbons, Eric M., Allie Greenman, Peter Norlander, and Todd Sørensen**, “Monopsony Power and Guest Worker Programs,” 2019.
- Goldin, Claudia**, “A grand gender convergence: Its last chapter,” *American Economic Review*, 2014, 104 (4), 1091–1119.
- Goolsbee, Austan and Chad Syverson**, “Monopsony Power in Higher Education: A Tale of Two Tracks,” 2019.
- Hazell, Jonathon and Bledi Taska**, “Downward Rigidity in the Wage for New Hires,” 2019.
- Hemphill, C Scott and Nancy L Rose**, “Mergers that harm sellers,” *Yale LJ*, 2017, 127, 2078.
- Hershbein, Brad, Claudia Macaluso, and Chen Yeh**, “Concentration in U.S. local labor markets: evidence from vacancy and employment data,” *Working Paper*, 2019.
- Hirsch, Barry T. and Edward J. Schumacher**, “Classic or new monopsony? Searching for evidence in nursing labor markets,” *Journal of Health Economics*, 2005, 24 (5), 969–989.
- Hovenkamp, Herbert**, “Whatever Did Happen to the Antitrust Movement,” *Notre Dame L. Rev.*, 2018, 94, 583.
- Hsieh, Chang-Tai and Esteban Rossi-Hansberg**, “The industrial revolution in services,” Technical Report, National Bureau of Economic Research 2019.
- Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin**, “Granular Search, Market Structure, and Wages,” 2019. Working paper.
- Johnson, Janna E and Morris M Kleiner**, “Is Occupational Licensing a Barrier to Interstate Migration?,” *American Economic Journal: Economic Policy*, 2020.
- Johnson, Matthew S, Kurt Lavetti, and Michael Lipsitz**, “The labor market effects of legal restrictions on worker mobility,” *Available at SSRN 3455381*, 2020.
- Kahn, Matthew E and Joseph Tracy**, “Monopsony in spatial equilibrium,” Technical Report, National Bureau of Economic Research 2019.
- Krueger, Alan B. and Eric A. Posner**, “A Proposal for Protecting Low-Income Workers from Monopsony and Collusion,” 2018.
- Lipsius, Ben**, “Labor Market Concentration does not Explain the Falling Labor Share,” 2018. Working paper.
- Macaluso, Claudia**, “Skill remoteness and post-layoff labor market outcomes,” 2019.
- 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.
- Marinescu, Ioana Elena and Eric A. Posner**, “Why Has Antitrust Law Failed Workers?,” 2020.

- , **Ivan Ouss**, and **Louis-Daniel Pape**, “Wages, Hires, and Labor Market Concentration,” 2019.
- , **Yue Qiu**, and **Aaron Sojourner**, “Wage Inequality and Labor Rights Violations,” 2020.
- Martins, Pedro S.**, “Making their own weather? Estimating employer labour-market power and its wage effects,” 2018.
- Matsudaira, Jordan D.**, “Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations,” *The Review of Economics and Statistics*, 2014, 96 (1), 92–102.
- Molloy, Raven, Christopher L. Smith, and Abigail Wozniak**, “Internal migration in the United States,” *Journal of Economic Perspectives*, 2011, 25 (3), 173–196.
- Naidu, Suresh and Eric A. Posner**, “Labor Monopsony and the Limits of the Law,” 2020.
- , – , and **E. Glen Weyl**, “Antitrust Remedies for Labor Market Power,” *Harvard Law Review*, 2018.
- , **Yaw Nyarko**, and **Shing-Yi Wang**, “Monopsony power in migrant labor markets: evidence from the United Arab Emirates,” *Journal of Political Economy*, 2016, 124 (6), 1735–1792.
- 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.
- Phillips, Noah Joshua**, “Prepared Statement before the U.S. House of Representatives on Antitrust and Economic Opportunity: Competition in Labor Markets,” 2019.
- Prager, Elena and Matthew Schmitt**, “Employer Consolidation and Wages: Evidence from Hospitals,” 2019.
- Qiu, Yue and Aaron Sojourner**, “Labor-Market Concentration and Labor Compensation,” 2019.
- Ransom, Michael R. and David P. Sims**, “Estimating the Firm’s Labor Supply Curve in a “New Monopsony” Framework: Schoolteachers in Missouri,” *Journal of Labor Economics*, 2010, 82 (2), 331–355.
- Rinz, Kevin**, “Labor Market Concentration, Earnings Inequality, and Earnings Mobility,” *U.S. Census Bureau Carra Working Paper Series*, 2018, 2018-10.
- Robinson, Joan**, *The Economics of Imperfect Competition*, Springer, 1933.
- Rose, Nancy**, “Thinking Through Anticompetitive Effects of Mergers on Workers,” Technical Report, American Bar Association 2019 Antitrust Spring Meeting 2019.
- Shambaugh, Jay, Ryan Nunn, Audrey Breitwieser, and Patrick Liu**, “The State of Competition and Dynamism: Facts about Concentration, Start-Ups, and Related Policies,” 2018.
- Shaw, Kathryn L**, “Occupational change, employer change, and the transferability of skills,” *Southern Economic Journal*, 1987, pp. 702–719.
- Sokolova, Anna and Todd Sorensen**, “Monopsony in Labor Markets: A Meta-Analysis,” 2020.
- Staiger, Douglas O., Joanna Spetz, and Ciaran S. Phibbs**, “Is there monopsony in the labor market? Evidence from a natural experiment,” *Journal of Labor Economics*, 2010, 28 (2), 211–236.
- Stansbury, Anna and Lawrence H Summers**, “The Declining Worker Power Hypothesis: An explanation for the recent evolution of the American Economy,” *Brookings Papers on Economic Activity*, 2020.
- Starr, Evan, James J Prescott, and Norman Bishara**, “Noncompetes in the US labor force,” *U of Michigan Law & Econ Research Paper*, 2019, (18-013).
- Steinbaum, Marshall and Maurice E Stucke**, “The Effective Competition Standard: A New Standard for Antitrust,” *University of Chicago Law Review*, 2020.
- Webber, Douglas**, “Firm market power and the earnings distribution,” *Labour Economics*, 2015.
- Xu, Ming**, “Understanding the Decline in Occupational Mobility,” *Working Paper*, 2018.

Figures and Tables

Table 1: Summary statistics: BGT occupational mobility data

Percentile (occ.)	1	5	10	25	50	75	90	95	99
<i>Panel A: Number of obs. in the BGT occ. mobility data in '000s, by occ. (2002-2015)</i>									
Observations	0.6	1.1	1.6	4.9	20.8	112.3	466.8	853.9	3,471.9
<i>Panel B: Share leaving job and occupation, by occ. (2002-2015)</i>									
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
<i>Panel C: Share of occupational transitions which cross SOC 2d boundary (2002-2015)</i>									
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

Notes: We exclude occupations with <500 observations in the BGT resume data. In Panel A, an observation is a person-year unit that is also observed in the data the following year. Panel B shows the share of workers observed in a new job or new occupation from one year to the next, and the “leave share”, defined in section 4 as the share leaving their occupation conditional on leaving their job. Panel C shows the share – by origin occupation – of all SOC 6-digit occupational transitions which also span SOC 2-digit boundaries. The percentiles refer to percentiles across occupations, such that (for example) the median occupation in our data has 20,800 observations (Panel A).

Table 2: Summary statistics: main data set

Percentile (occ.-city)	1	5	10	25	50	75	90	95	99
<i>Panel A: Employer concentration HHI (2016)</i>									
HHI	35	97	167	390	965	2,222	5,000	7,222	10,000
HHI (emp-wt)	10	22	33	81	221	536	1,357	2,256	5,556
<i>Panel B: Outside-occupation option index oo^{occs} (2016)</i>									
oo^{occs} (2016)	1.4	2.1	2.6	3.5	4.8	6.6	8.8	10.6	16.1
$\frac{oo^{occs}}{wage}$	0.03	0.06	0.09	0.15	0.23	0.34	0.45	0.53	0.74
$\frac{oo^{occs}}{wage}$, emp-wt	0.06	0.11	0.14	0.23	0.34	0.45	0.55	0.63	0.77
<i>Panel C: Occupation-city wages and employment (2016)</i>									
Employment	30	40	50	90	220	670	1,980	3,920	14,410
Mean hourly wage	9.05	10.50	11.94	15.40	21.42	31.08	44.07	53.68	90.50
Wage, emp-wt	8.97	9.94	10.99	13.39	18.33	30.28	45.10	56.42	80.50
<i>Panel D: national hourly wage distribution (2016) from BLS OES</i>									
Hourly wage	–	–	9.27	11.60	17.81	28.92	45.45	–	–

Notes: Panels A, B, and C show summary statistics for our main data set in 2016, calculated over all occupation-city-year cells for which we have wage data, a vacancy HHI, and an outside-occupation option index. This comprises 103,300 occupation-by-city labor markets and 109,366,900 workers (according to the BLS OES data). Panel D shows the national 10th, 25th, 50th, 75th, and 90th percentile of the hourly wage distribution according to the full BLS OES data set, for comparison.

Table 3: Regression of wage on HHI and oo^{occ} , full sample

<i>Dependent variable:</i>	Log wage			
	(a) OLS	(b) OLS	(c) 2SLS IV	(d) 2SLS IV
Log HHI	-0.010 (0.001)	-0.007 (0.001)	-0.014 (0.003)	-0.011 (0.003)
Log outside-occ. options		0.106 (0.007)		0.095 (0.009)
Vacancy growth			-0.001 (0.001)	-0.001 (0.001)
Predicted vacancy growth			-0.011 (0.010)	-0.013 (0.010)
Exposure control			0.008 (0.008)	0.004 (0.007)
Observations	184,411	184,411	184,411	184,411
F-Stat			705	379

Notes: 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. All regressions feature occupation-by-year and MSA-by-year fixed effects. Columns (a) and (b) show OLS regressions and columns (c) and (d) show 2SLS IV regressions (where both the log HHI and log outside-occ. option index are instrumented). The reported F-stat for the 2SLS IV regressions is the Kleibergen-Paap Wald F statistic. The vacancy growth and predicted vacancy growth variables are *rescaled* by dividing by 10 (so that 1% vacancy growth in a local area corresponds to a value of 0.001, rather than 0.01), such that the coefficient estimates can be seen in the table for most specifications. See text for detailed explanation of instruments and controls.

Table 4: Regression of wage on HHI and oo^{occ} , by quartile of occupation leave share

<i>Dependent variable:</i>	Log wage			
	(a) OLS	(b) OLS	(c) 2SLS IV	(d) 2SLS IV
Log HHI \times Q1 occ mobility	-0.015 (0.002)	-0.014 (0.002)	-0.024 (0.004)	-0.025 (0.005)
Log HHI \times Q2 occ mobility	-0.013 (0.001)	-0.007 (0.001)	-0.016 (0.003)	-0.010 (0.003)
Log HHI \times Q3 occ mobility	-0.011 (0.001)	-0.005 (0.001)	-0.013 (0.003)	-0.007 (0.004)
Log HHI \times Q4 occ mobility	-0.002 (0.001)	0.002 (0.001)	-0.005 (0.003)	0.000 (0.004)
Log outside-occ. options \times Q1 occ mobility		0.083 (0.009)		0.060 (0.011)
Log outside-occ. options \times Q2 occ mobility		0.109 (0.007)		0.100 (0.009)
Log outside-occ. options \times Q3 occ mobility		0.115 (0.007)		0.107 (0.010)
Log outside-occ. options \times Q4 occ mobility		0.114 (0.007)		0.109 (0.010)
Vacancy growth			-0.001 (0.001)	-0.001 (0.001)
Predicted vacancy growth			-0.012 (0.010)	-0.016 (0.010)
Exposure control			0.008 (0.008)	0.004 (0.008)
Observations	184,411	184,411	184,411	184,411
F-stat			164	87

Notes: 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. All regressions feature occupation-by-year and MSA-by-year fixed effects. Columns (a) and (b) show OLS regressions and columns (c) and (d) show 2SLS IV regressions. The reported F-stat for the 2SLS IV regressions is the Kleibergen-Paap Wald F statistic. The vacancy growth and predicted vacancy growth variables are *rescaled* by dividing by 10 (so that 1% vacancy growth in a local area corresponds to a value of 0.001, rather than 0.01), such that the coefficient estimates can be seen in the table for most specifications. Independent variables labelled “ \times Q i outward mobility” show the coefficient on an interaction term between the HHI or outside-occupation option index (respectively) with an indicator variable which takes the value 1 if the occupation in question is in the i th quartile of outward occupational mobility (where “Q1” represents the least outwardly mobile occupations, and so on). See text for detailed explanation of variables.

Table 5: Counterfactual wage effects of setting HHI to 200 (& number affected)

	2,500< HHI <10,000	1,500< HHI <2,500	500< HHI <1,500	200< HHI <500	0< HHI <200
Lowest mobility	7.8% (1.2m)	5.8% (1.5m)	3.6% (5.2m)	1.1% (5.2m)	0 (8.0m)
Q2 mobility	3.1% (1.3m)	2.3% (1.2m)	1.4% (5.2m)	0.4% (7.6m)	0 (18.3m)
Q3 mobility	2.2% (1.1m)	1.6% (0.9m)	1.0% (4.3m)	0.3% (9.4m)	0 (19.0m)
Q4 mobility	0 (1.4m)	0 (1.1m)	0 (4.6m)	0 (7.2m)	0 (5.6m)

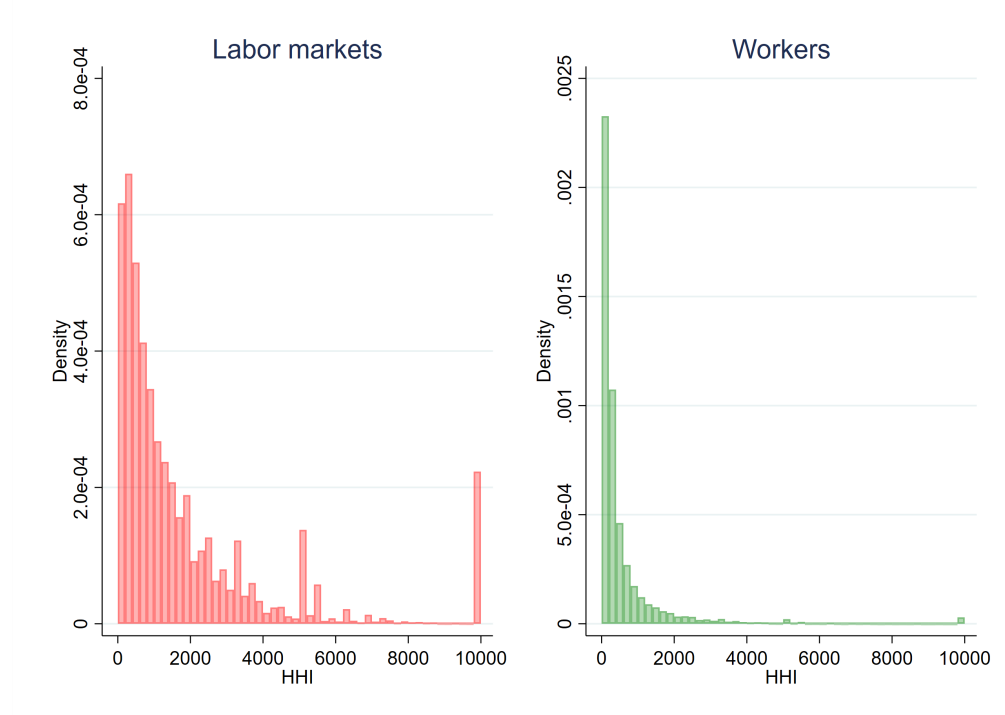
Notes: This table shows the estimated wage impact (and number of people affected in parentheses) of lowering the HHI to 200 in all occupation-city cells where it was greater than 200 in 2016. 200 is roughly the median HHI as experienced by workers. The estimated wage impact is calculated as the difference between the *actual* log HHI and the log of 200, multiplied by the estimated coefficient in our wage-HHI regressions (with the coefficient used corresponding to the appropriate quartile of occupational outward mobility, as estimated in Table 3 column (d)). The impact number in each cell in the table is the average impact across all workers in that cell: so, for example, for the 1.2 million workers in our data who are in occupations in the lowest quartile of outward mobility, and who are in occupation-city labor markets with an HHI greater than 2500, the *average* estimated wage impact of employer concentration on their wage is 7.8%. Note (1) this exercise implicitly holds productivity constant, and (2) our data set covers around 110 million workers in total, from the BLS OES occupation-by-metropolitan area employment and wage data.

Table 6: Twenty-five occupations with most people affected by employer concentration (based on a predicted occupation-CBSA wage effect of 2% or greater)

Occupation	National employment	Share of occupation w/ estimated effect 2% or greater	Number in occupation w/ estimated effect 2% or greater
Security guards	1,103,120	.88	969,230
Registered nurses	2,857,180	.19	545,520
Nursing assistants	1,443,150	.34	488,040
Hairdressers, hairstylists, and cosmetologists	352,380	.88	311,500
Nonfarm animal caretakers	187,360	.86	161,540
Fitness trainers and aerobics instructors	257,410	.62	158,310
Childcare workers	569,370	.25	144,290
Licensed practical and licensed vocational nurses	702,400	.19	131,840
Radiologic technologists	200,650	.64	128,600
Pharmacists	305,510	.38	117,360
Emergency medical technicians and paramedics	244,960	.46	112,290
Medical and clinical laboratory technologists	166,730	.55	92,510
Phlebotomists	120,970	.75	90,780
Pharmacy technicians	398,390	.21	84,360
Medical assistants	623,560	.13	81,300
Respiratory therapists	126,770	.62	78,360
Management analysts	637,690	.11	68,230
Lawyers	619,530	.11	68,200
Librarians	129,350	.52	66,730
Dentists, general	105,620	.62	65,730
Software developers, applications	794,000	.082	64,990
Bakers	180,450	.36	64,300
First-line supervisors of personal service workers	190,420	.32	61,720
Real estate sales agents	151,840	.4	60,190
Massage therapists	95,830	.61	58,820

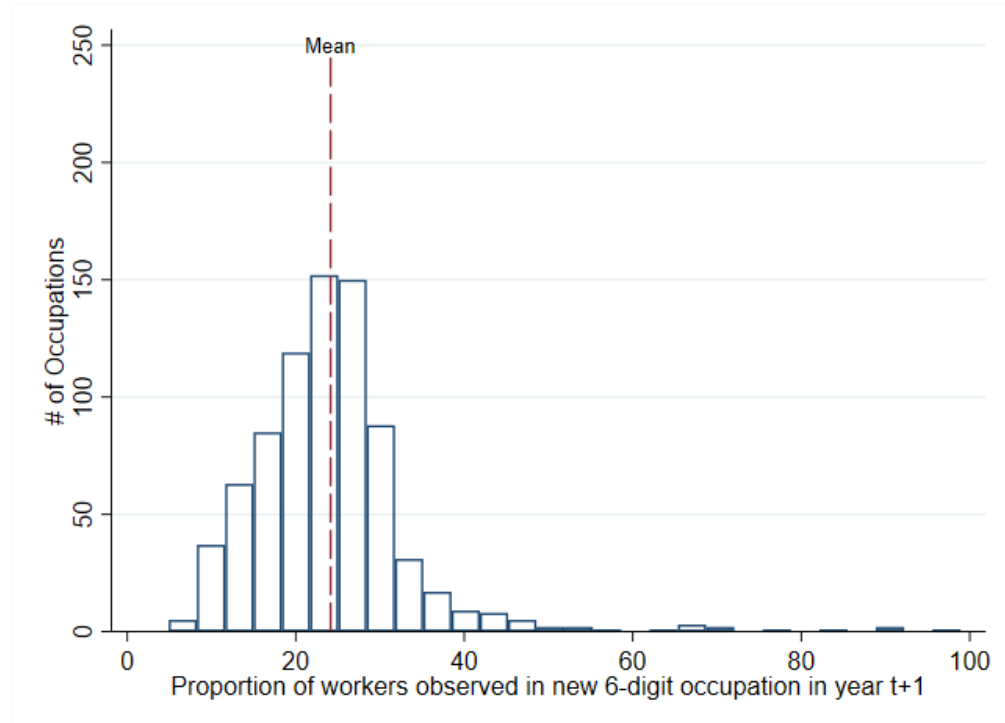
Notes: This table lists the twenty-five occupations with the largest number of workers who experience an estimated wage impact of 2% or more as a result of employer concentration (see Table 5 for description of how this effect is calculated). We exclude “teachers and instructors, all other” since the classification of this group in the BGT data relative to OES may not be fully comparable, and since there is a large share of public sector employment of teachers. We also exclude occupations that are very under-represented in the BGT vacancy data relative to overall employment (with a cutoff with represented-ness < 0.5, or around the 33rd percentile) – which, from this list, includes Personal care aides, Bartenders, Farmworkers and laborers (crop, nursery, and greenhouse), Dental hygienists, Lifeguards and ski patrol, Operating engineers and other construction equipment operators, Information and record clerks (all other), and Business operation specialists (all other), as well as some overwhelmingly public sector occupations (School bus drivers, Firefighters, Police officers, and Postal service mail carriers). In Appendix Table A23 we list the degree of represented-ness of each of these occupations in the BGT vacancy data.

Figure 1: Histogram of employer HHI across occ-city labor markets and across workers, 2016



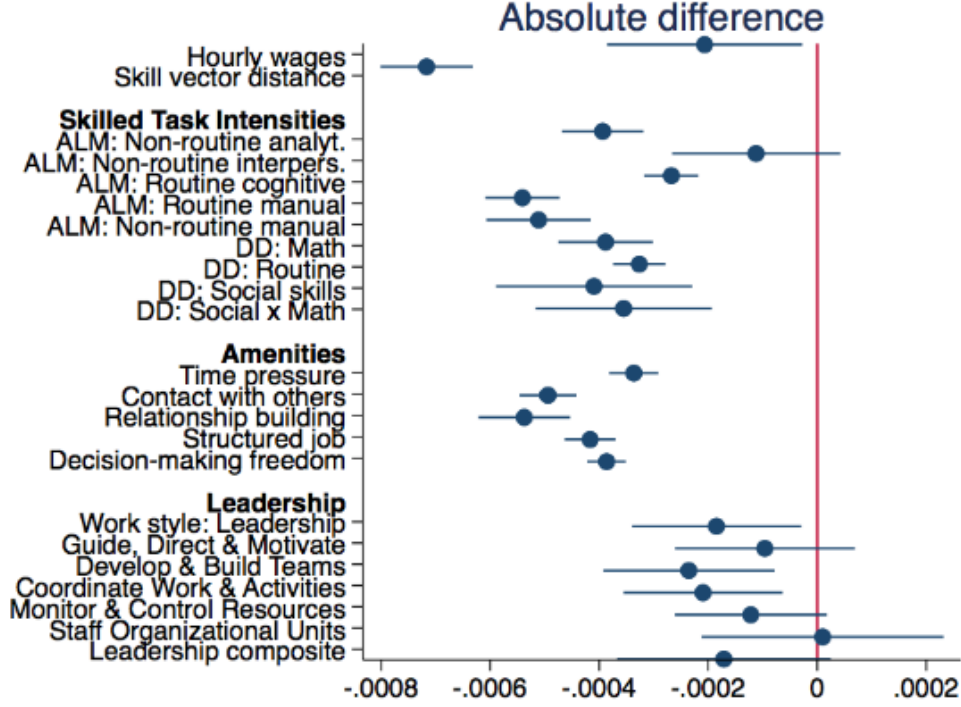
Note: HHI is measured using Burning Glass Technologies vacancy data, at the level of a SOC 6-digit occupation by metro area labor market. Our data covers occupation-metro area labor markets which include 110m of the 140m workers in the U.S. labor market in 2016. Left panel shows the distribution of HHIs across occ-city labor markets in 2016. Right panel shows the distribution of HHIs across workers in 2016 (i.e. the distribution of HHIs across occ-city labor markets, weighted by employment in each of these labor markets).

Figure 2: Outward occupational mobility from SOC 6-digit occupations



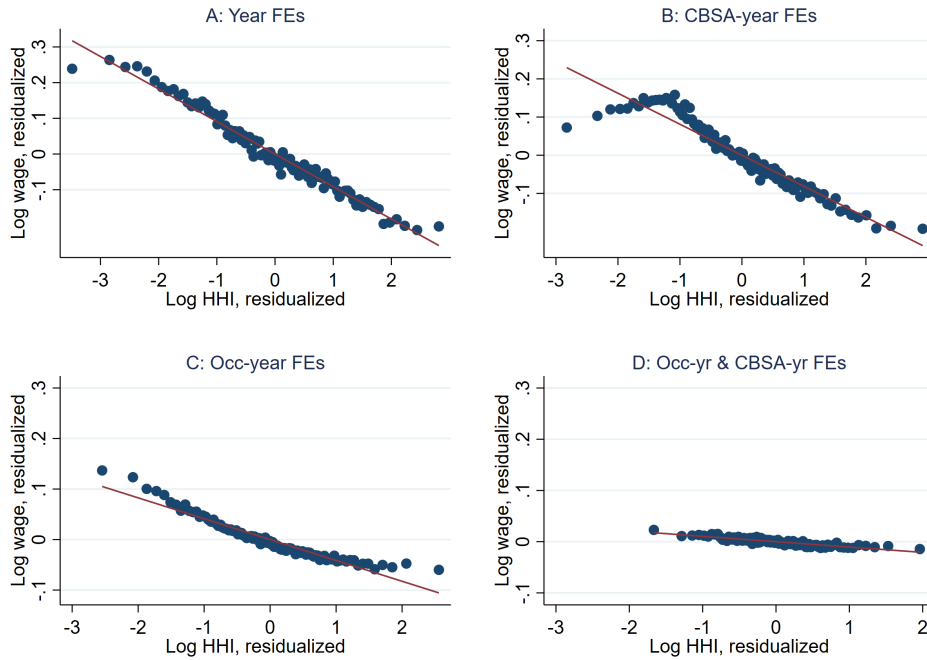
Distribution of the “occupation leave share” – the probability that a worker will leave their occupation conditional on leaving their job – by occupation. Occupation leave share is calculated from BGT resume data for 2002-2015 period. Histogram shows 786 occupations, with dashed line indicating the sample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

Figure 3: Occupational transitions and occupational characteristic similarity



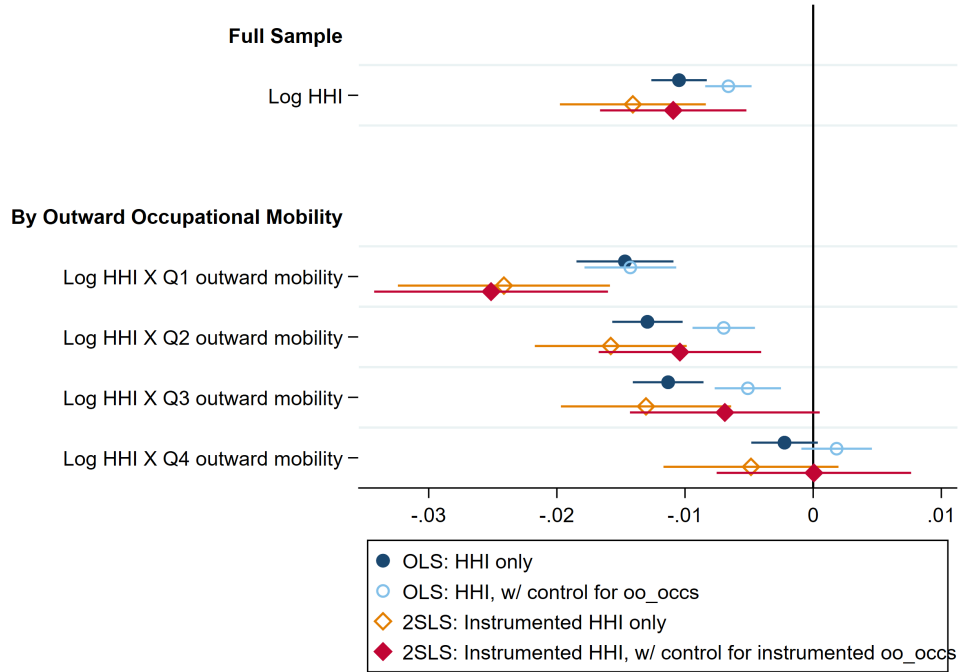
Note: This plot shows coefficients and 95% confidence intervals from regressions of occupation transition shares $\pi_{o \rightarrow p}$ on occupational characteristics: $\pi_{o \rightarrow p} = \alpha_o + \beta f(X_{occ \ o \rightarrow p}) + \gamma f(\Delta w_{o \rightarrow p}) + \epsilon_{op}$, where $f(\cdot)$ represents the absolute difference in characteristic X between occupation o and p , and α_o is occupation o fixed effect. Regressions also include absolute avg. hourly wage differences (except for amenities regressions). Standard errors are clustered by origin occupation. Appendix Figure A11, shows an analogous exercise with the *relative* difference in characteristics.

Figure 4: Correlations between wage and HHI



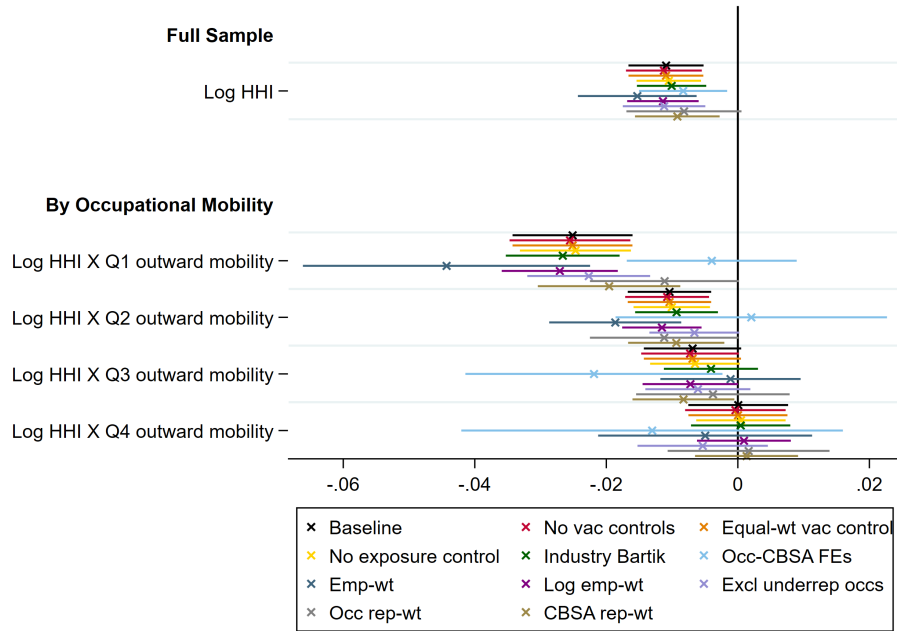
Note: Figure shows binned scatter plots of log wages and log within-occupation HHI index for occupation-CBSA cells over 2013–2016, residualized on different combinations of fixed effects. Slopes for the line of best fit on each graph are: A: -0.09, B: -0.08, C: -0.04; D: -0.01.

Figure 5: Coefficients on wage-HHI regressions



Note: Coefficients on log HHI and 95% confidence intervals from our baseline regressions of occupation-CBSA wages on employer HHI. Navy and light blue represent OLS regression coefficient of wages on HHI, with and without control for outside-occupation job options respectively. Orange and red represent 2SLS regression coefficient of wages on instrumented HHI with and without control for (instrumented) outside-occupation job options respectively. Top panel presents coefficients for the full sample (as in Table 3); bottom panels present the coefficients estimated separately by quartile of outward occupational mobility (as in Table 4). Regressions use annual data for occupation-by-CBSA labor markets over 2013-2016, and include occupation-year and CBSA-year fixed effects as well as controls described in the text. Standard errors clustered at CBSA level.

Figure 6: Coefficients on wage-HHI regressions: robustness checks



Note: Coefficients on log HHI and 95% confidence intervals from our baseline 2SLS IV regressions of occupation-CBSA wages on instrumented employer HHI, across various robustness checks (as in Appendix Tables A9-A12).

APPENDIX

A Appendix: Conceptual framework: more detail

This section expands on the conceptual framework presented briefly in section 2 in the main body of the paper.

Conceptual framework: effect of concentration on wages

Timing. Each period has two phases: the hiring phase and the production phase. During the hiring phase, workers exit or enter the labor market, employed workers bargain with their employer, and new workers are hired. During the production phase, employed workers produce at the firms they were hired at, receiving the wage determined in the hiring phase, and unemployed workers receive unemployment benefit b . At the start of the next period (in the next hiring phase), employed workers may leave their firms and/or renegotiate their wage with their current employer, and unemployed workers may search for a job. More details of each step in the process follow. Note that in this initial simple framework, we consider a clearly-defined labor market where all workers and all jobs are perfect substitutes. We relax this assumption later.

Firms. There are N firms in the labor market. Each firm j can employ up to n_j workers and has a constant returns to scale production technology with labor productivity p_j over a period. Firm size and productivity are both exogenously determined from the perspective of this model, and do not change over the period we are considering. In addition, no new firms can enter.⁷¹

Labor market exit and entry. At the start of each period, during the hiring phase, a fraction ξ of workers from each firm ‘die’ – that is, they leave their jobs *and* the labor market, for exogenous reasons (for example, family reasons, relocation, retirement, ill health, or death). These workers are replaced by an equal number of workers who are ‘born’ – i.e., they are new to the local labor market (perhaps they have moved, or newly entered or re-entered the labor force), who enter as job seekers.

Wage bargaining. Each worker who is currently employed at the start of the period Nash-bargains with her employer i over the wage.⁷² The outcome is a wage w_i , equal to the

⁷¹We choose this extremely simple, static set-up to illustrate as cleanly as possible the impact of employer concentration on wages. With the possibility for firm entry, and for firm growth – with some entry and/or adjustment costs – the general intuition of a negative impact of employer concentration on outside options and therefore on wages would persist, but the exact impact would depend on the degree to which new firms can enter, the speed at which incumbent firms can grow, and the distribution of production technologies across new and incumbent firms, and small and large firms.

⁷²The Nash bargaining outcome can be derived as the outcome of a bargaining problem where the firm

value of the worker’s outside option if bargaining breaks down and she leaves her job, oo_i , plus a share β – reflecting worker bargaining power – of the match surplus:

$$w_i = \beta(p_i - oo_i) + oo_i = \beta p_i + (1 - \beta)oo_i \quad (13)$$

If wage bargaining breaks down, the worker leaves her job and becomes a job seeker. We assume that all workers at a given firm have the same set of outside options in expectation, such that all workers at any given firm will in equilibrium receive the same wage.

Vacancies. After labor market exit and wage bargaining have happened, firms post vacancies to fill the positions which have been vacated by either worker exit or the breakdown of wage bargaining. There is no cost to post vacancies to fill existing positions, but firms cannot post vacancies for new positions (i.e. there is no firm growth).⁷³ A firm can only post one vacancy for each position they wish to fill. Vacancies are posted as take-it-or-leave-it wage offers, with the posted wage equal to the wage the firm is paying to its other workers (which was decided on in the wage bargaining process described above). This constraint – that similar workers are paid similar amounts – is often observed in practice in firms. It may be motivated by fairness concerns or explicit internal pay hierarchies or bargaining agreements.⁷⁴

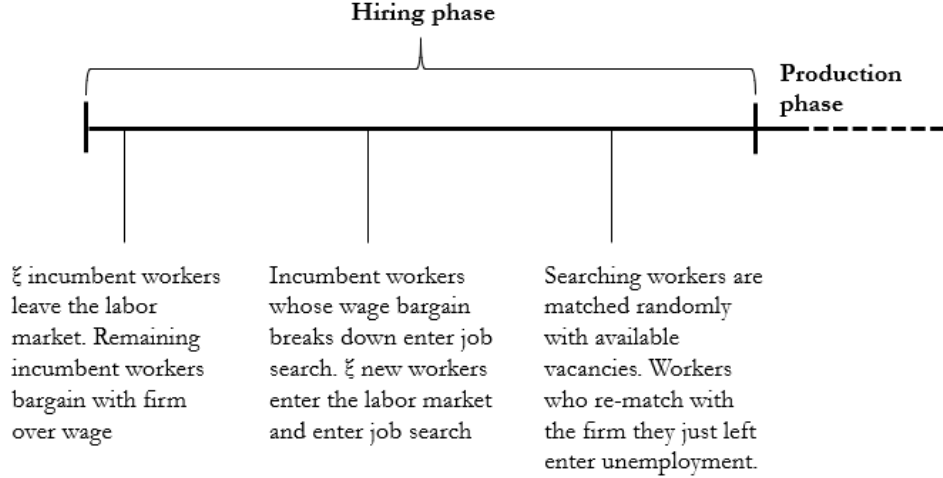
Job seekers. Job seekers are comprised of workers who have newly entered the labor market, workers who have left their previous job because their match broke down at the start of the period, and workers who were unemployed in the previous period. Job search can only take place at the start of the period, during the hiring phase. There is no on-the-job-search: only workers who have left their current employer may search for a new job. Each job seeker applies to all local employers $j \in N$, but search and matching frictions mean that each job seeker will receive exactly one job offer during each hiring phase. Specifically, a ball-urn matching procedure randomly matches pairs of job seekers and firms (in a procedure similar

and worker both wish to maximize their joint surplus from the match, where the surplus generated is the difference between the product of the match p_i and the worker’s outside option oo_i . The specific bargaining problem which generates the Nash outcome in our framework is one where the wage satisfies $w_i = \argmax_w (w_i - oo_i)^\beta (p_i - w_i)^{(1-\beta)}$, as shown in Jaeger, Schoefer, Young and Zweimueller (forthcoming) or Manning (2011). (This is a particularly simple formula which arises in part from the assumption that the firm’s outside option is zero).

⁷³This assumption is not necessary for the qualitative direction of our conclusions but keeps things simple.

⁷⁴There is a large literature on the role of internal vs. external factors in determining the wages of new hires as compared to incumbent workers. For some examples: Bewley (1999), in interviews with firms in New England, assembles a range of evidence that conceptions of internal fairness and equity are extremely important in wage setting. Galuscak, Keeney, Nicolitsas, Smets, Strzelecki and Vodopivec (2012) use firm-level data in 15 EU countries and find strong evidence that firms do not wish to differentiate between the wages of newly hired workers and similarly qualified incumbents even if external labor market conditions change. Our assumption that new hires receive a ‘take-it-or-leave-it’ wage offer conforms to the survey evidence of Hall and Krueger (2012), who find in a survey of 1,300 US workers that two thirds of workers considered their offers to be ‘take-it-or-leave-it’ and did not bargain over the wage.

to the framework developed by Jarosch, Nimczik and Sorkin (2019)). For a given job seeker i , define the probability of receiving an offer from each firm j as α_{ij} . The probability of receiving no offers is therefore $1 - \sum_j \alpha_{ij}$. If job seekers do not accept the job offer they receive, or if they receive no job offers, they remain unemployed for the period, receiving unemployment benefit b . We assume b is strictly lower than the wages offered by any feasible employer, such that a job seeker who receives an offer always accepts it. In our model this condition will always be satisfied as long as the productivity of all jobs is greater than the unemployment benefit: this is because the wage offered to each worker by firm j is equivalent to the wage bargained by the existing workers at firm j , which itself is strictly greater than unemployment benefit b .⁷⁵ The timing of the wage bargain and job search process is illustrated in the figure below.



Outside option value for employed workers. The wage at each firm is determined by the bargain with employed workers, which in turn depends partly on the value of the outside option for these employed workers. What is this outside option value? The outside option for an employed worker bargaining with her employer is to leave her current job and become a job seeker. She does not know with certainty what her outcome will be as she will be matched with at most one feasible job if she leaves her current job. Her expected wage if she leaves her job is therefore a weighted average of the wages paid by each firm

⁷⁵The analysis above assumes that workers care only about money. If we consider instead utility (which may include a value of leisure or work), the productivity of all jobs must be greater than the money-equivalent utility value of unemployment (including the unemployment benefit and any utility or disutility of unemployment relative to work) for all workers.

j , w_j , weighted by the probability of being matched with each firm j , α_j , as well as the unemployment benefit b multiplied by the probability of receiving no job offers $1 - \sum_{j \neq i}^N \alpha_j$:

$$oo_i = \sum_{j \neq i}^N \alpha_j \cdot w_j + \left(1 - \sum_{j \neq i}^N \alpha_j\right) \cdot b \quad (14)$$

We assume that each firm-worker match has weakly positive surplus, such that the bargained wage is always weakly greater than the outside option value. This means that, in equilibrium, no bargaining session will break down.⁷⁶

Equilibrium wage. The bargained wage for workers at firm i (which is also paid to new hires at firm i) satisfies:

$$\begin{aligned} w_i &= \beta p_i + (1 - \beta) oo_i \\ &= \beta p_i + (1 - \beta) \left(\sum_j \alpha_j \cdot w_j + \left(1 - \sum_j \alpha_j\right) \cdot b \right) \end{aligned} \quad (15)$$

What do the probabilities of being matched with each feasible firm, α_j , correspond to? Since each job seeking worker is randomly matched with one vacancy, and there are equal numbers of vacancies and job searchers, the probability of the offer a job seeker receives being from a particular firm j is proportional to the share of vacancies posted by that firm j , as a share of all vacancies in the labor market, σ_j . This is in the same spirit as Burdett and Mortensen (1980) and Jarosch et al. (2019) who use the employment share to proxy for the likelihood of getting a job offer from a given firm.

In a labor market with infinitesimally small firms, this procedure would lead to every job seeker receiving a match from a feasible employer each period. However, in a labor market with some large employers, a worker who left firm j at the start of the period has some non-zero chance of being re-matched with firm j in the job search process (in fact, her chance of being re-matched with her former employer is σ_j). Knowing this, during the wage bargain process a large employer can threaten *not* to re-employ their worker if that worker leaves the firm but is re-matched with them in the job search process. We assume that this threat is credible and is exercised: if bargaining breaks down between a given worker and her employer j , and if the random job search process re-matches that worker to her former employer j , employer j will refuse to re-hire her (as in Jarosch et al., 2019) and she will instead have to move to unemployment for the period.⁷⁷ These assumptions imply the following wage

⁷⁶Note that we assume complete information about the outside option for both the worker and the firm.

⁷⁷In a single-period game, this strategy would not be time-consistent for the employer, as they are unable to fill the vacancy for the period and so lose the potential for positive surplus from re-hiring the worker. However, in a multiple period game with repeated interactions between firms and workers, and/or firm reputations, it can be in the firm's interest to avoid re-hiring a worker who has quit the firm in order to

equation for a worker at firm i :

$$w_i = \beta p_i + (1 - \beta) \left(\sum_{j \neq i} \sigma_j \cdot w_j + \sigma_i \cdot b \right) \quad (16)$$

Note that the outcome of the wage bargain reached by workers at firm i depends on the outcome of the wage bargain reached by workers at all other local firms j , but the wage bargained by workers at firms j depends on the wage bargained by workers at firm i . To solve this “reflection problem”, we iteratively substitute for w_j . The wage expression becomes

$$\begin{aligned} w_i = & \beta p_i + \beta(1 - \beta) \sum_{j \neq i} \sigma_j p_j + \beta(1 - \beta)^2 \sum_{j \neq i} \sigma_j \sum_{k \neq j} \sigma_k p_k + \dots \\ & + (1 - \beta) \sigma_i b + (1 - \beta)^2 \sum_{j \neq i} \sigma_j^2 b + (1 - \beta)^3 \sum_{j \neq i} \sigma_j \sum_{k \neq j} \sigma_k^2 b + \dots \end{aligned} \quad (17)$$

expanded out to the third order, where the ellipses (...) signify further expansions to higher orders.

Since we are interested in the *average* wage in the labor market, we then take the average wage across all firms in the labor market, $\bar{w} = \sum_i \sigma_i w_i$. This gives us the expression:

$$\begin{aligned} \bar{w} = & \beta \left(\sum_i \sigma_i p_i + (1 - \beta) \sum_i \sigma_i \sum_{j \neq i} \sigma_j p_j + (1 - \beta)^2 \sum_i \sigma_i \sum_{j \neq i} \sigma_j \sum_{k \neq j} \sigma_k p_k + \dots \right) \\ & + (1 - \beta) b \left(\sum_i \sigma_i^2 + (1 - \beta) \sum_i \sigma_i \sum_{j \neq i} \sigma_j^2 + (1 - \beta)^2 \sum_i \sigma_i \sum_{j \neq i} \sigma_j \sum_{k \neq j} \sigma_k^2 + \dots \right) \end{aligned} \quad (18)$$

To simplify this expression, define $\bar{p} = \sum_i \sigma_i p_i$ as the average productivity across firms, and denote $\tilde{p}_j = p_j - \bar{p}$ as the difference between firm j ’s productivity and the market average. In addition, define the r th order concentration index Ω_r as the concentration index with r “steps” as in the expression

$$\Omega_r = \underbrace{\sum_i \sigma_i \sum_{j \neq i} \sigma_j \sum_{k \neq j} \sigma_k \dots \sum_{m \neq n} \sigma_m \sum_{p \neq m} \sigma_p^2}_{\text{with } r \text{ summation terms or “steps” in the expression}} \quad (19)$$

such that the first order concentration index Ω_1 is the sum of the squared employer shares (the HHI: $\Omega_1 = \sum_i \sigma_i^2$), the second order concentration index is $\Omega_2 = \sum_i \sigma_i \sum_{j \neq i} \sigma_j^2$, and so

maintain the firm’s reputation in future bargaining rounds with this or other workers - even if this comes at the cost of an unfilled vacancy in the current period. A less stylized setting producing a similar outcome could be a setting with on-the-job search, where workers can receive job offers from firms other than their own, and when they do so, they can use these as outside options to bargain for a higher wage. In this setting, workers at the largest firms will receive the fewest outside offers.

on. Also define $\Omega_0 = 0$. We can then rewrite the average wage equation (18) as

$$\begin{aligned} \bar{w} = & \beta \bar{p} \left(1 + \sum_{n=1}^{\infty} (1-\beta)^n \left(1 - \sum_{r=1}^n \Omega_r \right) \right) + (1-\beta)b \left(\sum_{n=1}^{\infty} (1-\beta)^{n-1} \Omega_n \right) \\ & - \beta \left((1-\beta) \sum_i \sigma_i^2 \tilde{p}_i + (1-\beta)^2 \sum_i \sigma_i \sum_{j \neq i} \sigma_j^2 \tilde{p}_j + (1-\beta)^3 \sum_i \sigma_i \sum_{j \neq i} \sigma_j \sum_{k \neq j} \sigma_k^2 \tilde{p}_k + \dots \right) \end{aligned} \quad (20)$$

That is, the average wage in a given labor market is a function of three terms: the average productivity in the labor market \bar{p} , multiplied by a function of worker bargaining power and employer concentration; the employment benefit b , multiplied by a function of worker bargaining power and employer concentration; and a third term which reflects the interaction between employer share σ_j and employer relative productivity \tilde{p}_j .

The first and second terms of our wage expression illustrate that the wage declines as average employer concentration increases: as different aspects of average employer concentration increase, Ω_r increases, meaning that less weight in the wage equation falls on average productivity (\bar{p}) and more weight falls on the unemployment benefit (b). This in turn is because large employers can credibly threaten not to re-hire workers who quit, reducing these workers' bargaining power by making their outside option worse (i.e. making it more likely that they will enter unemployment if they quit their firm). So, higher employer concentration suppresses wages by worsening workers' outside option in the wage bargain.

Note also that the relationship between employer shares σ_j and firm productivity p_j factors into the wage in both the first and third terms. Since average productivity \bar{p} is determined by the productivity of each employer, and the share of that employer in the labor market, the first term illustrates that average productivity will be higher and therefore the wage will be higher if the high productivity firms are also the large firms. However, the third term mitigates this effect somewhat: it reflects the fact that, if it is the largest firms which are the most productive, the passthrough of average productivity to average wages via the outside option channel will be lower than if it is the smallest firms which are the most productive, because the largest firms are in fewer workers' outside option set.

To obtain the average wage expression we cite in the main body of the paper, we take a second order approximation of 20 in the employer shares σ_j (i.e. removing all terms with σ_j^n where $n > 2$). This reduces the expression to become a function of the squares of the employer shares – and, therefore, the commonly-used Herfindahl Hirschmann Index or HHI:

$$\bar{w} = (1 - (1-\beta)HHI) \bar{p} + (1-\beta)HHIb - \beta(1-\beta) \sum_i \sigma_i^2 \tilde{p}_i \quad (21)$$

That is, this expression suggests that higher employer concentration increases the weighting

on the unemployment benefit b , and decreases the weighting on the productivity of labor p , relative to a world with no employer concentration. If all firms have the same productivity $p_i = \bar{p} \forall i$, or if the correlation of market shares and firm productivity is small, the last term is close to zero and the wage is simply a concentration- and bargaining power-weighted average of productivity p and unemployment benefit b .⁷⁸

Note, then, that one might mean two different things when one asks “what is the effect of employer concentration on wages?”. First, one might be asking “what is the effect of this labor market having become more concentrated, relative to the past?”. On the one hand, rising concentration exerts downward pressure on wages by worsening worker outside options. On the other hand, rising concentration may well exert upward pressure on wages by increasing average productivity:

On the other hand, one might be asking “what is the effect of the employer concentration in this labor market, in terms of suppressing workers’ wages below their productivity?”. In that case, the answer to this question takes the degree of productivity in the labor market *as given* and looks only at the effect of employer concentration in reducing wages below that level of productivity by worsening outside options. In a sense, this is trying to *isolate* the effect of employer concentration on outside options from its potential effect on productivity. This *latter* question is the one we are focusing on in this paper.

Incorporating outside-occupation options

The conceptual framework explained above, however, assumes that all jobs and workers are perfectly substitutable in a clearly delineated labor market. As discussed in this paper, this is rarely the case in practice. Workers can switch between jobs in different occupations and locations, but differentially so for different options. In this paper, we focus on occupations within a given metro area. One can easily extend our framework to incorporate also the option to move metro area.⁷⁹

⁷⁸Note: in our framework, for simplicity the *only* way the worker can end up unemployed is if they do not get matched with a firm, and the only way they do not get matched with a firm is if their ‘match’ in the random matching process is their current employer, who refuses to re-hire them. This means that in an unconcentrated labor market ($HHI = 0$) in this model, there would be no unemployment and the average wage would equal the average product $\bar{w} = \bar{p}$. On the other hand in a labor market with only one firm ($HHI = 1$) the wage would be $\beta p + (1 - \beta)b$, the Nash bargaining formula when the only outside option for workers is unemployment. In our framework, therefore, the only reason there is a markdown of the wage from the marginal product is because there is employer concentration, which results in some probability of becoming unemployed if bargaining breaks down with the worker’s own firm. This simplification is for clarity of exposition only; one could extend this framework to incorporate steady state unemployment even in an unconcentrated labor market, but where higher employer concentration increases the probability of unemployment if a match breaks down.

⁷⁹While the outside option to move location certainly matters, the data suggests that for most workers this will be less important than the outside option to get a job in another occupation in workers’ current

Ideally, we would be able to delineate which firms in outside occupations are in a worker's feasible labor market for any given worker, and estimate the probability that a worker would receive a job in that outside occupation. In practice, we cannot. To work with commonly-available data definitions and publicly-available data, we must instead work with the occupational definitions from the SOC 6-digit classification scheme. We therefore extend our framework above by defining the primary labor market for workers as their local occupation o and incorporating an outside option term reflecting the value of moving to jobs in other occupations. Refer back to our equation (16) for the bargained wage in firm i ,

$$w_i = \beta p_i + (1 - \beta) \left(\sum_{j \neq i} \sigma_j \cdot w_j + \sigma_i \cdot b \right)$$

but now edit this equation to reflect that some feasible jobs are in workers' own occupation o whereas some are in other occupations p (Where employer share $\sigma_{j,o}$ now refers to firm j 's share of vacancies *within* occupation o):

$$w_{i,o} = \beta p_{i,o} + (1 - \beta) \left(\underbrace{\text{own occupation options } oo^{own}}_{\text{own occupation options } oo^{own}} + \underbrace{\sum_{p \neq o}^{N_{occs}} \text{outside occupation options } oo^{occs}}_{\text{outside occupation options } oo^{occs}} + \underbrace{\text{unemployment}}_{\text{unemployment}} \right) \quad (22)$$

This expression states that the bargained wage for workers in firm i in occupation o is a function of their productivity at firm i $p_{i,o}$, the outside option value of moving to other firms j in their own occupation o , the outside option value of moving to other firms l in other occupations p , and the outside option value of moving to unemployment and receiving benefit b . Note that the expression for the probabilities of a worker from firm i in occupation o matching with firm l in occupation p if she becomes a job seeker is now the product of firm l 's vacancy share in occupation p , $\sigma_{l,p}$, as well as the probability of the worker getting her job offer from some firm in occupation p conditional on having left her job in occupation o , $Prob(o \rightarrow p)$.⁸⁰

city. Occupational mobility is substantially higher than geographic mobility: only around 3% of U.S. workers move between metropolitan areas each year (according to IRS county-to-county migration data), and over 80% of job applications by workers are sent to jobs within their metropolitan area Marinescu and Rathelot (2018). In addition, geographic mobility has declined over time (Molloy, Smith and Wozniak, 2011), and the proliferation of state-level occupational licensing has made geographic mobility more difficult for many workers (Johnson and Kleiner, 2020). There are, however, a subset of more mobile workers – like highly educated professionals – for whom the failure to consider job options outside their own city may be more problematic.

⁸⁰Note also that this expression assumes implicitly that each firm only employs workers of one occupation (or, alternatively, that the worker's own firm i can only refuse to re-employ that worker if she is re-matched

To take this expression to the data, we note that if employers' vacancy shares are relatively similar to their current employment shares, $\sum_l^{N_p} \sigma_{l,p} \cdot w_{l,p}$ can be approximated simply the average wage in local occupation p . This eliminates the need for us to consider the reflection problem that wages in occupation o affect wages in occupation p and vice versa - instead, we can use data on the average wage in each local occupation p to control directly for the effect of wages in occupation p on occupation o (as we do in our empirical implementation). This gives us the wage expression

$$w_{i,o} = \beta p_{i,o} + (1 - \beta) \left(Prob(o \rightarrow o) \sum_{j \neq i} \sigma_{j,o} w_{j,o} + \sum_{p \neq o}^{N_{occs}} Prob(o \rightarrow p) \bar{w}_p + Prob(o \rightarrow o) \sigma_i b \right) \quad (23)$$

We define $(Prob(o \rightarrow o) \sum_{j \neq i} \sigma_{j,o} w_{j,o})$ as our outside-occupation option index oo_o^{occs} . As in the simple framework, we iteratively substitute for wages in other local firms j in the same occupation o , rearrange, take the average wage in the local occupation, and write the resulting wage expression in terms of our higher order concentration indices $\Omega_{r,o}$ (where subscript o denotes that this is the concentration index for employers in local occupation o). This gives us an expression for the average wage in occupation o :

$$\begin{aligned} \bar{w}_o = & (\beta \bar{p}_o + (1 - \beta) oo_o^{occs}) \left(1 + \sum_{n=1}^{\infty} (1 - \beta)^n Prob(o \rightarrow o)^n \left(1 - \sum_{r=1}^n \Omega_{r,o} \right) \right) \\ & + b \left(\sum_{n=1}^{\infty} (1 - \beta)^n Prob(o \rightarrow o)^n \Omega_{n,o} \right) \\ & - \beta \left((1 - \beta) Prob(o \rightarrow o) \sum_i \sigma_i^2 \tilde{p}_{i,o} + (1 - \beta)^2 Prob(o \rightarrow o)^2 \sum_i \sigma_i \sum_{j \neq i} \sigma_j^2 \tilde{p}_{j,o} + \dots \right) \end{aligned} \quad (24)$$

Once again taking a second order approximation in employer shares – that is, only considering concentration index $\Omega_{1,o} = \sum_i \sigma_{i,o}^2 = HHI_o$ – we can simplify the expression for the average wage in occupation o to become

$$\begin{aligned} \bar{w}_o = & (1 - (1 - \beta) Prob(o \rightarrow o) HHI_o) (\alpha \bar{p}_o + (1 - \alpha) oo_o^{occs}) + (1 - \beta) Prob(o \rightarrow o) HHI \cdot b \\ & - \beta (1 - \beta) Prob(o \rightarrow o) \sum_i \sigma_i^2 \tilde{p}_{i,o} \end{aligned} \quad (25)$$

where $\alpha = \frac{\beta}{1 - Prob(o \rightarrow o)(1 - \beta)}$.

Ignoring the final term, which is very small if the average productivity of individual firms is not strongly correlated with their vacancy shares, the average wage in occupation o is a

with a job in her initial occupation o , but not if she is re-matched with firm i with a job in a new occupation p .

weighted average of the average productivity in occupation o , \bar{p}_o , the value of jobs outside occupation o , oo_o^{occs} , and unemployment benefit b . The weights are a function of worker bargaining power β , the probability that workers are matched with another firm in their own occupation if they leave their job $Prob(o \rightarrow o)$, and employer concentration HHI_o .

As before, employer concentration within workers’ own occupation increases the relative likelihood of workers ending up unemployed if they quit their job, increasing the weighting on b the unemployment benefit in the wage bargain and reducing the weighting on other jobs (productivity \bar{p}_o and the value of moving to outside options outside workers’ own occupation oo_o^{occs}). In addition, though, there is now an interaction with $Prob(o \rightarrow o)$, the likelihood of the worker staying in her occupation if she leaves her job. The more likely she is to stay in her own occupation if she leaves her job – i.e., the less likely she is to be able to find a job in a different occupation – the more employer concentration in her own occupation matters for her wage. Finally, as before, note that there is an interaction with worker bargaining power β . The more bargaining power a worker has over the match surplus, the less the outside option matters in the wage bargain and therefore the less employer concentration matters for the wage.⁸¹

B Appendix: Burning Glass Technologies Vacancy Posting Data

This section contains further information about the vacancy posting data set from Burning Glass Technologies (“BGT”), which we use to construct our employer concentration index (as discussed briefly in Section 4). (We also use a different data set from BGT – the resume data set – to construct our measures of occupational mobility. We discuss the BGT resume data set in more detail in Appendix section C.)

Burning Glass Technologies is an analytics software company that provides real-time data on job growth, skills in demand, and labor market trends. They frequently collaborate with academic researchers by providing data. The BGT vacancy data on online job postings has been used in several other academic papers, including Azar, Marinescu, Steinbaum and Taska (2020) and Hazell and Taska (2019).

Vacancy posting data overview

Burning Glass Technologies constructs its vacancy database by collecting online job postings from about 40,000 websites, capturing the near-universe of online US job vacancies. They only measure *new* vacancy postings, so do not re-capture a given vacancy if it is left open

⁸¹Note that if there is no possibility of finding a job in another occupation (i.e. $Prob(o \rightarrow o) = 1$, and $oo_o^{occs} = 0$), this expression becomes identical to the expression for the average wage in the first part of this section, equation (21).

for a long time.⁸² They use proprietary algorithms to de-duplicate vacancies (for example if the same vacancy is posted on different websites).

We use BGT’s vacancy data for the years 2013–2016. Over this four year period, we have data on 74 million vacancies which had been assigned a SOC 6-digit occupation and metropolitan area by BGT. Of these, about one third or 24.8 million had no information about the employer, while 49.2 contained employer names (with a total of 1.02 million different employers).

Defining the employer and calculating the HHI

A key aspect for our purposes is how an “employer” is defined in the data. BGT’s algorithm attempts to group together name variants for employers into a standard set, counting for example “Lowe’s” or “Lowes” as the same employer. However, there may be some instances where employers which are in reality the same have not been detected by the algorithm due to large differences in spelling, punctuation, or naming conventions. We therefore carry out an additional layer of grouping by removing punctuation, spacing, and capitalization, and adjusting for common spelling differences or acronyms. We also used the Agency for Healthcare Research Quality’s “Compendium of US Health Systems” database for 2016 to link hospitals to the health systems which own them where possible, treating a health system as a single employer rather than a specific hospital. This match was not always perfect: there are several cases where we have not necessarily succeeded in matching all hospitals to their owner, because of the presence of multiple hospitals in our database with the same name. We also manually scanned several thousand of the largest employers in the database to group together different employer names which were evidently part of the same ultimate employer.

This means that we for the most part treat vacancies as offered by the same employer if the *name* listed by the employer on the vacancy is sufficiently similar, or if there is a well-known or easily-identifiable relationship between a parent and subsidiary company with different names (such as “Alphabet” and “Google”, or two hospitals which are part of the same health system).

We do not capture relationships where one company owns another company but the names are not similar enough to identify this easily: this means that in some cases we will understate employer concentration by attributing vacancies to different employers. On the other hand, our employer categorization means that individual establishments of an employer – or even franchises of a brand – will be treated as the same employer, which may

⁸²To capture vacancies which firms keep online to hire workers continually for a given job, BGT consider a vacancy to be “new” if the identical vacancy is still online after 60 days (Carnevale, Jayasundera and Repnikov, 2014).

overstate employer concentration if pay decisions are made at the level of the establishment or franchise rather than the overall firm or brand group. It is not entirely conceptually clear whether employer concentration should be measured at the level of the establishment or the firm. On the one hand, individual establishments often have independent hiring policies; but on the other hand, multi-establishment firms often have common internal pay scales meaning they effectively operate as one employer across establishments. Similarly, it is not entirely conceptually clear whether franchises of the same brand should be considered as separate employers. On the one hand, they are independent businesses; on the other hand, franchisees’ human resources policies are often at least partly dictated by the franchisor (Weil, 2014), and there have been a number of prominent cases where franchisors have required franchisees not to ‘poach’ each others’ employees (with (Krueger and Ashenfelter, 2018) estimating that over half of major franchisors have no-poaching agreements in their franchise contract). We view the question of the appropriate *level* at which to calculate employer concentration – taking into account ownership structures across firms, as well as establishment structures within firms – as a fruitful avenue for further research.

How do we treat the one third of vacancies which do not include an employer name? When we calculate our HHI statistics for each occupation-metropolitan area-year cell we assume that each vacancy listing by an employer with no name information in the database is a *separate employer* (as do Azar et al. (2020)). This will lead us to mechanically underestimate the HHI, as it is likely that at least some of these different vacancy postings where no name information is available come from the same employer in practice (Azar et al. (2020) note that the vacancy postings without employer name information are often due to staffing companies not disclosing on whose behalf they are posting a given job).

Summary statistics

Here, we provide summary statistics for the 49.2 million vacancies which contain employer names. As one might expect given the skewed distribution of employment, the large majority of these vacancies are accounted for by a small group of large employers: 841 employers each posted more than 10,000 vacancies online over 2013–2016, and these 841 employers are responsible for a total of 32.4 million vacancies. While many of the small employers in our data are only present in the data for a subset of the four years 2013–2016, a subset of large employers are present for all four years (as shown in Table A1): as a result more than 75% of all vacancies in our database are listed by employers which are present in all four years of the sample. If employers hire a lot in any one year, they also tend to hire a lot in other years: the correlation of local occupation-specific vacancies by employer from one year to

the next is 0.77.

Vacancy postings, job vacancies, and employment

A natural question is how our data on vacancy postings relates to total job vacancies and to total employment. In theory, when calculating an HHI of employer concentration, one would either like to use data on the share of job vacancies or the share of employment accounted for by each employer. Instead, we have the share of job *postings* accounted for by each employer at the level of each SOC 6-digit occupation, metropolitan area, and year.

BGT estimates that its vacancy data covers the near-universe of online job postings. The Bureau of Labor Statistics' JOLTS database (Job Openings and Labor Turnover Survey) collects data on job *openings*, where each opening represents a specific position that the firm is actively recruiting to fill. The conceptual difference between a job posting and a job opening is that one job posting (a job advertisement) could be used to fill multiple job openings, if the firm needs to hire several people for a job with the same title, job description, and location at the same time. This may be a particular concern when measuring employer concentration, as a large employer may hire more workers per job posting than a small employer, and so we would systematically underestimate concentration in labor markets with a highly skewed distribution of employer size, relative to labor markets with more symmetric distributions of employer size. For example, when hiring for warehouse laborers, a large warehousing company like Amazon might hire several workers under a job ad for a "Warehouse Associate".⁸³ On the other hand, for occupations where there is a high degree of granularity of individual job titles and job requirements within an occupation, we may be more likely to observe a one-to-one mapping between job *postings* and job *openings*. One might expect, therefore, that our measures of employer concentration will be less reliable for occupations for which there are many large employers who hire a lot of workers who are not required to be much differentiated in their job tasks, job titles, and qualifications or skills. If an occupation has a particularly low ratio of job postings to job openings, one would expect it to be underrepresented in our data relative to its employment in the general workforce: As discussed later in the 'representativeness' section, our data appears to be underrepresentative particularly for certain large low-wage occupations like laborers, cashiers, and food serving and preparation workers, for whom this might be a particularly common phenomenon. Ideally, we would be able to calculate employer concentration at the level of true job openings/vacancies, or employment, rather than vacancy postings, but we are

⁸³In the extreme case, where each firm only posts one vacancy per occupation that it is hiring for, our measure of the HHI will actually be a measure of $1/N$ where N is the number of firms hiring for that occupation in that local area.

not aware of a data set that enables us to observe firm-level local occupational employment or vacancies in the US.

Representativeness

To what extent is the online job *posting* data representative of all job *openings*? Carnevale et al. (2014) estimated as of 2014 that between 60 to 70 percent of all job openings could be found in the BGT online vacancy posting data. They do this by comparing the number of new job postings (as measured by BGT) to the number of active job openings as measured by the JOLTS database (inflating the BGT job postings number by the new jobs to active jobs ratio in the Help Wanted Online database to take account of the fact that BGT only captures new postings while JOLTS captures all active job postings). Azar et al. (2020), using the same methodology, estimate that the share of job openings online as captured by BGT is roughly 85% of total job openings as measured by the JOLTS database in 2016, and the jobs that are not online are usually offered by small businesses and union hiring halls.

The BGT vacancy data has been used in several other academic papers in recent years, which have carried out detailed analyses of its representativeness. We provide a brief summary of the representativeness of the BGT vacancy data here and refer the interested reader to Carnevale et al. (2014), Hershbein and Kahn (2018), and Azar et al. (2020) for more details. Note in particular that Azar et al. (2020) use the BGT vacancy data for the same purposes as we do: to calculate employer HHI concentration indices at the level of local SOC 6-digit occupations.

Hershbein and Kahn (2018) compare the distribution of BGT vacancies across major industry groups to the distribution of job vacancies in the Bureau of Labor Statistics' JOLTS database. While BGT is overrepresented in health care and social assistance, finance and insurance, and education, and underrepresented in accommodation and food services, public administration/government, and construction, the differences are mostly small in magnitude. Hershbein and Kahn (2018) also compare the distribution of BGT vacancies by occupation to both the stock and flow of employment in the United States, showing that BGT vacancy data has a much larger than average representation of computer and mathematical occupations, management, healthcare, and business and financial operations, and lower representation in transportation, food preparation and serving, production, and construction. This degree of representativeness does not change much over time in the BGT sample.

To analyze representativeness by occupation systematically, we calculate a measure we call 'represented-ness': the share of all vacancies in our data represented by each SOC 6-digit occupation, divided by the share of all employment in the BLS occupational employment

statistics which is represented by each SOC 6-digit occupation. Note that our ‘representedness’ measure captures three dimensions: one is the degree to which the BGT vacancy *posting* data is representative of the totality of vacancy postings in the US, one is the degree to which vacancy *postings* are representative of true vacancies (job openings), and one is the degree to which individual occupations have high or low turnover (and as a result, a high or low ratio of vacancies to employment). We are interested primarily in the first two of these three, and would ideally compare the representativeness of our BGT vacancy data to a data set of the universe of online *and* offline vacancies by occupation, but this is not available. We show a scatter plot of the share of vacancies each occupation accounts for in our data, relative to the share of employment that occupation accounts for in the BLS OES, in Appendix Figure A1.

Of the largest occupations in the data, sales occupations are relatively equally represented in BGT data as compared to the BLS OES; registered nurses, truck drivers, and computer and software occupations are overrepresented, while laborers, cashiers, waiters, janitors, personal care aides, and food preparation and serving workers are substantially underrepresented in the BGT vacancy data. This pattern of underrepresentativeness may not be surprising. These underrepresented occupations are all occupations which tend to have a higher share of their employment accounted for by self-employment, households, or small employers, who may be more likely to advertise through local advertisement channels (posted, for example, on physical job boards, or hired through local agents) or through networks, referrals, or word-of-mouth. In addition, some of these underrepresented occupations may be more likely to have a high ratio of job openings to job postings (a high number of workers hired per job posting).

Similarly, zooming in on the next tier of occupations by size, we see overrepresentation of financial, information, management, and healthcare occupations, relatively even representation of sales, delivery, and mechanical occupations, and underrepresentation of workers in occupations with a large share of self-employment (construction, plumbing, landscaping), employment by individual households (maids and housekeeping cleaners, home health aides), or employment where firms may run single job ads for many workers, or which may advertise informally (dishwashers, cooks, food preparation workers, receptionists).

For our purposes, we have two potential representativeness concerns. One concern might be that the representativeness of our data is correlated in some way with factors which would affect both employer concentration and the wage. This concern is only relevant for the *estimated effect of concentration in our regressions* if our database systematically underrepresents low-wage occupation-city labor markets even when controlling for occupation and fixed effects: that is, that within a given occupation, the lower-wage cities are under-

represented and within a given city, the lower-wage occupations are underrepresented. For our normative conclusions in terms of estimating the *aggregate number* of workers who are affected by employer concentration, and creating a ranking of which occupations are more or less affected, underrepresentativeness of the data is more of a concern: if some occupations are underrepresented in the BGT resume data, they may appear more concentrated when in fact, it is simply the case that online vacancy postings reflect fewer of the true vacancies available in the labor market for that occupation. As such, we take care when drawing these conclusions not to isolate specific occupations which appear to be severely underrepresented in our data.

C Appendix: Burning Glass Technologies Resume Data

The Burning Glass Technologies resume data set is a new proprietary data set of 16 million unique US resumes spanning years over 2002–2018. Resumes are sourced by BGT from a variety of BGT partners, including recruitment and staffing agencies, workforce agencies, and job boards. Using the raw resumes, BGT populates a database which contains observations for each individual, denoting their education, jobs, and years in which they worked in each job. BGT’s proprietary occupation parser assigns SOC 6-digit occupation codes to each job title listed on each resume. With this data set, we are able to observe 16 million unique workers’ job histories and education up until the point where they submit their resume, effectively making it a longitudinal data set (spanning different segments of the 2002–2018 period for different workers). In this paper, we use the resume data to construct occupational transition matrices between SOC 6-digit occupations at a highly granular level. We describe the data set and our methods further below.

Construction of occupation transition matrices

Before calculating occupation transition matrices, we apply a number of filters to the raw BGT data:

- Reduce the number of mis-parsed job or resume observations in our data set: eliminate all jobs listed as having lasted more than 70 years, and eliminate any resumes submitted by workers whose imputed age is less than 16 or greater than 100.⁸⁴
- Eliminate all jobs held before 2001.
- Eliminate all resumes with non-US addresses.

⁸⁴See the next subsection for more details on how we impute ages to the resumes.

- Eliminate any jobs which are listed as having lasted less than 6 months, to ensure that we are only capturing actual jobs rather than short-term internships, workshops etc.

The final number of resumes that contain at least two sequential years of job data under these restrictions is 15.8 million.

From each of these resumes, we extract a separate observation for each job a worker was observed in, in each year they were observed in that job. (We define a ‘job’ as a unique job title-employer-occupation combination, meaning that a worker can in theory switch job but remain at the same employer and/or in the same occupation.) For each job, we retain information on the SOC 6-digit occupation code. This gives us a data set of 80.2 million worker-job-occupation-year observations, where each worker might be observed in multiple jobs in the same year (either if jobs were held concurrently or the worker switched from one job to another within a given year).

To identify occupational transitions from year to year, we match all sequential pairs of worker-job-occupation-year observations. For instance, if a worker had a job as a Purchasing Manager in the period 2003-2005, and a job as a Compliance Officer in 2005-2007, we would record sequential occupation patterns of the form shown in the table below.

Illustrative example of sequential job holding data.

Year:	2004	2005	2006
<i>Occ. in year t</i>	<i>Occ. in year $t+1$</i>		
Purchasing Mgr. (11-3061)	11-3061 13-1040		
Compliance Off. (13-1040)		13-1040	13-1040

This matching of sequential job-year coincidence pairs results in 178.5 million observations (including year-to-year pairs where workers are observed in the same occupation in both years). We use these sequential job-year 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 sequential job-year 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).⁸⁵ We then aggregate them, reweighting by the average proportion of employment in each of these age categories in that occupation in the US 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, reducing the potential for bias arising from the skewed age distribution of our sample.

Summary statistics

Below, we describe the characteristics of the BGT resume data and how it compares to other data sets. All statistics refer to the final set of 15.8 million filtered resumes, or 178.5 million observations of sequential job-year coincidence pairs (‘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 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 A2 shows more information on the distribution of job incidences and job durations on our resumes.

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 be more likely to be male, and 50% are imputed to be 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 A2). It is, however, to be expected that the sample of the resumes which *provide* educational information are biased

⁸⁵Where we impute age based on the year in which the worker finished either college or high school, as described in the next section.

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 not be a reason for too much 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 A3 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, 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.⁸⁶

Figure A4 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.

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 (if any address is provided). 115.4 million of our observations come from resumes that list an address in the 50 US states or District of Columbia. The broad patterns of the demographic distribution of populations across the US is reflected in our data. By Census region, 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: that is, our sample is very close to representative for the Midwest and South regions, somewhat

⁸⁶The occupations with more than 2 million observations are: General and Operations Managers; Sales Managers; 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.

overweights the Northeast, and underweights workers from the West region. Zooming in on US states (Figure A5), we see that New Jersey, Maryland and Delaware, for instance, are 1.5-2x as prevalent in our data as they are in the overall US 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 US labor force.

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 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). 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.⁸⁷

⁸⁷In 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 (although we do not take advantage of this feature of the data in this paper).

Caveats and concerns

The BGT resume data set 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 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 executives in banking and finds that 51 responding executives were jointly aware of a total of 17 instances of meaningfully falsified job titles, which seems small given the presumably large number of resumes that these executives would have processed during their careers. 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, Friedmann, Nagler and Dinlenc (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. We exclude these from the data, but there may be other parsing errors we are unable to detect).

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.

Comparability with CPS occupational mobility

The average occupation “leave share” in our BGT resume data is 23%. This is roughly the probability that a worker leaves their SOC 6-digit occupation when they leave their job. 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.

To what extent is our measure similar to measures of occupational mobility constructed from the CPS? Our measure is not strictly comparable to the concept of annual occupational mobility estimated from the CPS by Kambourov and Manovskii (2008) and Xu (2018) for two reasons. First, the occupation categorization is different: we use SOC 6-digit occupations (of which there are a total of 840 in US data) and the CPS uses Census occupation codes, which are slightly broader. Second, because of the nature of our resume data, we cannot measure annual occupational mobility (share of workers whose main job was in occupation o on date d in year t whose main job was no longer in occupation o on date d in year $t + 1$). Instead, our measure of the average share of workers leaving their occupation in any given

year (11%) reflects the total number of workers who are observed in occupation o in year t who are *not* observed in occupation p at any point in year $t + 1$. This makes it a slightly more conservative measure of occupational mobility than the annual occupational mobility concept commonly constructed from the CPS.

With these caveats in mind: our measure of occupational mobility – the share of workers leaving their occupation being 11% from one year to the next – 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. Our measure is, however, 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 might 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.

Our “occupation leave share” represents not *unconditional* annual occupational mobility but rather the degree of outward occupational mobility *conditional* on leaving the worker’s initial job. We find that 46% of workers in our data are observed in some new job from one year to the next. This is consistent with the average length of a job in our data being 2 years. Note that according to the definition of a job we have chosen to work with, leaving your job does not necessarily entail leaving your firm: moving occupation or job title at the same firm would entail leaving your job. The CPS reports that median employee tenure at their firm 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.

D Appendix: OES Occupational Code Crosswalk

In our analysis of the effect of outside-occupation options on wages, we run some regressions over a longer period of 1999–2016. To construct our data set of wages and employment at the occupation-CBSA level over this period, 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.⁸⁸ 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 2010.⁸⁹

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.⁹⁰

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.

E Appendix: Alternative approaches to estimating occupational similarity

In Section 4 of this paper, we define workers' baseline labor market as a SOC 6-digit occupation within a metropolitan area.⁹¹

We then use occupational transitions to identify workers' outside options. There are two other possible methods of estimating occupational similarity to infer which jobs are good options for workers' outside their occupation: skill-and task-based similarity measures, and demographic- and qualification-based similarity measures. Why do we use occupational mobility?

⁸⁸Detailed breakdown of the affected codes available at: https://www.bls.gov/oes2010_and_2011_oes_classification.xls

⁸⁹This was the case for the following OES 2010 codes: 11-9013, 15-1799, 51-9151

⁹⁰This 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.

⁹¹We choose local SOC 6-digit occupations as our baseline labor market, rather than industries, since research on human capital specificity suggests that occupations are a more accurate approximation of the set of jobs open to workers (Kambourov and Manovskii, 2009; Sullivan, 2010). We choose a metropolitan area as an approximation of the jobs that are available to workers without having to move. A Commuting Zone would be a better geographic measure than a metropolitan area, but unfortunately the BLS data does not include wages by SOC 6-digit occupation at the Commuting Zone level.

To answer this question, we ask: What makes jobs in a given occupation a good outside option? Good outside option jobs should be both *feasible* in the sense that the worker can relatively easily become as productive as an average worker in that job, and should be at least somewhat *desirable* to work in (relative to the worker’s current job). We show that occupational mobility measures capture the underlying feasibility of a job transition, in the sense that they represent moves that people actually made. This means that they can capture many dimensions of feasibility of a transition – including task, skill, and amenity similarity, but also including other constraints that prevent moves in practice but may not be observed in task or skill data (e.g. regulation, occupational licensing barriers, etc.). Since occupational transitions also reflect moves people have (mostly) chosen to make, they also incorporate the desirability of moves between different occupations.

Skill- and task-based occupational similarity measures define two occupations as more similar, the more similar the skills and tasks are that they require. 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. A skill- or task-based measure of the similarity between two occupations does indeed capture many dimensions of the feasibility of an occupational transition. However, it has a number of weaknesses relative to a transition-based measure.

First, a skill- or task-based similarity measure cannot capture non-skill-related aspects which affect the feasibility of moving from one occupation to another occupation, such as occupational licensing or certification barriers between two occupations which may have similar skill requirements. Second, a skill- or task-based similarity measure 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 (for example, long or unpredictable hours being required may make an occupation less desirable for parents of young children) – 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. Fourth, skill- or task-based similarity measures require both the ability to *measure* the underlying skill and task requirements for each occupation with some accuracy *and* substantial assumptions as to how skill and task data should be combined to create a similarity measure. Skill- and task-based similarity measures can be highly sensitive

to these assumptions. In contrast, a transition-based measure has the advantage of being non-parametric. 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 (Bajari, Benkard and Levin, 2007).

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. (This is 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). 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 determining 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.

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.

F Appendix: Determinants of occupational mobility

In section 4 we showed that empirical occupational transitions reflect underlying similarity in occupations’ task and skill requirements and in their amenities. We explain this analysis in more detail here.

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). 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).

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, Levy and Murnane (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.⁹²

⁹²We update the task composites from Deming (2017) by using the latest source for task contents on

Job amenities. We measure similarity in the “temporal flexibility” of different occupations 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.⁹³ 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. 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.

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 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. Since this is a new composite measure of an important occupational characteristic, we outline it in more detail here.

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;

O*Net, and computing the composites at the level of SOC 2010 occupational codes.

⁹³The 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.

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.

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:

$$\pi_{o \rightarrow p} = \alpha_o + \beta^{abs} |X_{occ\ p} - X_{occ\ o}| + \gamma |\Delta w_{o \rightarrow p}| + \epsilon_{op}. \quad (26)$$

where $\pi_{o \rightarrow p}$ 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,⁹⁴ 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 3.⁹⁵

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 (26), but now using the *relative* (target

⁹⁴Amenities are most likely to be priced into wages (Goldin, 2014) and controlling for the latter would therefore be inappropriate.

⁹⁵Our findings build on Macaluso (2019), who showed that greater skill distance between SOC 2-digit occupations 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.

minus origin) difference in occupational characteristics as the independent variable:

$$\pi_{o \rightarrow p} = \alpha_o + \beta^{rel}(X_{occ\ p} - X_{occ\ o}) + \gamma \Delta w_{o \rightarrow p} + \epsilon_{op}. \quad (27)$$

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 (27) for the different measures are shown in Figure A11. (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.)

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;⁹⁶ 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 $\pi_{o \rightarrow p}$ 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 low.

The failure of skill similarity measures to explain many occupational transitions can be illustrated by a few cases from 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 vs. Medical & clinical laboratory technologists; Carpenters vs. Dental assistants; Travel agents vs. 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

⁹⁶These 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, for example, Acemoglu and Autor (2011), and the increasing demand for social skills documented by Deming (2017).

requirements. Second, consider another pair of occupations which are very similar on the skill distance metric (again, in the lowest distance decile): Pediatricians vs. 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.

G Appendix: IV analysis

Identification assumptions for concentration instrument

This section provides more formal details on the assumptions required for the IV identification of the effects of labor market concentration on wages. Our instrument can be interpreted as a type of granular IV following Gabaix and Koijen (2020), where market-level trends are instrumented for using idiosyncratic firm-level shocks (for details on the granular IV identification approach see Gabaix and Koijen (2020)). Or, it can be seen through the lens of the Bartik or shift-share IV approach, following Borusyak, Hull and Jaravel (2018), with exogenous ‘shocks’ in the form of differential national hiring patterns for large firms, and initial squared employer shares of each firm in a given local labor market determining the exposure to those shocks.

We can rewrite the concentration instrument as

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

where $\tilde{G}_{j,o,k,t} = \frac{(1 + \tilde{g}_{j,o,t})^2}{(1 + \tilde{g}_{o,k,t})^2} - 1$ is the predicted firm-level excess local vacancy growth relative to the average predicted local occupation vacancy growth - the time-varying shock - and $\sigma_{j,o,k,t-1}^2$ is the exposure of the local concentration index to that shock.⁹⁷

As noted in the main text, we add three controls to our baseline specification. To control for any effects on local labor demand of differential exposure to large national firms’ hiring, we control for (1) the growth rate of local vacancies in the occupation-city labor market ($g_{o,k,t} = \sum_j \sigma_{j,o,k,t-1} g_{j,o,k,t}$), and (2) the predicted growth rate of local vacancies based on large firms’ national growth ($\tilde{g}_{o,k,t} = \sum_j \sigma_{j,o,k,t-1} \tilde{g}_{j,o,t}$). To control for differential initial exposure to non-local firms, we introduce our “exposure control” $e_{o,k,t} = \sum_j \sigma_{j,o,k,t-1}^2 \cdot \mathbb{1}[\tilde{g}_{j,o,t} \neq 0]$.

⁹⁷For simplicity of exposition, we assume here that employer concentration and outside-occupation options are not correlated – but the logic of this argument does not depend on this assumption.

In our fixed effects IV estimation of equation (9), the exclusion restriction for the instrument on the HHI concentration index is then equivalent to

$$Cov[Z_{o,k,t}^{HHI}, \xi_{o,k,t} | \Gamma_{kt}, \Gamma_{ot}, g_{o,k,t}, \tilde{g}_{o,k,t}, e_{o,k,t}] = \mathbb{E} \left[\sum_{t=1}^T \sum_o^{N^{occs}} \sigma_{j,o,k,t-1}^2 \tilde{G}_{j,o,k,t}^\perp \xi_{o,k,t} \right] \rightarrow 0$$

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 our three control variables $g_{o,k,t}$, $\tilde{g}_{o,k,t}$, $e_{o,k,t}$.

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} , actual and predicted average local vacancy growth $g_{o,k,t}$ and $\tilde{g}_{o,k,t}$, and initial exposure to non-local firms $e_{o,k,t}$. That is,

$$\mathbb{E} \left[\sum_j \sigma_{j,o,k,t-1}^2 \tilde{G}_{j,o,k,t} | \xi_{o,k,t}, \Gamma_{kt}, \Gamma_{ot}, g_{o,k,t}, \tilde{g}_{o,k,t}, e_{o,k,t} \right] = \tau_1 \Gamma_{kt} + \tau_2 \Gamma_{ot} + \tau_3 g_{o,k,t} + \tau_4 \tilde{g}_{o,k,t} + \tau_5 e_{o,k,t} \quad \forall o \in N^{occs}$$

for some constant parameters τ_1 , τ_2 , τ_3 , τ_4 , and τ_5 . 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.⁹⁸

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

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

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

The first assumption requires that the local size-squared-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 the hypothetical example from the main text, which considered insurance sales agents in Bloomington, Illinois and in Amarillo, Texas. In each city, there are several insurance companies who employ insurance sales agents. Assume that in Bloom-

⁹⁸In a robustness check, we also include a control for average vacancy growth across firms within a local occupation, with each firm weighted equally, $\frac{1}{N} \sum_j g_{j,o,k,t}$. This is suggested by Gabaix and Koijen (2020) as an appropriate control for local demand effects in a granular IV setting, as it controls for the increase in vacancies experienced commonly across all firms in the local labor market. The identification assumptions in the specification with this control would require that local squared exposure to excess national firm-level growth is random in expectation conditional on this proxy for local labor demand (alongside the other controls and fixed effects already discussed).

ington, State Farm has a large share of local insurance sales agent employment, while in Amarillo employment is more dispersed amongst a number of insurance companies. In years where State Farm grows substantially faster than other major insurance companies nationwide, under most combinations of the distribution of that growth across cities and the initial distribution of employer shares in each city, employer concentration of insurance sales agents will grow by more in Bloomington IL than in Amarillo TX. Moreover, our granular IV identification approach controls for local growth rates of overall insurance sales agent employment in both cities. Thus, it allows for each city to be exposed differently to overall trends in the demand for insurance sales agents. The identification only requires that once we account for overall city exposure to insurance sales agent demand, whether that demand was driven by the city’s major employer or smaller employers is not correlated with local idiosyncratic wage shocks for insurance sales agents.

How does the first-stage assumption work? The first stage of our regression holds if, when firm j grows nationally, local occupation-city labor markets with a higher share of vacancies accounted for by firm j in year $t - 1$ see a larger increase in employer concentration. A sufficient condition for this to be the case under *most* initial employer share distributions is if firm j ’s new vacancies are allocated evenly across occupation-city labor markets, such that each occupation-city labor market sees the same growth rate in its firm j vacancies as the national average.⁹⁹ However, this condition is not necessary: in fact, the first stage can be valid even if the growth rate of firm j ’s new vacancies in low-initial-employment-share occupation-city labor markets is higher than in high-initial-employment-share labor markets, as long as this relationship is not too strong. In our data, for a given employer, there is a negative relationship between the initial vacancy share in an occupation-city labor market and the next year’s vacancy growth rate, but this relationship is not sufficiently strong to invalidate our first stage (for each one percentage point increase in the initial vacancy share, there is roughly a 1.4 percentage point lower vacancy growth rate from one year to the next). Empirically, our first stage holds for occupation-city labor markets with HHIs above all but

⁹⁹Note that this is not the case – i.e. the first stage might not hold – for *all* possible combinations of the distribution of employment growth and initial employer shares. For example, consider a world in which there is a labor market for where Employer X has 80% of the market in one city, and the rest of the market is comprised of atomistic firms; and Employer X has 65% of the market in another city, with the rest of the market comprised of atomistic firms. If Employer X grows by 10% in both locations in a given year, and the other firms do not grow at all, employer concentration will actually increase by more in the latter than the former market. This circumstance, however, only occurs when comparing two labor markets which both have extremely high levels of employer concentration already, and so is not relevant for the vast majority of the labor markets in our data (only 6% of which have HHIs of greater than 6,800 – and only 2% of workers in our data face HHIs of greater than 5,000). In practice our first stage regressions are positive and strongly significant even when segmenting our data to analyze only cells with high HHIs of 2,500 or more, meaning this concern is not hugely relevant in practice for our empirical analysis.

very low levels, as shown in the main text (and in Appendix Tables A7 and A13).

Identification assumptions for outside-occupation option index instrument

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.

As described in Section 5, our instrument for the oo^{occs} index, Z^{oo} , is the weighted average of national leave-one out mean wages in occupation p , $\bar{w}_{p,k,t}$, where the weights are the product of the year 1999 relative employment share in each of those occupations in the worker's own city, $\frac{s_{p,k,1999}}{s_{p,1999}} \cdot 100$ and the national occupation transition shares from the worker's occupation o to each of the other occupations, $\pi_{o \rightarrow p}$:

$$Z_{o,k,t}^{oo} = \sum_p^{N_{occs}} \left(\pi_{o \rightarrow p} \cdot \frac{s_{p,k,1999}}{s_{p,1999}} \cdot \bar{w}_{p,k,t} \right) \quad (28)$$

To make the assumptions transparent under which this wage instrument identifies the coefficient on our outside-occupation option index in equation (9), we again follow the framework presented in Borusyak et al. (2018).¹⁰¹ Note that we can write the instrument as

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

where $s_{okp} = \pi_{o \rightarrow p} \cdot \frac{s_{p,k,1999}}{s_{p,1999}}$ is a measure of predicted local exposure to the shock. In our fixed effects IV estimation of equation (9), the exclusion restriction for the instrument for outside-occupation options is then equivalent to

$$Cov[Z_{o,k,t}^{oo}, \xi_{o,k,t} | \Gamma_{kt}, \Gamma_{ot}, g_{o,k,t}, \tilde{g}_{o,k,t}, e_{o,k,t}] = \sum_{t=1}^T \sum_{p=1}^{N_{occs}} \bar{s}_{okp} w_{p,k,t}^\perp \phi_{pt}^{oo} \rightarrow 0$$

where $\bar{s}_{okp} = \mathbb{E}[s_{okp}]$ is the average exposure to occupation p , and $\phi_{pt}^{oo} \equiv \mathbb{E}[s_{okp} \xi_{o,k,t}] / \mathbb{E}[s_{okp}]$ is an exposure-weighted expectation of the structural wage residuals. Moreover, $w_{p,k,t}^\perp$ represents $\bar{w}_{p,k,t}$ 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 and control variables.

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

¹⁰⁰Or the first year in the data, if there is no data for the occupation-city cell in 1999.

¹⁰¹For simplicity, assume that the outside-occupation option index and the concentration index are not correlated - but the intuition for the identification does not depend on that.

conditional on local exposure to structural wage shocks ϕ_{pt} , the fixed effects Γ_{kt} and Γ_{ot} , and the control variables. That is,

$$\mathbb{E}[\bar{w}_{p,k,t} | \phi_{pt}^{oo}, \Gamma_{kt}, \Gamma_{ot}, g_{o,k,t}, \tilde{g}_{o,k,t}, e_{o,k,t}] = \tau_1 \Gamma_{kt} + \tau_2 \Gamma_{ot} + \tau_3 g_{o,k,t} + \tau_4 \tilde{g}_{o,k,t} + \tau_5 e_{o,k,t} \quad \forall p \in N^{occs}$$

for some constant parameters τ_1 through τ_5 . Second, there needs to be a large number of independent occupational shocks, that is,

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

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

The first assumption requires that the national leave-one-out mean wage $\bar{w}_{p,k,t}$ in outside option occupation p is correlated with the local wage of occupation p 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 $Z_{o,k,t}^{oo}$. 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 .¹⁰² The inclusion of these fixed effects increases our confidence that the assumptions for instrument validity hold.

Industry Bartik shock

One possible concern with the identification assumptions required for our outside-occupation index – which may not entirely be picked up by our occupation-year, or city-year fixed effects – is that industry-level wage trends may differentially impact local occupations based on their city’s direct exposure to those industries, rather than only based on indirect exposure through outside occupation job options. As discussed in the text, an example of this could be the following. Imagine that the finance industry and the tech industry employ both accountants and data scientists to a disproportionate degree relative to other occupations, and that San Francisco has a large share of employment in tech while New York has a large share of employment in finance. Imagine further that being a data scientist is a good outside option occupation for an accountant. In years where the tech industry is booming nationwide, this will impact San Francisco more than New York. Accountants in San Francisco will see wages rising by more than accountants in New York – partly driven by the increase in the outside option value of becoming a data scientist, but partly simply because more accountants in SF already work in the tech industry, as compared to accountants in NY, and so they will

¹⁰²As an example, note that national-level correlation in the wages of a pair of occupations (e.g. Compliance Officers and Financial Analysts), perhaps due to common industry shocks, 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.

see their wages rise by more.

To control for this possible omitted variable bias, we incorporate an industry Bartik shock in a robustness check for our baseline regressions. We construct the industry Bartik shock for each occupation-city-year cell. The industry Bartik shock for occupation o in city k in year t is defined as

$$\sum_{\iota} \text{industries} \frac{\frac{emp_{\iota,o,t-1}}{emp_{o,t-1}} \cdot \frac{emp_{\iota,k,t-1}}{emp_{k,t-1}} \cdot \bar{w}_{\iota,t}}{\frac{emp_{\iota,o,t-1}}{emp_{o,t-1}} \cdot \frac{emp_{\iota,k,t-1}}{emp_{k,t-1}}}$$

where ι denotes each NAICS 4-digit industry. The shock to each industry is the national industry wage $\bar{w}_{\iota,t}$, and the exposure of each local occupation to that shock is determined by (1) the national share of employment in that occupation which is in industry ι , $\frac{emp_{\iota,o,t-1}}{emp_{o,t-1}}$ and (2) the share of employment in that city which is industry ι , $\frac{emp_{\iota,k,t-1}}{emp_{k,t-1}}$. The exposure measures are lagged by one year to avoid the possibility of endogenous responses of employment to the industry-level shock in question. The Bartik instrument relies on the assumption that national industry-level wage shocks are uncorrelated with local occupation-level wage trends, except to the extent that the former causes the latter. We use data on employment by NAICS 4-digit industry and SOC 6-digit occupation from the Bureau of Labor Statistics Occupational Employment Statistics to construct the employment shares of each occupation by industry, and we construct data on industry employment shares by metropolitan statistical area (“city”) from Eckert, Fort, Schott and Yang (2020)’s county-by-industry employment data, which was constructed from the County Business Patterns database. We report our baseline regression results, controlling for this industry Bartik shock, in Panel A of Table A9.

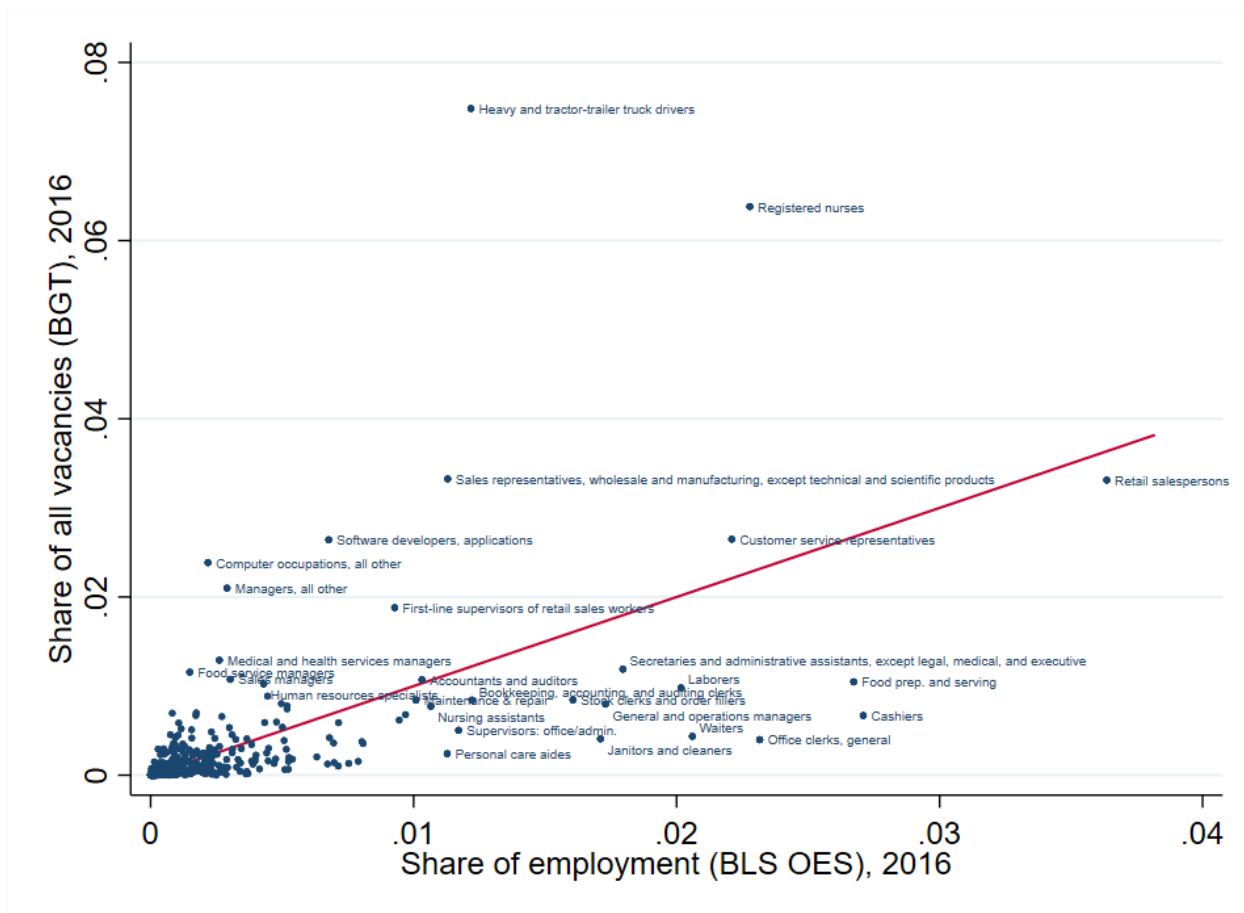
In other unreported specifications, we construct the industry Bartik shock with, variously, national wage growth in each industry ι , national employment growth in each industry ι , or the national wage bill growth in each industry ι . In all these specifications, our coefficients on the outside-occupation option index and HHI index (instrumented) remain very similar to those in the baseline specification.

H Appendix: Stata commands

In our estimation, we used a number of user-written Stata commands: *reg2hdfe* (Guimaraes and Portugal, 2010), *reghdfe* (Correia, 2016), *ivreg2hdfe* (Bahar, 2014), *binscatter* (Stepner, 2013), *binscatter2* (Droste, 2019), and *coefplot* (Jann, 2013).

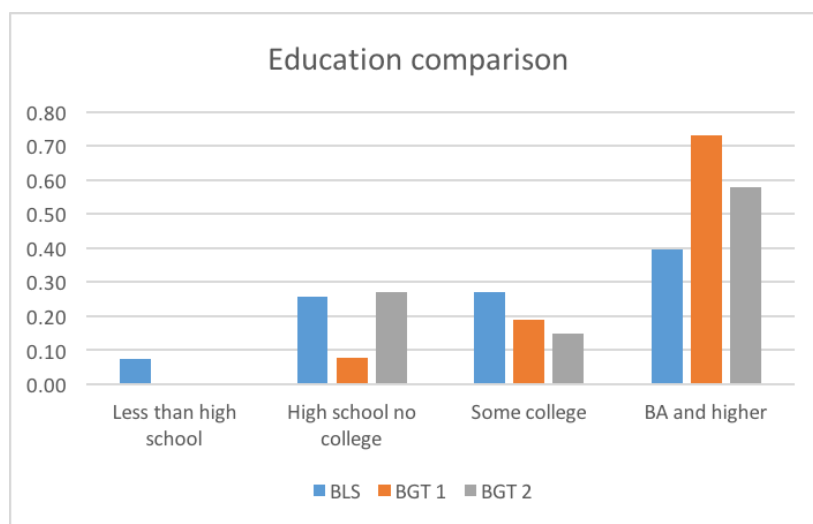
I Appendix: Figures

Figure A1: BGT Vacancy Data: representedness of occupations, relative to BLS OES



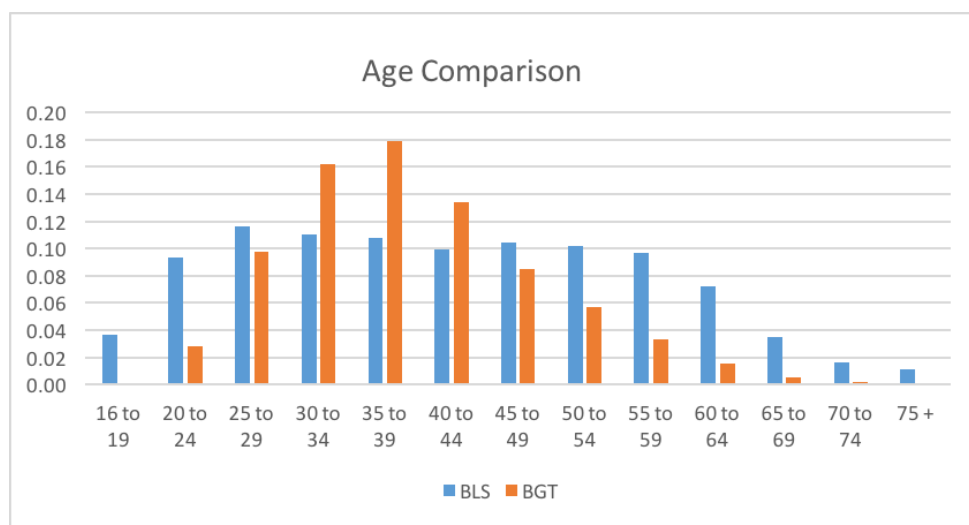
Note: Comparison of distribution of share of each SOC 6-digit occupation in the BGT vacancy data, relative to its share in the BLS occupational employment statistics, with occupations comprising greater than 1% share of either data set labeled. Red line is the 45 degree line. The vacancy data is discussed in detail in Appendix Section B.

Figure A2: BGT Resume Data: education relative to 2018 labor force



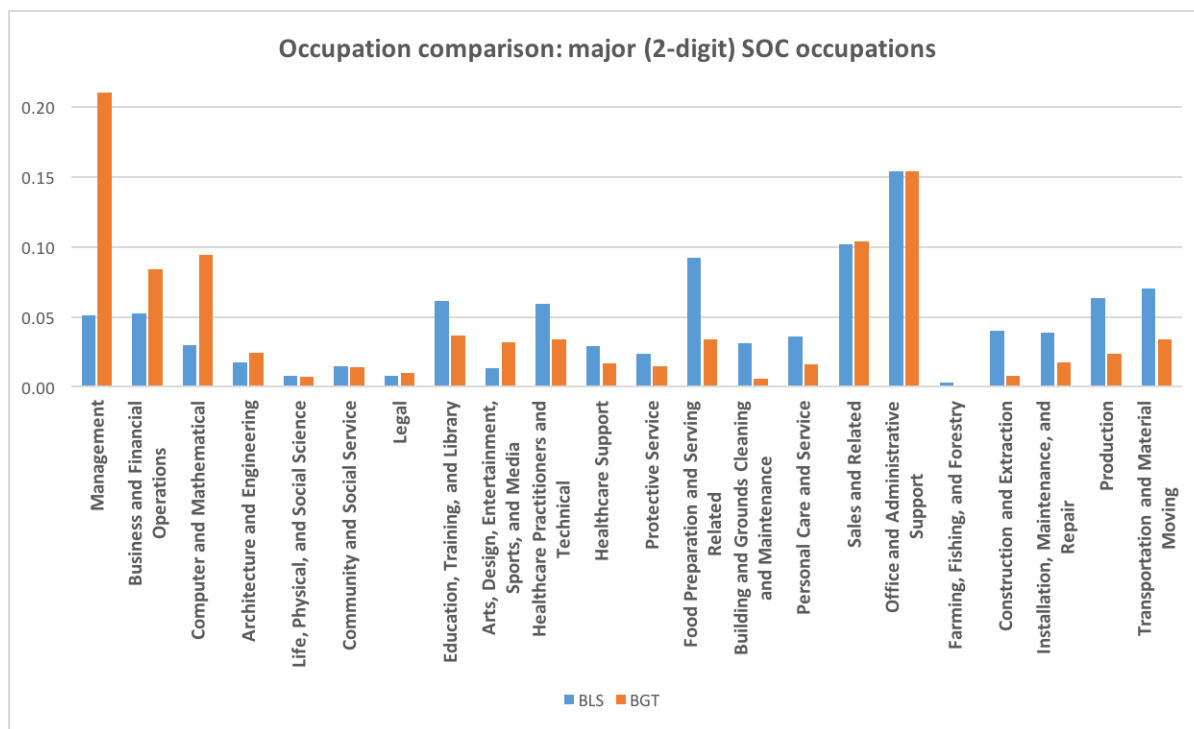
Note: Comparison of distribution of highest educational attainment in the labor force, according to BLS data, to distribution in BGT resume data. Two versions are shown: BGT 1 excludes all resumes missing educational information, while BGT 2 assumes all resumes missing educational information have high school education but no college. The resume data is discussed in detail in Appendix Section C.

Figure A3: BGT Resume Data: age distribution relative to 2018 labor force



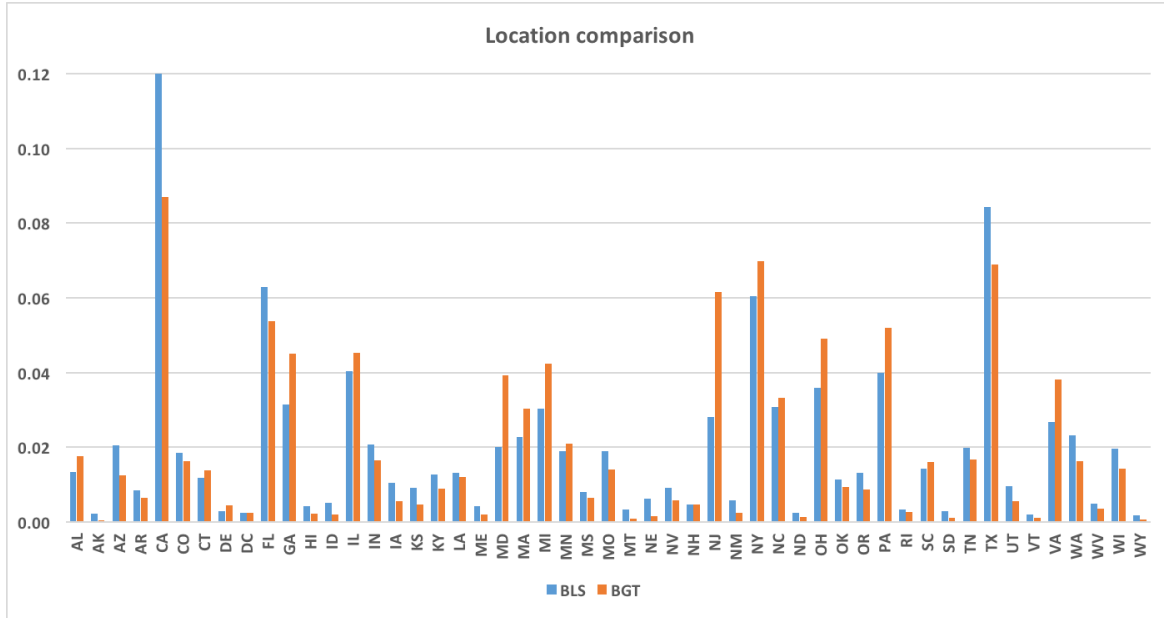
Note: Comparison of distribution of age in the labor force, according to 2018 BLS data, to distribution of imputed worker ages in BGT resume data. The resume data is discussed in detail in Appendix Section C.

Figure A4: BGT Resume Data: occupations relative to 2017 labor force



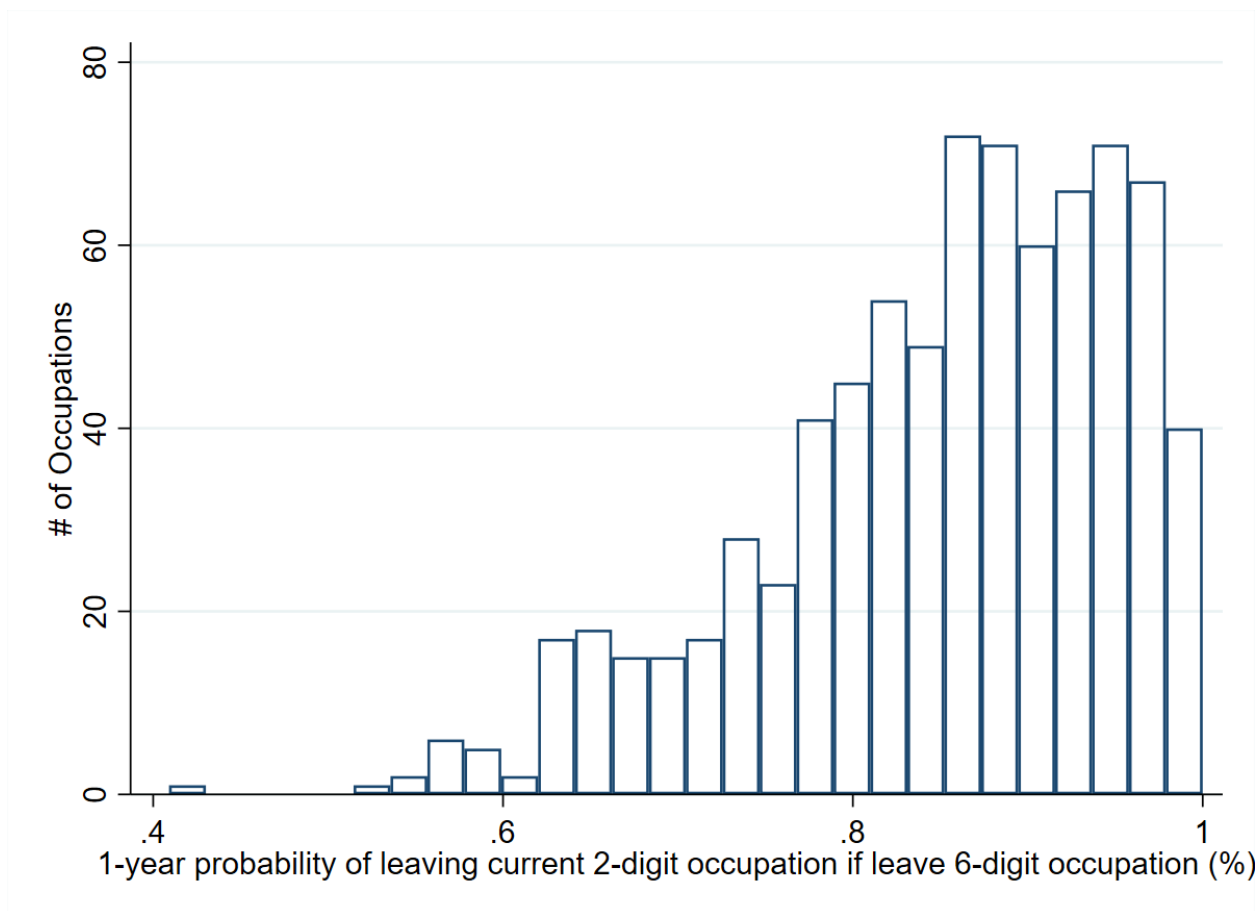
Note: Comparison of distribution of 2-digit SOC occupations in the labor force, according to 2017 BLS data, to distribution of occupations in BGT resume data. The resume data is discussed in detail in Appendix Section C.

Figure A5: BGT Resume Data: locations relative to 2017 labor force



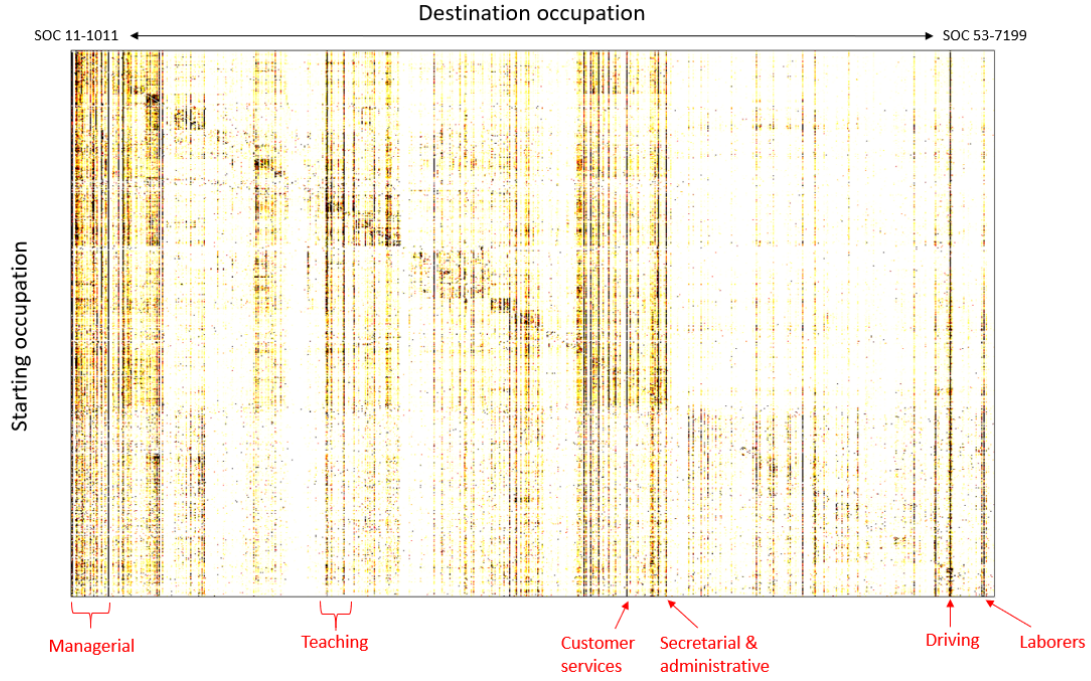
Note: Comparison of distribution of employment by U.S. state, according to 2017 BLS data, to distribution of resume addresses in BGT resume data. Graph shows share of total in each state. The resume data is discussed in detail in Appendix Section C.

Figure A6: Occupational mobility: SOC 6-digit moves that are also 2-digit moves



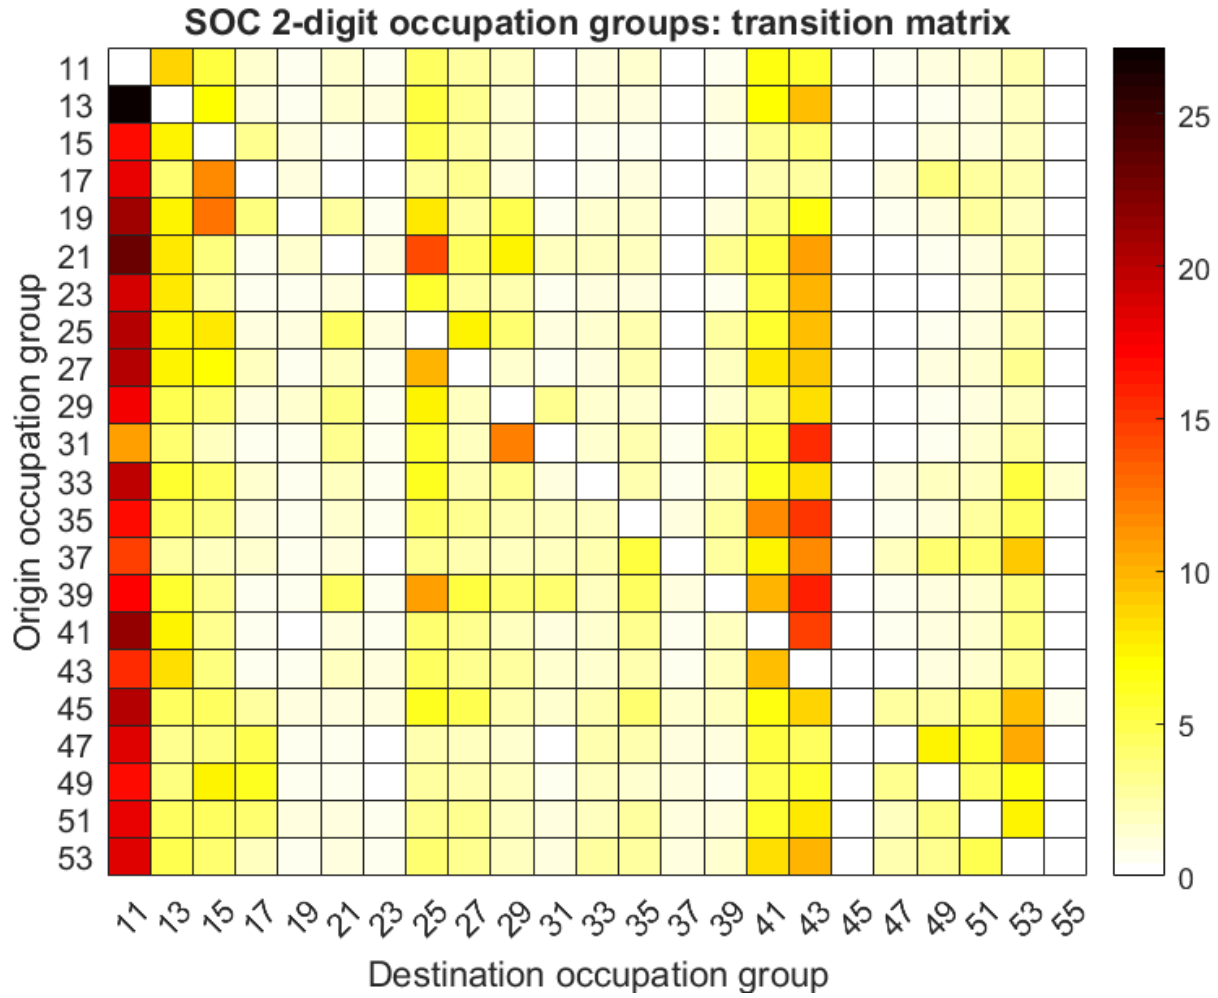
Note: Distribution of the proportion of workers moving 6-digit SOC occupation who *also* move 2-digit SOC occupation, by occupation, calculated from BGT resume data for 2002-2015 period. Histogram shows 786 occupations. The resume data is discussed in detail in Appendix Section C.

Figure A7: 6-digit SOC occupational transition matrix



Note: 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. The presence of a darker line along the diagonal suggests that workers commonly transition to occupations which are close to their own according to the numerical order of SOC codes. The resume data is discussed in detail in Appendix Section C.

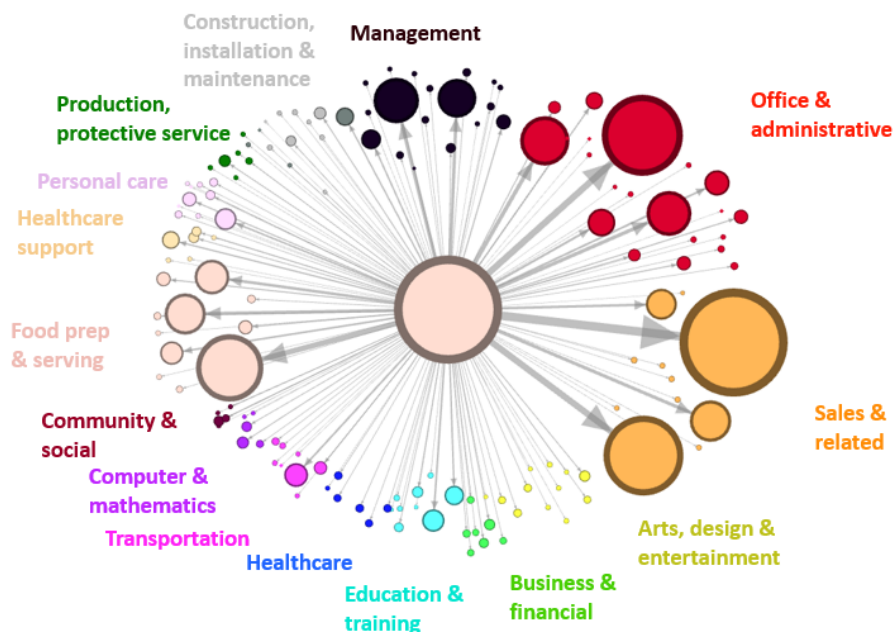
Figure A8: 2-digit SOC occupational transition matrix



Note: 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. The resume data is discussed in detail in Appendix Section C.

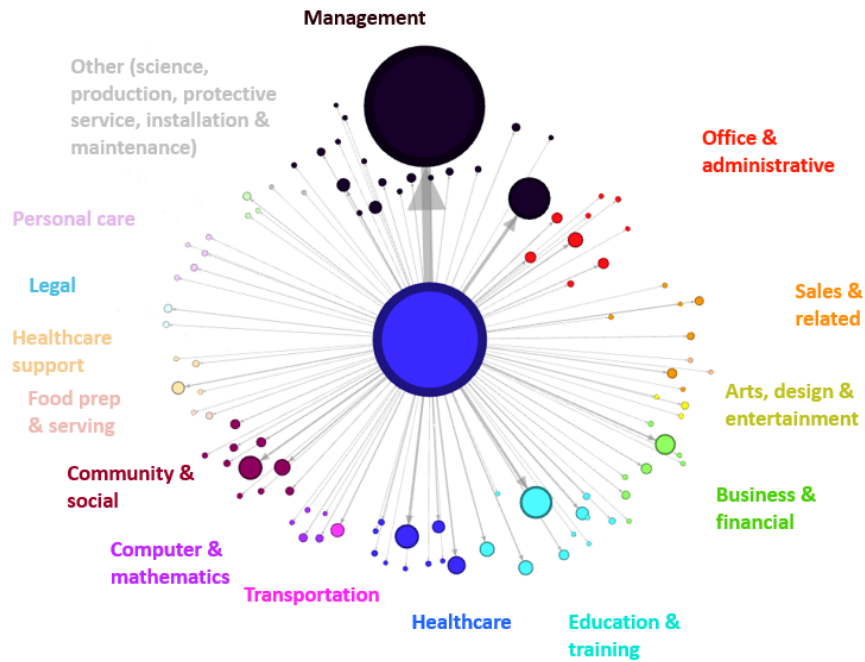
Figure A9: Examples of probabilistic labor markets: counter attendants

Which occupations do counter attendants (in food service) go to?



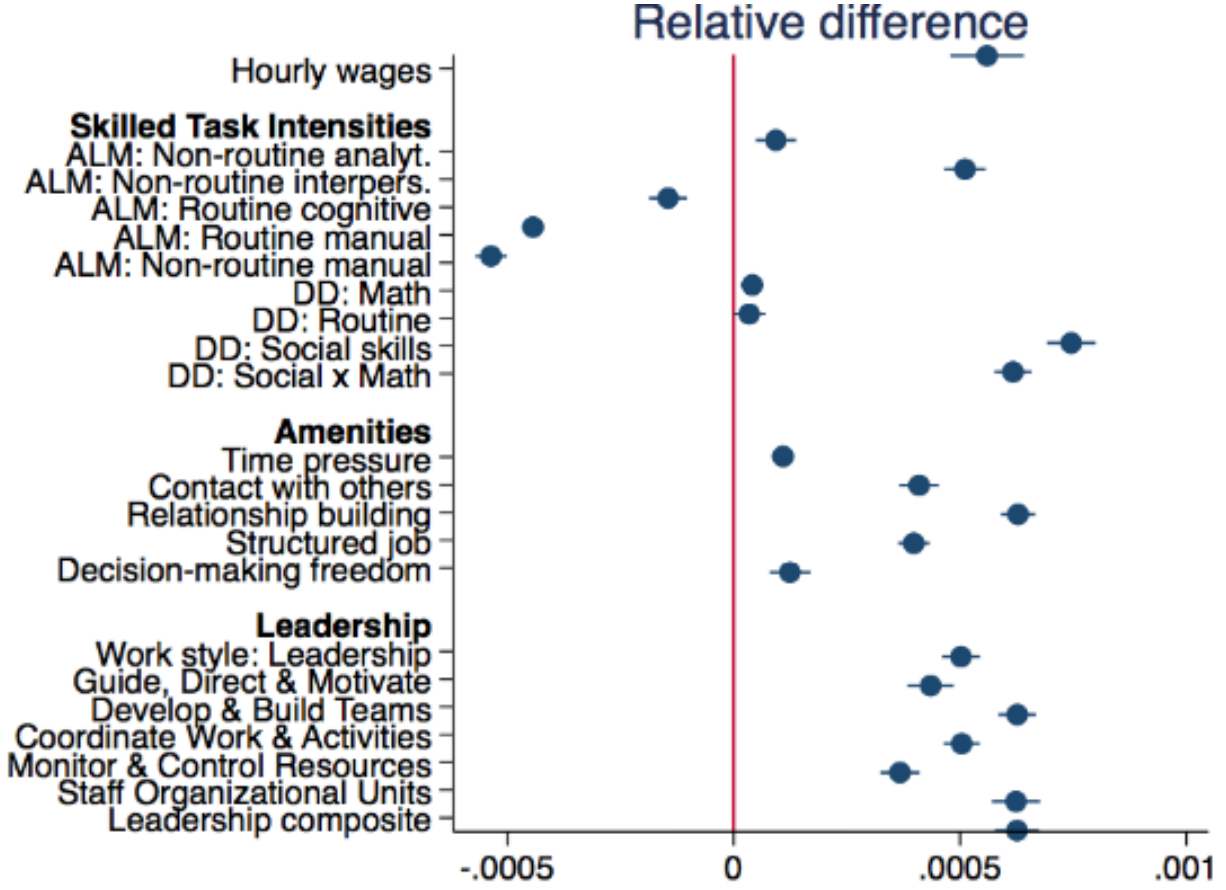
Note: Example visualization of occupational transitions for counter attendants in the food industry. Each bubble is a SOC 6-digit occupation, and the colors represent SOC 2-digit occupational groups. The size of each bubble is proportional to the share of counter attendants in the BGT data who switch occupation, who are observed in each destination occupation in the following year. The resume data is discussed in detail in Appendix Section C.

Figure A10: Examples of probabilistic labor markets: registered nurses
Which occupations do registered nurses go to?



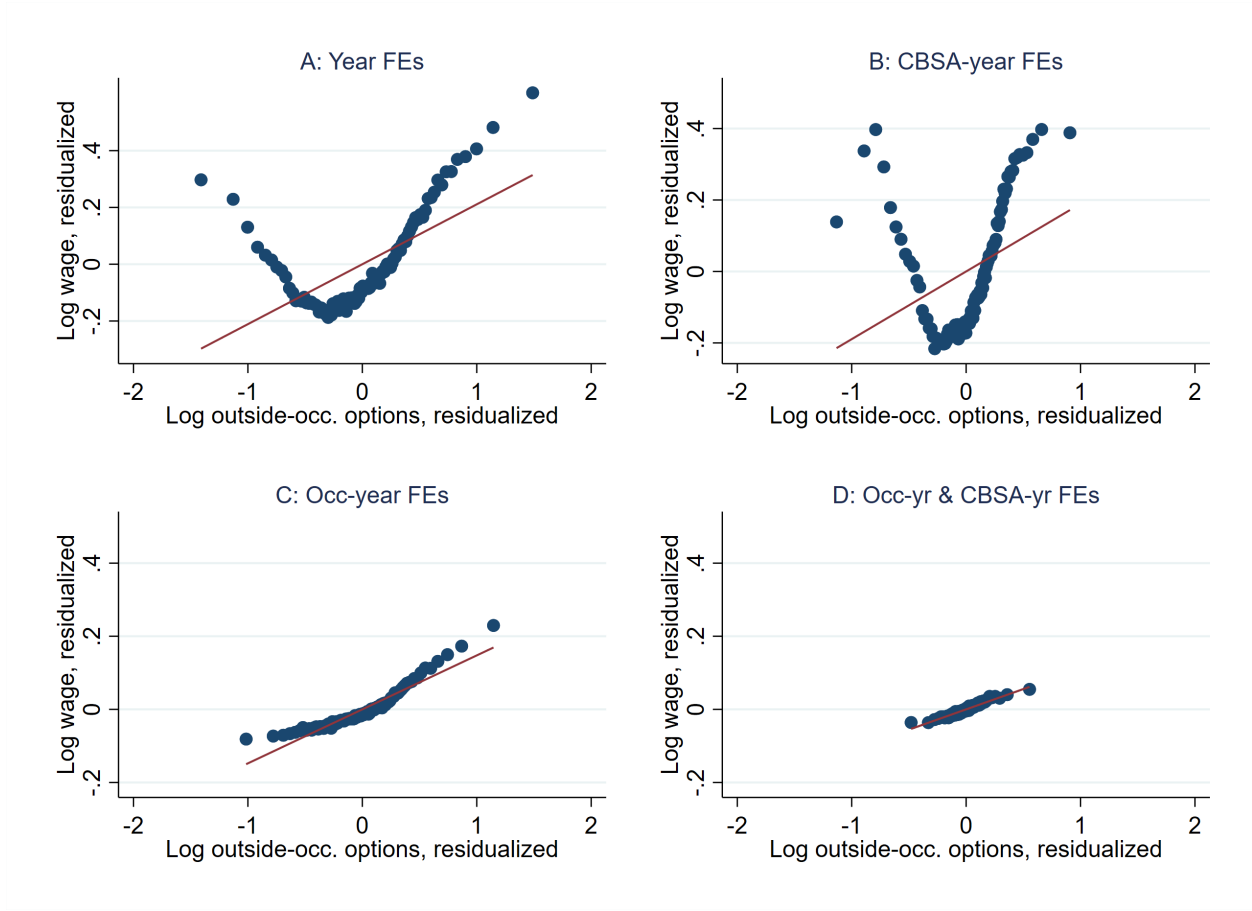
Note: Example visualization of occupational transitions for registered nurses. Each bubble is a SOC 6-digit occupation, and the colors represent SOC 2-digit occupational groups. The size of each bubble is proportional to the share of registered nurses in the BGT data who switch occupation, who are observed in each destination occupation in the following year. The resume data is discussed in detail in Appendix Section C.

Figure A11: Determinants of occupational mobility



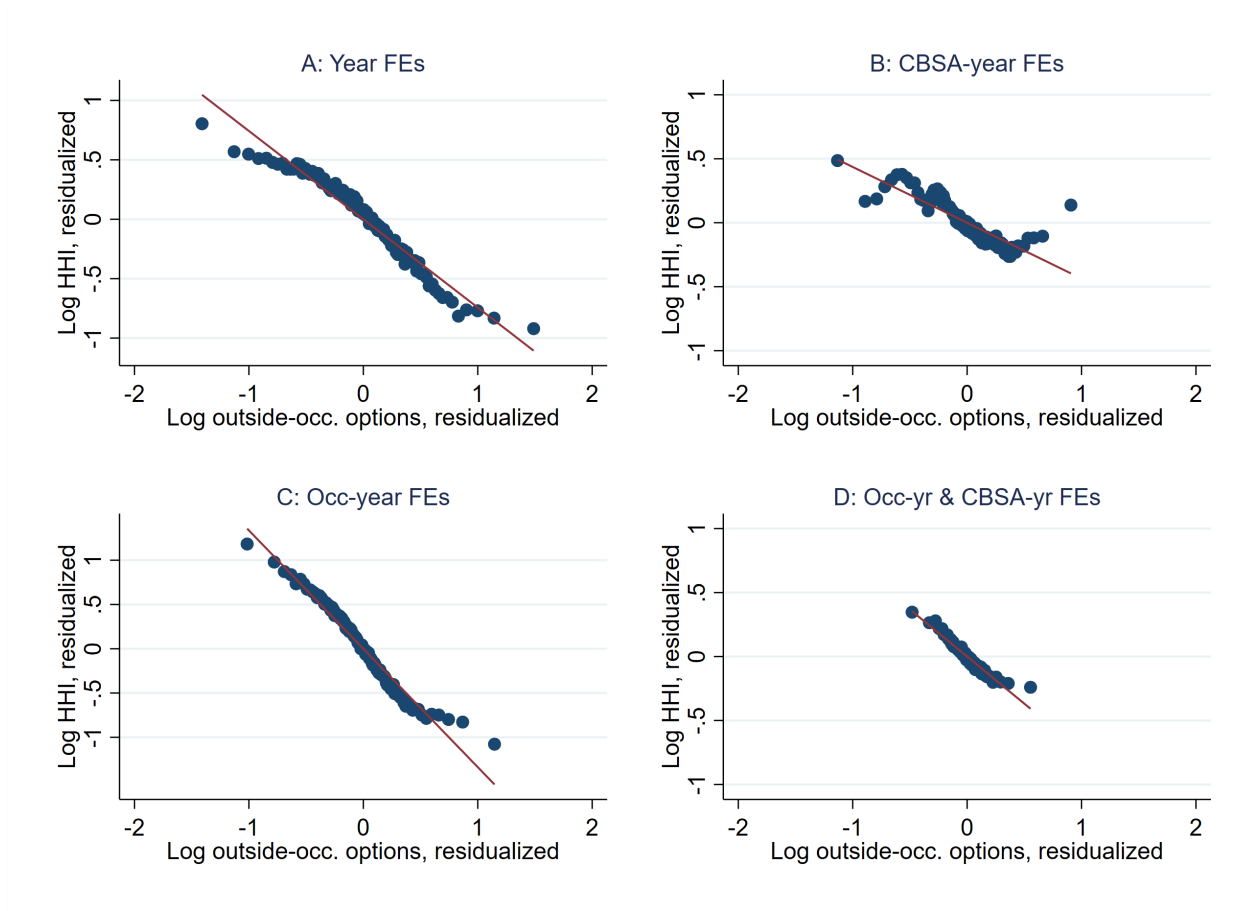
Note: This plot shows the coefficients and 95% confidence intervals from regressions of occupation transition shares $\pi_{o \rightarrow p}$, calculated from Burning Glass Technologies Resume data, on *relative* differences in occupational characteristics: $\pi_{o \rightarrow p} = \alpha_o + \beta f(X_{occ\ o \rightarrow p}) + \gamma f(\Delta w_{o \rightarrow p}) + \epsilon_{op}$ where the function $f(\cdot)$ represents the difference in characteristic between starting occupation o and destination p , and α_o is occupation o fixed effect. Regressions also include absolute avg. hourly wage differences (except for the amenities regressions). Standard errors are clustered at the origin occupation level. This is the analog of Figure 3, which shows coefficients on the regression of occupation transition shares on the *absolute* difference in characteristics between the pairs of occupations. These analyses are discussed in more detail in Appendix Section F.

Figure A12: Correlations between wage and outside-occupation option index



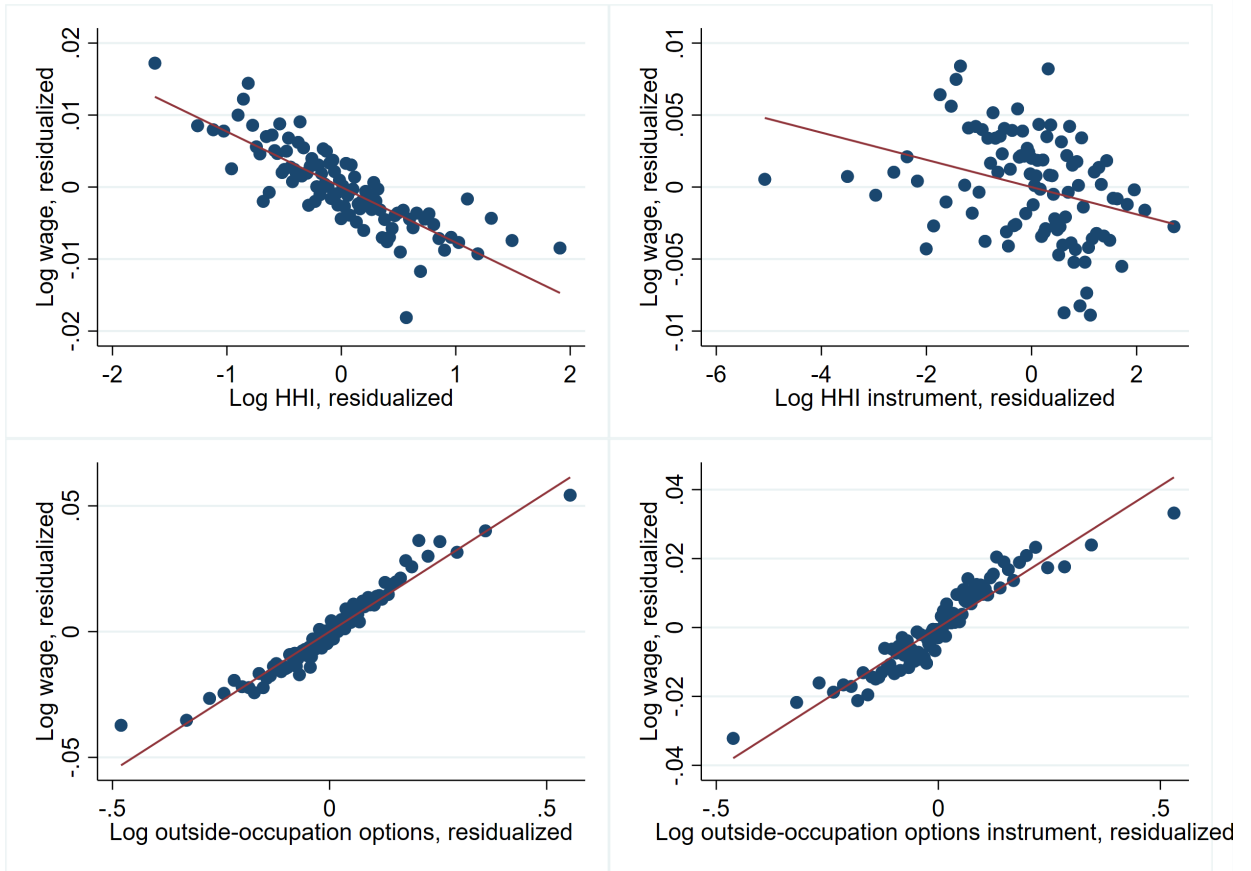
Note: Figure shows binned scatter plots of the correlation between average log wages and log outside-occupation option index for occupation-CBSA cells over 2013–2016, residualized on different combinations of fixed effects (as described by the panel titles). Regression coefficients for the line of best fit on each graph are: A: 0.16, B: 0.12, C: 0.15; D: 0.11. The non-linear shape of the figures without occupation fixed effects (panels A and B) is explained by healthcare occupations which tend to have both low outward mobility and high pay.

Figure A13: Correlations between HHI and outside-occupation option index



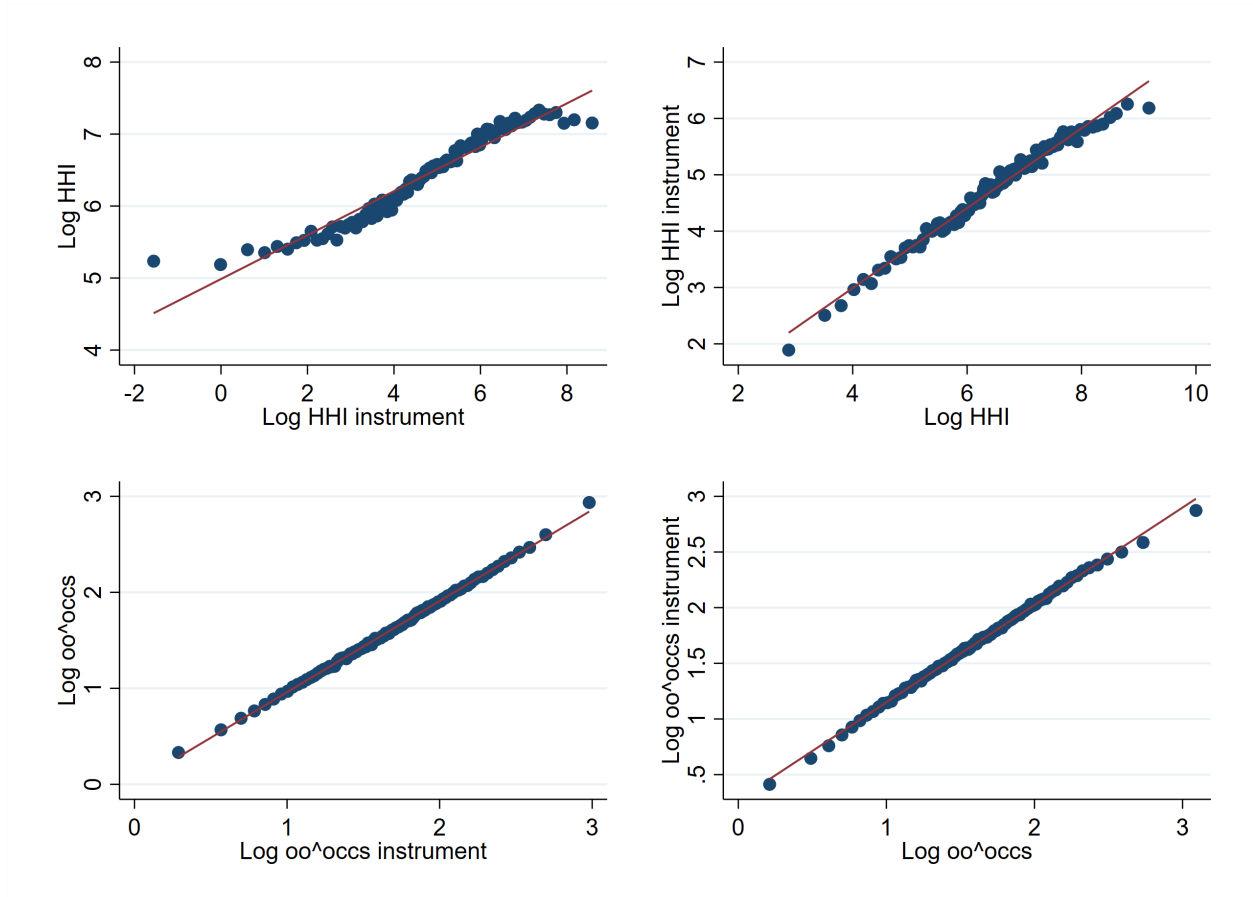
Note: Figure shows binned scatter plots of the correlation between average log HHI and log outside-occupation option index for occupation-CBSA cells over 2013–2016, residualized on different combinations of fixed effects (as described by the panel titles). Regression coefficients for the line of best fit on each graph are: A: -0.74, B: -0.44, C: -1.34; D: -0.74.

Figure A14: Visualization of baseline regression results



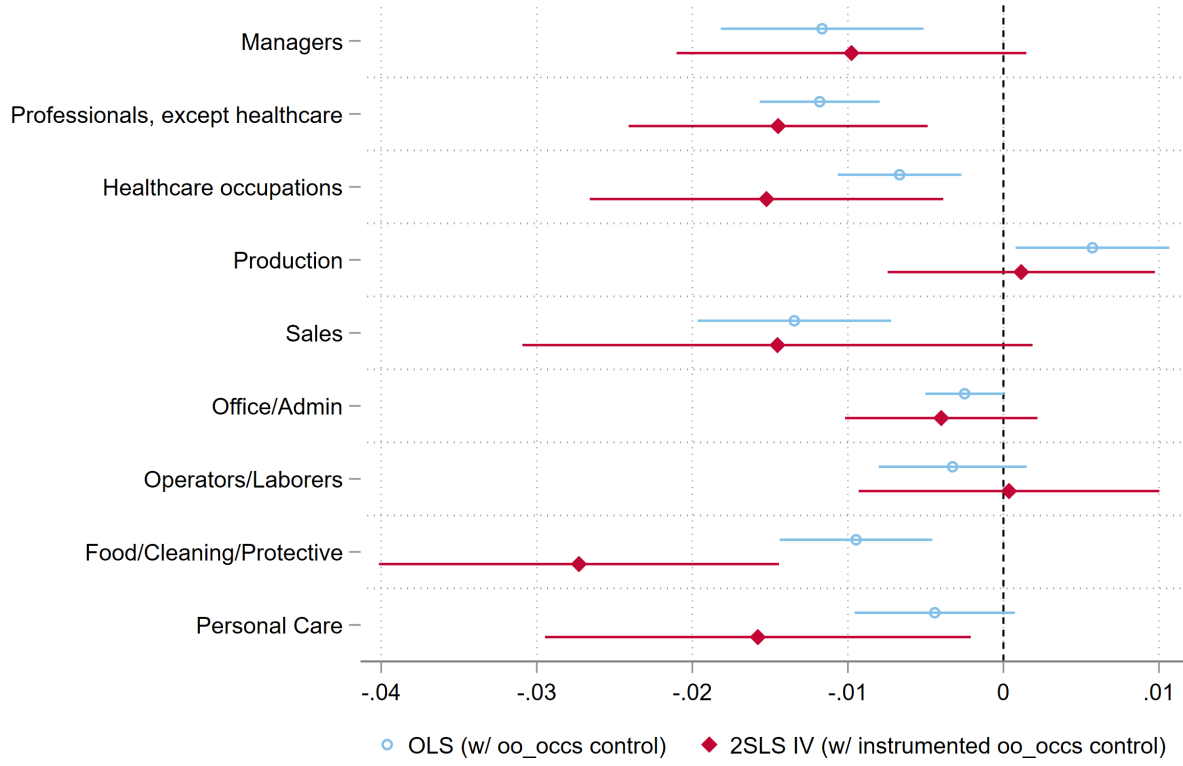
Note: Binned scatter plots of the log wage on log HHI and log outside-occupation options, and instrumented, including full controls as in baseline regression specification (i.e. the left panels correspond to coefficient estimates in Table 3 column (b), and the right panels correspond to the reduced form equivalent of the 2SLS IV coefficient estimates in Table 3 column (d)).

Figure A15: Visualization of first-stage regressions



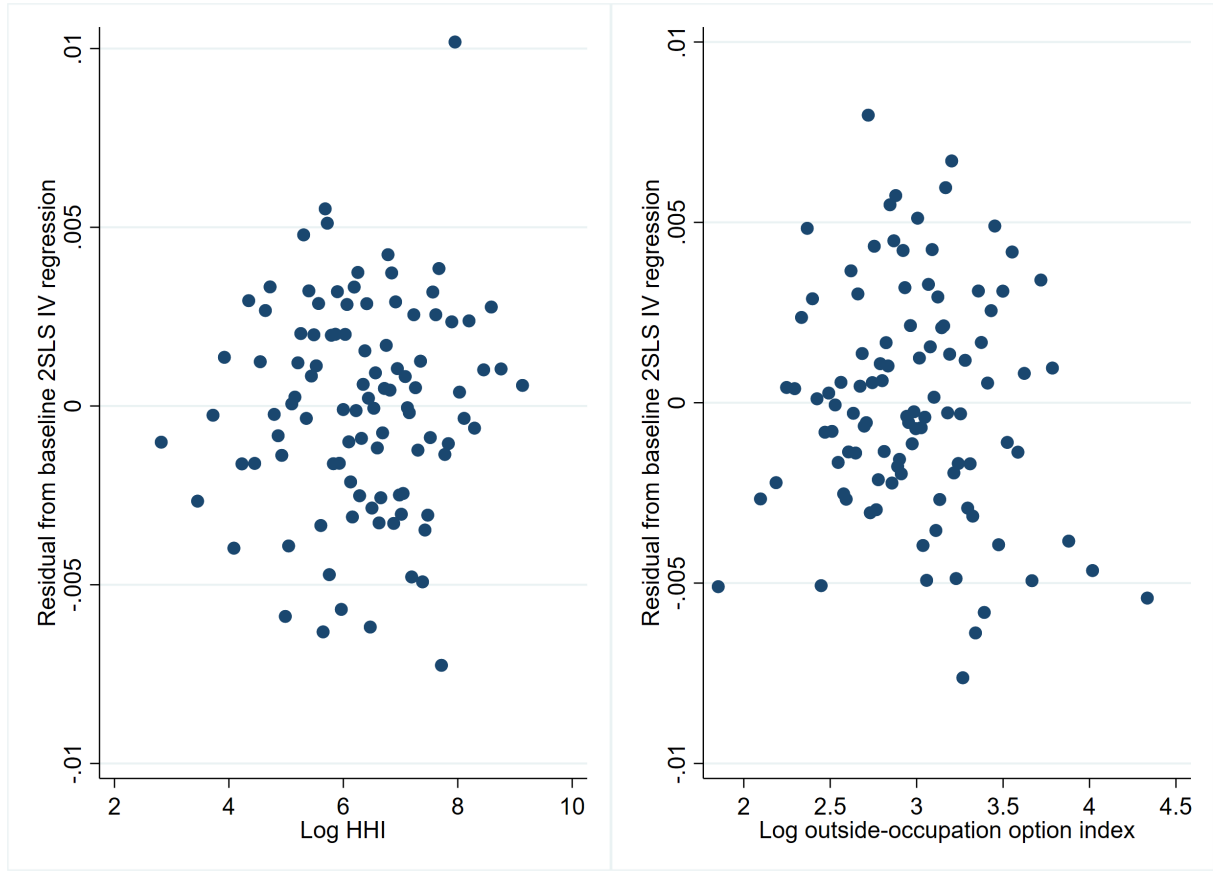
Note: Binned scatter plots of the correlation between the HHI instrument and raw variable (top panel) and outside-occupation option index instrument and raw variable (bottom panel) for occupation-CBSA cells in 2016.

Figure A16: Coefficients on wage-HHI regressions: by occupation group



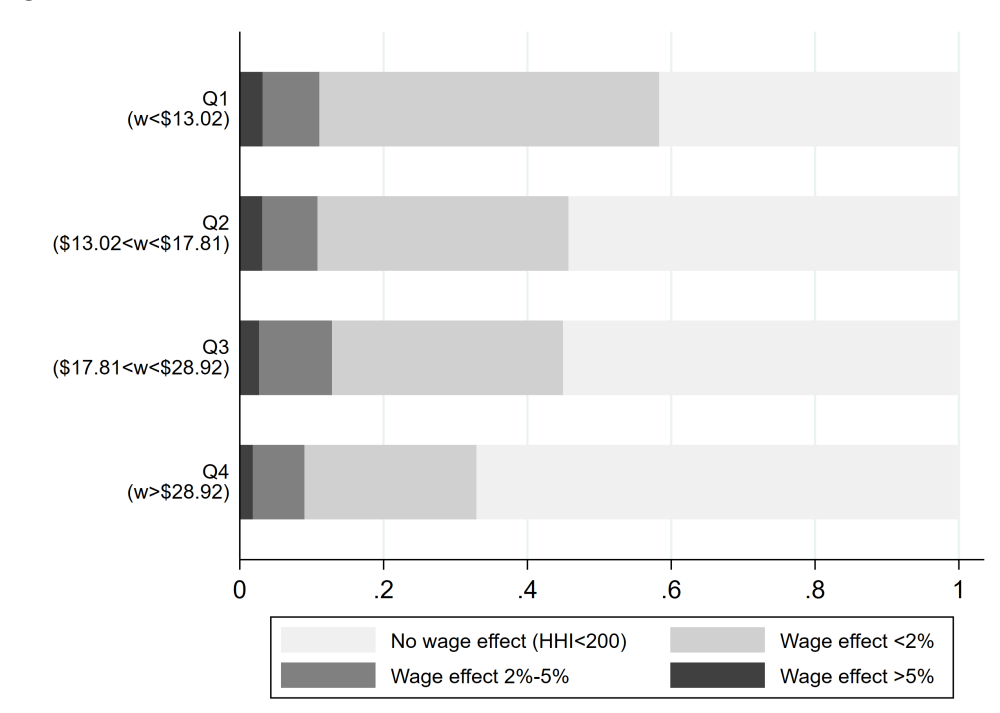
Note: Coefficients on HHI and 95% confidence intervals from regressions of occupation-CBSA wages on instrumented local employer HHI, controlling for (instrumented) outside-occupation job options, with coefficient on HHI allowed to vary by occupation group. Regressions span 2013-2016 and include occupation-year and CBSA-year fixed effects, as well as the other controls included in our baseline 2SLS IV regressions (described in Section 5). Standard errors are clustered at the CBSA level. Occupation groups are listed in descending order of average wage. Occupation groups map from SOC 2-digit occupations as defined in Appendix Table A6. Coefficient estimates correspond to those in Appendix Table A15 column (a).

Figure A17: Baseline regressions: residual plots



Binned scatter plots of residuals from baseline 2SLS IV regression of wage on log HHI and log outside-occupation option index, against the log HHI and log outside-occupation option index respectively. (The baseline regression results are reported in column (d) of Table 3).

Figure A18: Average estimated effect of employer concentration relative to HHI=200, by hourly wage



Note: This figure shows what share of workers may have experienced different degrees of wage suppression as a result of employer concentration, across the US hourly wage distribution, relative to a counterfactual where HHI was 200 (holding all else constant). We estimate the wage effect of employer concentration as described in section 8.

Figure A19: Average estimated effect of employer concentration relative to HHI=200, by metro area



Note: This figure shows the average estimated wage effect of concentration in each metro area relative to a counterfactual HHI of 200 (holding all else constant), plotted against the average hourly wage in that metro area in 2016 (left) and the employment in that metro area in 2016 (right) according to BLS OES data. Bubble size in the left hand graph represents metro area employment in 2016. We estimate the wage effect of employer concentration as described in section 8: we use our coefficient estimates for the effect of the HHI on the wage, by quartile of outward occupational mobility, to calculate a counterfactual wage for each occupation-city labor market *if* the HHI had been 200.

J Appendix: Tables

Table A1: Summary statistics in BGT vacancy data

	p5	p10	p25	p50	p75	p90	p95
Total vacancies posted by employer (by employer)	1	1	1	2	5	19	58
No. of years employer present (by employer)	1	1	1	1	2	3	4
No. of years employer present (vacancy-weight)	2	3	4	4	4	4	4
Occ. share relative to BLS OES (by occ.)	0.11	0.18	0.37	0.81	1.80	4.28	7.22
Occ. share relative to BLS OES (emp.-weight)	0.16	0.21	0.34	0.62	1.14	2.10	2.83
Metro area share relative to BLS OES (by MSA)	0.60	0.67	0.76	0.90	1.12	1.35	1.54
Metro area relative to BLS OES (emp.-weight)	0.60	0.63	0.75	0.90	1.03	1.17	1.34

Note: This table shows some summary statistics from the BGT vacancy data for 2013–2016 inclusive. ‘No. of years employer present’ refers to the number of years in which a given employer posted a vacancy, with a maximum of 4. The vacancy-weighted version of this statistic weights each observation by the number of vacancies an employer posted. ‘Occ. (or metro area) share relative to BLS OES’ refers to the share of each SOC 6-digit occupation (/metro area) in our vacancy data, relative to the share of that SOC 6-digit occupation (metro area) in the BLS OES data for the entire country (calculated for each year 2013–2016 then averaged across the four years). The employment-weighted version of this statistic weights each occupation-metro area cell by employment in that cell in 2016.

Table A2: Distribution of number of jobs on resume and duration of jobs in BGT resume data set.

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

Note: This table shows some summary statistics from the BGT resume data: the distribution of the number of jobs in each resume (across all 16 million resumes in our data set), and the distribution of average job duration in months (across all the jobs reported in our data set).

Table A3: Adjusted R-squared from regressions of occupational transitions on occupational similarity (based on tasks, skills, amenities, or wages)

<i>Dependent variable:</i>	$\pi_{o \rightarrow p}$	
<i>Included characteristic</i>	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

Note: Table shows adjusted R-squared from regressions of the form $\pi_{o \rightarrow p} = \kappa + \alpha_o + \beta \Delta X_{occ\ p-o} + \epsilon_{op}$. Here, $\pi_{o \rightarrow p}$ 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. These analyses are discussed in more detail in Appendix Section F.

Table A4: 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

Note: This table shows the twenty large occupations with the lowest and the highest occupation leave shares - defined as share of workers observed in one occupation in one year but not in the following year, divided by the share that leave their job over that period (see Section 4) - 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 A5: 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

Note: 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 (as in Section 4). Only occupations with at least 500 observations in the BGT data and 2017 OES employment data are shown.

Table A6: Assignment of SOC 2-digit occupation groups into 10 larger occupation categories

Occupation category	2-digit SOC occupation group	SOC code
Managers	Management occupations	11-0000
Managers	Business and financial operations occupations	13-0000
Professionals, except healthcare	Computer and mathematical operations	15-0000
Professionals, except healthcare	Architecture and engineering occupations	17-0000
Professionals, except healthcare	Life, physical, and social science occupations	19-0000
Professionals, except healthcare	Community and social service occupations	21-0000
Professionals, except healthcare	Legal occupations	23-0000
Professionals, except healthcare	Education, training, and library occupations	25-0000
Professionals, except healthcare	Arts, design, entertainment, sports, and media occupations	27-0000
Healthcare occupations	Healthcare practitioners and technical support occupations	29-0000
Healthcare occupations	Healthcare support occupations	31-0000
Food, Cleaning, and Protective Service	Protective service occupations	33-0000
Food, Cleaning, and Protective Service	Food preparation and serving related occupations	35-0000
Food, Cleaning, and Protective Service	Building and grounds cleaning and maintenance occupations	37-0000
Personal Care	Personal care and service occupations	39-0000
Sales	Sales and related occupations	41-0000
Office/Administrative	Office and administrative support occupations	43-0000
Operators / Laborers	Farming, fishing, and forestry occupations	45-0000
Operators / Laborers	Transportation and material moving occupations	53-0000
Production	Construction and extraction occupations	47-0000
Production	Installation, maintenance, and repair occupations	49-0000
Production	Production occupations	51-0000

Note: Table shows our allocation of 2-digit SOC occupational groups into 9 larger occupation categories for the purpose of analyzing whether the relationship between wages and HHI differs across occupation groups. These occupation categories are drawn primarily from Acemoglu and Autor (2011).

Table A7: First-stage regressions: HHI instrument

<i>Dependent variable: log vacancy HHI</i> <i>(segmented by quartile of occ mobility in cols (b)-(e))</i>					
	Full sample	By quartile of occ mobility			
	(a)	Q1 (b)	Q2 (c)	Q3 (d)	Q4 (e)
Log vacancy HHI instrument	0.081 (0.003)				
Log outside-occ. options instrument	-0.712 (0.044)				
Log HHI instrument X Q1 occ mobility		0.078 (0.004)			
Log outside-occ options instrument X Q1 occ mobility		-0.566 (0.049)			
Log HHI instrument X Q2 occ mobility			0.084 (0.003)		
Log outside-occ options instrument X Q2 occ mobility			-0.785 (0.056)		
Log HHI instrument X Q3 occ mobility				0.074 (0.003)	
Log outside-occ options instrument X Q3 occ mobility				-0.742 (0.052)	
Log HHI instrument X Q4 occ mobility					0.085 (0.004)
Log outside-occ options instrument X Q4 occ mobility					-0.818 (0.050)
Vacancy growth	-0.061 (0.044)	-0.461 (0.163)	-0.023 (0.012)	-0.097 (0.062)	-0.588 (0.064)
Predicted vacancy growth	-0.038 (0.100)	-0.473 (0.373)	-0.023 (0.219)	-0.099 (0.237)	0.085 (0.083)
Exposure control	1.424 (0.037)	1.578 (0.054)	1.335 (0.059)	1.464 (0.059)	1.432 (0.049)
Observations	184,411	46,300	46,189	46,080	45,840

Note: In column (a) we run a first-stage regression for our HHI instrument. In columns (b) through (e) we run separate first-stage regressions for our HHI instrument, segmenting our data into four quartiles by outward occupational mobility (the occupation “leave share” as defined in Section 4). 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. All regressions have occupation-year and CBSA-year fixed effects.

Table A8: First-stage regressions: outside-occ. options instrument

<i>Dependent variable: log outside-occ. options</i> <i>(segmented by quartile of occ mobility in cols (b)-(e))</i>					
	Full sample	By quartile of occ mobility			
	(a)	Q1 (b)	Q2 (c)	Q3 (d)	Q4 (e)
Log outside-occ. options instrument	0.783 (0.018)				
Log HHI instrument	-0.000 (0.000)				
Log outside-occ options instrument X Q1 outward mobility		0.689 (0.018)			
Log HHI instrument X Q1 outward mobility		-0.001 (0.000)			
Log outside-occ options instrument X Q2 outward mobility			0.819 (0.019)		
Log HHI instrument X Q2 outward mobility			0.000 (0.000)		
Log outside-occ options instrument X Q3 outward mobility				0.816 (0.021)	
Log HHI instrument X Q3 outward mobility				-0.000 (0.000)	
Log outside-occ options instrument X Q4 outward mobility					0.810 (0.024)
Log HHI instrument X Q4 outward mobility					-0.000 (0.000)
Vacancy growth	0.001 (0.001)	-0.007 (0.005)	0.003 (0.000)	-0.004 (0.001)	0.004 (0.005)
Predicted vacancy growth	0.018 (0.017)	-0.010 (0.041)	-0.021 (0.023)	0.036 (0.032)	0.034 (0.022)
Exposure control	0.001 (0.005)	0.010 (0.010)	0.008 (0.008)	-0.006 (0.009)	-0.004 (0.005)
Observations	184,411	46,300	46,189	46,080	45,840

Note: In column (a) we run a first-stage regression for our outside-occupation option index instrument. In columns (b) through (e) we run separate first-stage regressions for our outside-occupation option index instrument, segmenting our data into four quartiles by outward occupational mobility (the occupation “leave share” as defined in Section 4). 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. All regressions have occupation-year and CBSA-year fixed effects.

Table A9: Regression of wage on HHI and outside-occupation options: robustness

<i>Dependent variable:</i>	Log wage				
	(a) No vac controls	(b) Equal-wt vac	(c) No exposure control	(d) Industry Bartik	(e) Occ-CBSA FEs
Log HHI, instrumented	-0.011 (0.003)	-0.011 (0.003)	-0.011 (0.002)	-0.010 (0.003)	-0.008 (0.003)
Log outside-occ. options, instrumented	0.095 (0.009)	0.095 (0.009)	0.095 (0.009)	0.085 (0.009)	-0.002 (0.007)
Exposure control	0.004 (0.007)	0.004 (0.007)		0.003 (0.008)	-0.001 (0.004)
Vacancy growth		-0.001 (0.001)	-0.001 (0.001)	-0.007 (0.004)	-0.003 (0.003)
Predicted vacancy growth		-0.013 (0.010)	-0.012 (0.010)	0.033 (0.030)	-0.008 (0.008)
Equal-weighted vacancy growth		-0.000 (0.000)			
Industry Bartik				0.146 (0.014)	
Observations	184,411	184,411	184,411	169,341	158,393
F-Stat	386	373	491	620	221
Fixed effects	Occ-year CBSA-year	Occ-year CBSA-year	Occ-year CBSA-year	Occ-year CBSA-year	Occ-CBSA Year

Note: These represent robustness checks for our baseline regression specification (reported in column (d) of Table 3). Each column is a different robustness check. All columns are 2SLS IV regressions estimated using our HHI and outside-occupation option instruments. Column (a) excludes our controls for local actual and predicted vacancy growth. Column (b) includes an additional control for equal-weighted vacancy growth of local firms in the relevant occupation. Column (c) excludes our HHI exposure control. Column (d) includes an industry Bartik shock to control for correlated industry shocks across occupation-metro area cells. All these specifications feature occupation-year and city-year fixed effects. Column (e) runs our baseline regression specification, but with occupation-city and year fixed effects. Other regression info: 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. F-Stat is Kleibergen-Paap Wald F statistic.

Table A10: Regression of wage on HHI and outside-occupation options: robustness (2)

<i>Dependent variable:</i>	Log wage				
	(a) Emp weight	(b) Log emp weight	(c) Drop low rep. occs	(d) Occ rep. weight	(e) CBSA rep. weight
Log HHI	-0.015 (0.005)	-0.011 (0.003)	-0.011 (0.003)	-0.008 (0.004)	-0.009 (0.003)
Log outside-occ. options	0.118 (0.027)	0.103 (0.009)	0.093 (0.009)	0.108 (0.011)	0.094 (0.009)
Vacancy growth	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.000)	-0.001 (0.001)	-0.001 (0.000)
Predicted vacancy growth	-0.018 (0.007)	-0.012 (0.008)	-0.018 (0.011)	0.020 (0.032)	-0.020 (0.009)
Exposure control	0.003 (0.017)	0.007 (0.008)	0.008 (0.009)	0.018 (0.013)	0.003 (0.008)
Observations	184,411	184,258	137,567	184,411	184,411
F-Stat	96	370	411	369	221

Notes: These represent robustness checks for our baseline regression specification (reported in column (d) of Table 3). Each column is a different robustness check. All columns are 2SLS IV regressions estimated using our HHI and outside-occupation option instruments. Columns (a) and (b) weight the regressions by employment and log employment of the occupation-MSA cell, respectively, with employment weights calculated as the average employment in that occupation-by-MSA labor market over 2013–2016 according to BLS OES. Column (c) drops all occupations with average represented-ness in the BGT vacancy data of 0.5 or less. Columns (d) and (e) weight the regressions by average represented-ness of the occupations, and MSAs, in the BGT vacancy data (respectively). Represented-ness by occupation (/CBSA) in the BGT vacancy data is calculated as the share of all vacancies accounted for by a given occupation (/CBSA) in the BGT vacancy data in a given year, divided by the share of employment accounted for by a given occupation (CBSA) in the BLS OES in that same year, averaged over 2013–2016. About one third of occupations in our data have occupation represented-ness of 0.5 or less in the BGT data. Other regression info: 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. All specifications feature occupation-year and city-year fixed effects. F-Stat is Kleibergen-Paap Wald F statistic.

Table A11: Regression of wage on HHI and outside-occupation options, by quartile of outward mobility: robustness checks

<i>Dependent variable:</i>	Log wage				
	(a) No vac controls	(b) Equal-wt vac	(c) No exposure control	(d) Industry Bartik	(e) Occ-CBSA FEs
Log HHI X Q1 outward mobility	-0.026 (0.005)	-0.025 (0.005)	-0.025 (0.004)	-0.027 (0.004)	-0.004 (0.007)
Log HHI X Q2 outward mobility	-0.011 (0.003)	-0.010 (0.003)	-0.010 (0.003)	-0.009 (0.003)	0.002 (0.010)
Log HHI X Q3 outward mobility	-0.007 (0.004)	-0.007 (0.004)	-0.007 (0.003)	-0.004 (0.004)	-0.022 (0.010)
Log HHI X Q4 outward mobility	-0.000 (0.004)	0.000 (0.004)	0.000 (0.003)	0.000 (0.004)	-0.013 (0.015)
Log outside-occ. options X Q1 outward mobility	0.060 (0.011)	0.060 (0.011)	0.060 (0.011)	0.051 (0.012)	0.003 (0.012)
Log outside-occ. options X Q2 outward mobility	0.100 (0.009)	0.100 (0.009)	0.100 (0.009)	0.090 (0.010)	-0.020 (0.017)
Log outside-occ. options X Q3 outward mobility	0.106 (0.010)	0.107 (0.010)	0.107 (0.010)	0.100 (0.010)	0.004 (0.016)
Log outside-occ. options X Q4 outward mobility	0.108 (0.010)	0.108 (0.010)	0.109 (0.010)	0.094 (0.011)	0.013 (0.024)
Exposure control	0.004 (0.008)	0.004 (0.008)		0.003 (0.008)	-0.002 (0.004)
Vacancy growth		-0.001 (0.001)	-0.001 (0.001)	-0.007 (0.004)	-0.003 (0.003)
Predicted vacancy growth		-0.015 (0.010)	-0.015 (0.010)	0.028 (0.030)	-0.009 (0.009)
Equal-weighted vacancy growth		-0.000 (0.000)			
Industry Bartik				0.143 (0.014)	
Observations	184,411	184,411	184,411	169,341	158,393
Fixed effects	Occ-year CBSA-year	Occ-year CBSA-year	Occ-year CBSA-year	Occ-year CBSA-year	Occ-CBSA Year

Note: This table repeats the robustness checks in Table A9, but allowing the coefficient on the HHI and outside-occupation option index to vary by quartile of outward occupational mobility (the occupation “leave share” defined as in Section 4). 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.

Table A12: Regression of wage on HHI and outside-occupation options, by quartile of outward mobility: robustness checks (2)

<i>Dependent variable:</i>	Log wage				
	(a) Emp weight	(b) Log emp weight	(c) Drop low rep. occs	(d) Occ rep. weight	(e) CBSA rep. weight
Log HHI X Q1 outward mobility	-0.044 (0.011)	-0.027 (0.004)	-0.023 (0.005)	-0.011 (0.006)	-0.020 (0.006)
Log HHI X Q2 outward mobility	-0.019 (0.005)	-0.012 (0.003)	-0.007 (0.003)	-0.011 (0.006)	-0.009 (0.004)
Log HHI X Q3 outward mobility	-0.001 (0.005)	-0.007 (0.004)	-0.006 (0.004)	-0.004 (0.006)	-0.008 (0.004)
Log HHI X Q4 outward mobility	-0.005 (0.008)	0.001 (0.004)	-0.005 (0.005)	0.002 (0.006)	0.001 (0.004)
Log outside-occ. options X Q1 outward mobility	0.093 (0.025)	0.067 (0.012)	0.059 (0.011)	0.067 (0.013)	0.069 (0.012)
Log outside-occ. options X Q2 outward mobility	0.135 (0.024)	0.108 (0.010)	0.105 (0.009)	0.118 (0.013)	0.098 (0.010)
Log outside-occ. options X Q3 outward mobility	0.125 (0.032)	0.115 (0.010)	0.105 (0.010)	0.116 (0.012)	0.099 (0.010)
Log outside-occ. options X Q4 outward mobility	0.101 (0.043)	0.117 (0.010)	0.102 (0.011)	0.116 (0.014)	0.105 (0.009)
Vacancy growth	0.000 (0.001)	-0.001 (0.001)	-0.001 (0.000)	-0.001 (0.001)	-0.001 (0.000)
Predicted vacancy growth	-0.020 (0.008)	-0.015 (0.008)	-0.020 (0.011)	0.016 (0.031)	-0.023 (0.009)
Exposure control	0.025 (0.021)	0.006 (0.008)	0.009 (0.009)	0.014 (0.013)	0.002 (0.009)
Observations	184,411	184,258	137,567	184,411	184,411

Notes: This table repeats the robustness checks in Table A10, but allowing the coefficient on the HHI and outside-occupation option index to vary by quartile of outward occupational mobility (the occupation “leave share” defined as in Section 4). 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. All specifications feature occupation-year and city-year fixed effects.

Table A13: Regression of wage on HHI and outside options: heterogeneity by quartile of HHI

<i>Specification:</i>	OLS	2SLS IV	First stage for HHI			
<i>Dependent variable:</i>	Log wage		Log HHI (by quartile of HHI)			
			Q1	Q2	Q3	Q4
Log HHI (instrument)	-0.008	-0.015	0.048			
X Q1 HHI	(0.002)	(0.005)	(0.003)			
Log outside-occ. options (instrument)	0.114	0.102	-0.608			
X Q1 HHI	(0.008)	(0.010)	(0.057)			
Log HHI (instrument)	-0.004	-0.010		0.021		
X Q2 HHI	(0.001)	(0.006)		(0.001)		
Log outside-occ. options (instrument)	0.099	0.087		-0.137		
X Q2 HHI	(0.007)	(0.009)		(0.016)		
Log HHI (instrument)	-0.004	-0.010			0.024	
X Q3 HHI	(0.001)	(0.005)			(0.002)	
Log outside-occ. options (instrument)	0.099	0.088			-0.082	
X Q3 HHI	(0.007)	(0.008)			(0.014)	
Log HHI (instrument)	-0.005	-0.011				0.031
X Q4 HHI	(0.001)	(0.005)				(0.003)
Log outside-occ. options (instrument)	0.104	0.091				-0.177
X Q4 HHI	(0.007)	(0.009)				(0.021)
Vacancy growth		-0.001	-0.891	-0.013	-0.551	-0.086
		(0.001)	(0.190)	(0.008)	(0.050)	(0.056)
Predicted vacancy growth		-0.013	0.455	-0.176	-0.086	0.213
		(0.010)	(0.698)	(0.217)	(0.106)	(0.104)
Exposure control		-0.000	10.547	1.725	0.829	0.785
		(0.007)	(0.838)	(0.096)	(0.039)	(0.022)
Observations	184,411	184,411	45,827	45,884	45,916	45,897

Notes: Column (a) repeats our baseline OLS regression and Column (b) repeats our baseline 2SLS IV regression (corresponding to Table 3 columns (b) and (d) respectively), but allow the coefficients on the HHI and outside-occupation option index to vary based on the average HHI of the occupation-city cell in question. Specifically, occupation-city labor markets are split into four quartiles based on their average HHI over the four year period in question 2013–2016. Quartile boundaries for the average HHI in the occupation-city labor market over 2013–2016 are 276 (p25), 622 (p50), and 1,340 (p75). The remaining columns (c) through (f) show the first stage HHI regressions for the HHI instrument, separately for each HHI quartile. Other info: 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. Regressions have occupation-year and city-year fixed effects.

Table A14: Regression of wage on HHI and outside options: heterogeneity by occupational wage quartile

<i>Specification:</i>	OLS	2SLS IV	First stage for HHI			
<i>Dependent variable:</i>	Log wage		Log HHI (by occ wage quartile)			
			Q1	Q2	Q3	Q4
Log HHI (instrument)	-0.003	-0.009	0.082			
<i>X</i> Q1 occ. wage	(0.002)	(0.004)	(0.004)			
Log outside-occ. options (instrument)	0.081	0.047	-0.357			
<i>X</i> Q1 occ. wage	(0.010)	(0.015)	(0.068)			
Log HHI (instrument)	-0.002	-0.004		0.097		
<i>X</i> Q2 occ. wage	(0.001)	(0.003)		(0.004)		
Log outside-occ. options (instrument)	0.109	0.091		-0.459		
<i>X</i> Q2 occ. wage	(0.008)	(0.012)		(0.048)		
Log HHI (instrument)	-0.004	-0.011			0.074	
<i>X</i> Q3 occ. wage	(0.001)	(0.004)			(0.003)	
Log outside-occ. options (instrument)	0.113	0.083			-0.647	
<i>X</i> Q3 occ. wage	(0.007)	(0.010)			(0.041)	
Log HHI (instrument)	-0.012	-0.014				0.071
<i>X</i> Q4 occ. wage	(0.002)	(0.004)				(0.003)
Log outside-occ. options (instrument)	0.095	0.089				-0.648
<i>X</i> Q4 occ. wage	(0.008)	(0.010)				(0.048)
Vacancy growth		-0.001	-0.350	-0.560	-0.032	-0.511
		(0.001)	(0.077)	(0.063)	(0.018)	(0.196)
Predicted vacancy growth		-0.013	0.175	0.124	-0.392	-0.258
		(0.010)	(0.054)	(0.125)	(0.208)	(0.233)
Exposure control		0.000	1.588	1.515	1.158	1.768
		(0.007)	(0.070)	(0.045)	(0.046)	(0.069)
Observations	184,411	184,411	19,428	40,181	64,770	60,031

Notes: Column (a) repeats our baseline OLS regression and column (b) repeats our baseline 2SLS IV regression (Table 3 columns (b) and (d) respectively), but allow the coefficients on the HHI and outside-occupation option index to vary based on the average wage of the occupation in question. Specifically, each occupation's average wage is calculated in our BLS OES data sample as of 2016, and occupations are split by employment-weighted quartiles of this distribution (such that all observations from one occupation across cities and years are always in the same quartile). The cutoff for the 25th percentile is \$13.01, the median is \$18.49 and the 75th percentile is \$31.74. The remaining columns (c) through (f) show the first stage HHI regressions separately for each occupation wage quartile. Other info: 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. Regressions have occupation-year and city-year fixed effects.

Table A15: Regression of wage on HHI and outside options: heterogeneity by occupation group

	OLS	2SLS IV	First stage (by occupation group)								
Log HHI (instrument)	-0.012	-0.010	0.057								
X Managers	(0.003)	(0.006)	(0.003)								
Log outside-occ. options (instrument)	0.098	0.107	-0.642								
X Managers	(0.010)	(0.013)	(0.067)								
Log HHI (instrument)	-0.012	-0.014	0.071								
X Professionals, excl managers	(0.002)	(0.005)	(0.004)								
Log outside-occ. options (instrument)	0.087	0.077	-0.849								
X Professionals, excl managers	(0.007)	(0.010)	(0.063)								
Log HHI (instrument)	-0.007	-0.015	0.058								
X Healthcare	(0.002)	(0.006)	(0.003)								
Log outside-occ. options (instrument)	0.024	-0.008	-0.207								
X Healthcare	(0.008)	(0.012)	(0.052)								
Log HHI (instrument)	0.006	0.001	0.085								
X Production	(0.003)	(0.004)	(0.004)								
Log outside-occ. options (instrument)	0.103	0.081	-0.817								
X Production	(0.010)	(0.013)	(0.058)								
Log HHI (instrument)	-0.013	-0.015	0.074								
X Sales	(0.003)	(0.008)	(0.006)								
Log outside-occ. options (instrument)	0.119	0.128	-1.069								
X Sales	(0.010)	(0.016)	(0.105)								
Log HHI (instrument)	-0.002	-0.004	0.088								
X Office/Admin	(0.001)	(0.003)	(0.005)								
Log outside-occ. options (instrument)	0.093	0.087	-0.730								
X Office/Admin	(0.007)	(0.011)	(0.085)								
Log HHI (instrument)	-0.003	0.000	0.101								
X Operators/Laborers	(0.002)	(0.005)	(0.007)								
Log outside-occ. options (instrument)	0.063	0.068	-0.162								
X Operators/Laborers	(0.010)	(0.014)	(0.138)								
Log HHI (instrument)	-0.009	-0.027	0.068								
X Food/Cleaning/Protective Service	(0.002)	(0.007)	(0.004)								
Log outside-occ. options (instrument)	0.069	0.022	-0.009								
X Food/Cleaning/Protective Service	(0.009)	(0.015)	(0.065)								
Log HHI (instrument)	-0.004	-0.016	0.091								
X Personal Care	(0.003)	(0.008)	(0.010)								
Log outside-occ. options (instrument)	0.074	0.054	-0.564								
X Personal Care	(0.010)	(0.015)	(0.109)								
Vacancy growth		-0.001	-0.931	-0.020	-0.374	-0.466	0.344	-0.051	-0.772	-0.397	-0.605
		(0.001)	(0.082)	(0.009)	(0.213)	(0.069)	(0.122)	(0.034)	(0.106)	(0.047)	(0.111)
Predicted vacancy growth		-0.009	-0.543	-0.516	0.017	-0.127	1.498	0.247	4.493	0.204	0.158
		(0.010)	(0.182)	(0.474)	(0.091)	(0.356)	(0.762)	(0.319)	(1.341)	(0.050)	(0.548)
Exposure control		-0.002	1.947	1.153	1.533	1.209	1.654	1.392	1.637	1.306	1.458
		(0.007)	(0.100)	(0.077)	(0.058)	(0.062)	(0.106)	(0.080)	(0.102)	(0.067)	(0.087)
Observations	184,411	184,411	27,955	25,917	28,326	26,802	11,924	20,977	9,181	18,373	6,611

Notes: Column (a) repeats our baseline OLS regression and column (b) repeats our baseline 2SLS IV regression (Table 3 columns (b) and (d) respectively), but allow the coefficients on the HHI and outside-occupation option index to vary for different aggregated occupation groups. The remaining columns show the first stage HHI regressions separately for each occupation group. Occupation groups map from SOC 2-digit occupations as defined in Appendix Table A6. Other info: 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. Regressions have occupation-year and city-year fixed effects.

Table A16: Regression of wage on outside-occupation options: 1999–2016

<i>Dependent variable:</i>	Log wage			
	(1)	(2)	(3)	(4)
Panel A: OLS regressions				
oo^{occs}	0.140 (0.010)	0.082 (0.004)	0.095 (0.005)	0.044 (0.006)
<i>Observations</i>	1,944,370	1,944,370	1,944,370	1,944,370
Panel B: 2SLS IV regressions				
oo^{occs} , instrumented	0.122 (0.010)	0.070 (0.005)	0.076 (0.006)	0.030 (0.006)
<i>Observations</i>	1,944,370	1,944,230	1,944,370	1,944,114
Fixed effects	Year	Occ-Year, City	Occ-Year, City-Year	Occ-Year, Occ-City

Notes: This table repeats our baseline regressions with the outside-occupation option index only, over a longer period (1999-2016). 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.

Table A17: Regressions of wage on outside-occupation option index: aggregated occupation codes, with different combinations of fixed effects

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

Notes: This table reports 2SLS IV regressions of the wage on outside-occupation option index with outside options defined at the level of 3-digit or 2-digit occupations (rather than SOC 6-digit). Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses. 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 A18: Regressions of wage on outside-occupation option index, sample split into three periods, with different combinations of fixed effects

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

Notes: This table shows 2SLS IV regressions of wages on outside-occupation options, splitting our sample over three periods. 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 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 A19: Regressions of wage on outside-occ. option index, *employment-weighted*, with different combinations of fixed effects

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

Note: This shows regressions of wages on outside-occupation options, with weighted by the average employment of their occupation-city over the sample period. 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. 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 A20: Regressions of wages on outside-occupation options, incorporating employment share

<i>Dependent var.:</i>	Employment share		Log wage
oo^{occs} , instrumented	-0.226 (0.027)	0.030 (0.006)	0.026 (0.006)
Empl. share			-0.015 (0.001)
Fixed effects	Occ-City, Occ-Year	Occ-City, Occ-Year	Occ-City, Occ-Year
<i>Observations</i>	<i>1,931,901</i>	<i>1,931,901</i>	<i>1,931,901</i>

Notes: 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.

Table A21: Regression of wage on HHI and oo^{occ} : Right-to-work and non right-to-work states

<i>Dependent variable:</i>	Log wage			
	(a) OLS	(b) OLS	(c) 2SLS IV	(d) 2SLS IV
Log HHI	-0.007	-0.004	-0.011	-0.008
X Non right-to-work	(0.001)	(0.001)	(0.003)	(0.003)
Log HHI	-0.014	-0.010	-0.018	-0.015
X right-to-work	(0.001)	(0.001)	(0.003)	(0.003)
Log outside-occ options		0.099		0.086
X Non right-to-work		(0.008)		(0.009)
Log outside-occ options		0.114		0.104
X right-to-work		(0.008)		(0.009)
Vacancy growth			-0.001	-0.001
			(0.001)	(0.001)
Predicted vacancy growth			-0.012	-0.013
			(0.010)	(0.010)
Exposure control			0.011	0.007
			(0.008)	(0.007)
Observations	184,411	184,411	184,411	184,411
F-Stat			377	201

Notes: This table shows our baseline regression specifications as in Table 3 but allowing the coefficients on the HHI and outside-occupation options to differ in right-to-work states and non right-to-work states. States are classified as right-to-work or non-right-to-work according to data from the National Conference of State Legislatures. States which passed right-to-work laws in 2015 or 2016 (Wisconsin and West Virginia) are coded as non-right-to-work in our sample, under the assumption that it would take some time for the passage of right-to-work statutes to affect labor market behavior. Other info: Heteroskedasticity-robust standard errors clustered at the city level shown in parentheses. Units of observation are 6 digit SOC occupation by city by year, for all observations with available data over 2013–2016 inclusive. All specifications include occupation-by-year and city-by-year fixed effects.

Table A22: Counterfactual wage effects of setting HHI to 200,
excluding occupations with a represented-ness<0.5 in our BGT vacancy data

	2,500< HHI <10,000	1,500< HHI <2,500	500< HHI <1,500	200< HHI <500	0< HHI <200
Lowest mobility	7.6% (0.7m)	5.7% (0.9m)	3.6% (3.6m)	1.1% (4.1m)	0 (7.3m)
Q2 mobility	3.0% (0.5m)	2.3% (0.6m)	1.4% (2.8m)	0.5% (3.7m)	0 (10.2m)
Q3 mobility	2.1% (0.4m)	1.6% (0.4m)	1.0% (2.5m)	0.3% (7.0m)	0 (13.8m)
Q4 mobility	0 (0.3m)	0 (0.4m)	0 (1.7m)	0 (1.5m)	0 (1.8m)

Notes: This table repeats the analysis in Table 5, but *dropping* any occupations with an average represented-ness in our BGT vacancy data of less than 0.5 (roughly the bottom third of occupations). This is because our concentration data on these occupations might be a substantial overestimate of the true degree of employer concentration, if our data is disproportionately sampled from large employers for these occupations.

Table A23: Twenty-five occupations with most people affected by employer concentration (based on a predicted occupation-CBSA wage effect of 2% or greater)

Occupation	Represented-ness in occupation in BGT vacancy data
Security guards	.63
Registered nurses	2.1
Nursing assistants	.75
Hairdressers, hairstylists, and cosmetologists	.71
Nonfarm animal caretakers	.69
Fitness trainers and aerobics instructors	.61
Childcare workers	1.3
Licensed practical and licensed vocational nurses	1.4
Radiologic technologists	.69
Pharmacists	1.2
Emergency medical technicians and paramedics	.59
Medical and clinical laboratory technologists	.98
Phlebotomists	1.4
Pharmacy technicians	.91
Medical assistants	.8
Respiratory therapists	.99
Management analysts	1.7
Lawyers	.83
Librarians	.5
Dentists, general	1.1
Software developers, applications	4.6
Bakers	1.1
First-line supervisors of personal service workers	.68
Real estate sales agents	1.9
Massage therapists	.83

Notes: This table lists the degree of represented-ness of each of these twenty-five occupations in the BGT vacancy data. Represented-ness is defined as the occupation's share of vacancy postings in the BGT database relative to the occupation's share of total employment (as per BLS OES). The twenty-five occupations in this table correspond to the occupations with the highest number of people affected by employer concentration, as listed in Table 6. The more underrepresented an occupation is in the BGT vacancy data, the more likely we are overestimating the degree of employer concentration in these occupations and therefore overestimating the effect of concentration. On the other hand, in better-represented occupations we might be more confident that we are accurately prioritizing these occupations.

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.
- 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.
- Azar, José A, Ioana Marinescu, Marshall I Steinbaum, and Bledi Taska**, “Concentration in US Labor Markets: Evidence From Online Vacancy Data,” *Labour Economics*, 2020, 66, 101886.
- Bahar, Dany**, “IVREG2HDFE: Stata module to estimate an Instrumental Variable Linear Regression Model with two High Dimensional Fixed Effects,” *Statistical Software Components S457841*, Boston College Department of Economics, 2014.
- Bajari, Patrick, C Lanier Benkard, and Jonathan Levin**, “Estimating dynamic models of imperfect competition,” *Econometrica*, 2007, 75 (5), 1331–1370.
- Bewley, Truman**, *Why Wages Don’t Fall During a Recession*, Harvard University Press. Cambridge, Massachusetts, and London, England, 1999.
- 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 Oren Danieli**, “Outside Options in the Labor Market,” 2018. Working paper.
- Carnevale, Anthony P., Tamara Jayasundera, and Dmitri Repnikov**, “Understanding Online Job Ads Data: A Technical Report,” 2014.
- Correia, Sergio**, “Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator,” Technical Report 2016. Working Paper.
- Deming, David J**, “The growing importance of social skills in the labor market,” *The Quarterly Journal of Economics*, 2017, 132 (4), 1593–1640.
- Eckert, Fabian, Teresa C. Fort, Peter K. Schott, and Natalie J. Yang**, “Imputing Missing Values in the US Census Bureau’s County Business Patterns,” 2020.
- Gabaix, Xavier and Ralph Koijen**, “Granular Instrumental Variables,” 2020. Working paper.
- Galuscak, Kamil, Mary Keeney, Daphne Nicolitsas, Frank Smets, Pawel Strzelecki, and Matija Vodopivec**, “The determination of wages of newly hired employees: Survey evidence on internal versus external factors,” *Labour Economics*, 2012, 19 (5), 802–812.
- 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.
- Guimaraes, Paulo and Pedro Portugal**, “A Simple Feasible Alternative Procedure to Estimate Models with High-Dimensional Fixed Effects,” *Stata Journal*, 2010, 10 (4).
- Hall, Robert E and Alan B Krueger**, “Evidence on the incidence of wage posting, wage bargaining, and on-the-job search,” *American Economic Journal: Macroeconomics*, 2012,

4 (4), 56–67.

- Hazell, Jonathon and Bledi Taska**, “Downward Rigidity in the Wage for New Hires,” 2019.
- Hershbein, Brad and Lisa B. Kahn**, “Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings,” *American Economic Review*, 2018, 108 (7), 1737–1772.
- Jaeger, Simon, Benjamin Schoefer, Sammy Young, and Josef Zweimueller**, “Wages and the Value of Nonemployment,” *Quarterly Journal of Economics*, forthcoming.
- Jann, Ben**, “COEFPLOT: Stata module to plot regression coefficients and other results,” *Statistical Software Components S457686*, Boston College Department of Economics, 2013.
- Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin**, “Granular Search, Market Structure, and Wages,” 2019. Working paper.
- Johnson, Janna E and Morris M Kleiner**, “Is Occupational Licensing a Barrier to Interstate Migration?,” *American Economic Journal: Economic Policy*, 2020.
- Kambourov, Gueorgui and Iouri Manovskii**, “Rising occupational and industry mobility in the United States: 1968–97,” *International Economic Review*, 2008, 49 (1), 41–79.
- and –, “Occupational specificity of human capital,” *International Economic Review*, 2009, 50 (1), 63–115.
- 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.
- Krueger, Alan B and Orley Ashenfelter**, “Theory and evidence on employer collusion in the franchise sector,” Technical Report, National Bureau of Economic Research 2018.
- Macaluso, Claudia**, “Skill remoteness and post-layoff labor market outcomes,” 2019.
- Manning, Alan**, “Imperfect competition in the labor market,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 973–1041.
- Marinescu, Ioana and Roland Rathelot**, “Mismatch Unemployment and the Geography of Job Search,” *American Economic Journal: Macroeconomics*, 2018.
- McFadden, D**, “Conditional Logit Analysis of Qualitative Choice Behavior,” *Frontiers in Econometrics*, 1974, pp. 105–142.
- Molloy, Raven, Christopher L. Smith, and Abigail Wozniak**, “Internal migration in the United States,” *Journal of Economic Perspectives*, 2011, 25 (3), 173–196.
- 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.
- Sloane, Arthur A**, “Countering resume fraud within and beyond banking: no excuse for not doing more,” *Labor Law Journal*, 1991, 42 (5), 303.
- Stepner, Michael**, “BINSCATTER: Stata module to generate binned scatterplots,” *Statistical Software Components S457709*, Boston College Department of Economics, 2013.
- Sullivan, Paul**, “Empirical Evidence on Occupation and Industry Specific Human Capital,” *Labour Economics*, 2010, 17 (3), 567–580.
- Weil, David**, *The Fissured Workplace* 2014.
- Xu, Ming**, “Understanding the Decline in Occupational Mobility,” *Working Paper*, 2018.